IPv6 Over the Silk Satellite Network

Peter T. Kirstein
University College London
Kirstein@cs.ucl.ac.uk

Wolfgang Fritsche
IABG
fritsche@iabg.de

Hans Frese & Vladimir Lavrov
DESY
Hans.Frese@desy.de
Vladimir Lavrov@desy.de

Abstract

The Silk network is a VSAT-based network connecting a number of NRENs in Central Asia and the Caucasus into the Internet. While most of the network uses IPv4, this paper discusses how it was possible to provide native IPv6 facilities onto the Network. Our initial experiences with that system are described.

1. Background

The Silk Project has been described in several previous publications [1]. It is basically funded mainly by NATO, with the management funded mainly under the SPONGE project by the European Commission, with several other sizable donations from Cisco and DESY. It employs a VSAT system to connect the National Research and Engineering Networks (NRENs) from the eight Newly Independent States (NIS) of the Southern Caucasus and Central Asia, plus Afghanistan, into the Internet. These NRENs will be called in the following *partner sites*. Figure 1 shows with a cross the location of the current remote earth stations.

Figure 1 Silk Earth Station Locations

The hub of the VSAT is in DESY at Hamburg, Germany. The dishes in the remote sites are 2.4 or 3.8m, and the hub station has a 5.6m dish. Each remote station uses the Single Channel Per Carrier (SCPC) technology on the return link; with the current 2W transmitter this can attain 1.5 Mbps. In the West-> East direction, a common DVB forward channel is used able to operate with 40 Mbps of DVB data.

Each site has the identical equipment of Fig. 2.

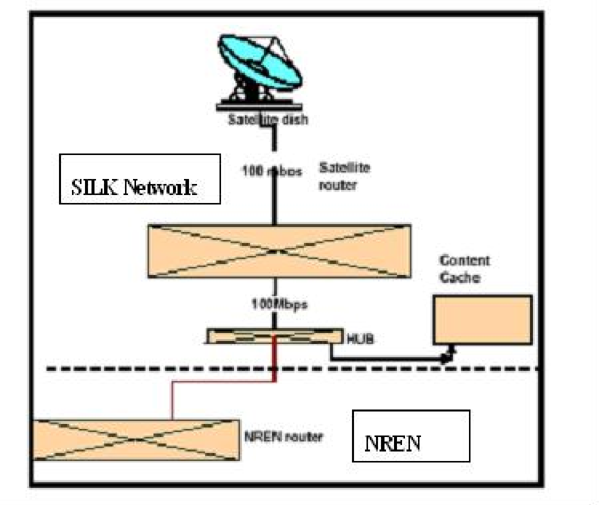

Fig. 2. Schematic of Equipment at each Remote

In Fig. 2 the satellite router, the NREN router and the Content Cache are all from Cisco. At DESY the NREN router is attached to their local network and hence, though the DFN, to GEANT. In the remote sites they are attached directly to the NRENs. Only the equipment above the dotted line counts as the Silk Network. Until recently, the network has been runonly in IPv4 mode.

The NRENs in the partner sites of Fig. 1 have expressed interest in finding out more about IPv6 – provided it does not interfere with their IPv4 service. With this in mind, two initiatives are being undertaken. First, we are running the Cisco routers dual stack; at the bandwidth we are using, there are no performance problems. Second, we have been experimenting with new equipment developed by GCS and being provided by ESA/IABG to run directly IPv6/DVB on a second, separate carrier. Without the GCS equipment, it is necessary to run the IPv6 in tunnel mode; this is because the transmitter cards in the DESY hub automatically put in IPv4/DVB formatting. Both of these mechanisms are described in this paper.

2. Normal Dual Stack Operation

Most of the operation of the Silk Network is dual stack. This link is the principal lifeline for their academic communities, so that the operators of the remote NRENs cannot risk any speculative activities interfering with their service traffic. Moreover, many of the applications being used have not been ported to IPv6, so that there is no question of the network running only IPv6.

In view of the above, it was necessary to change the routers of Fig.2 to run dual stack. Then it became possible to run IPv6 applications between IPv6 hosts in the partner sites and other IPv6 Hosts in the West. All such communications must, of course, pass through the DESY hub. DESY has configured some IPv6 hosts, and also has IPv6 connectivity via the German NREN (DFN) and the European academic backbone (GEANT).

Eventually the Partner sites expect to put up some IPv6 facilities on their NRENs. So far, however, they have had only very limited experience with IPv6, and only one host which sits near the partner earth stations on LANs connected directly to the routers of Fig. 2. Eventually they expect to experiment also with various transition aids like NAT-PT. This is, however, very early days in their experience.

When using the normal Silk equipment, they run IPv6/IPv4 tunnels between the Cisco routers in the partner sites and those at DESY. While this puts in some additional overhead, it is not very significant – compared to other overhead cause by using the special equipment discussed below.

3. Running Native IPv6 over the Silk Network

In order to understand the operation of running native IPv6 over the Silk network, it is necessary to discuss the operation of the network in greater detail.

The network uses Single Channel per Carrier (SCPC) on the return channel between the partner sites and DESY. Under this mechanism, any data format can be carried. Thus it is possible to send either IPv4 or IPv6 without any problem. The routers at DESY can understand the version of IP being used by the transmitters in the partner sites, because they are running dual stack. While this mechanism is reasonably efficient, it has one big drawback. Each partner earth-station must be allocated a particular frequency range, so that it can send its data. It is possibly to change the frequency allocation, but this is not usually a dynamic process. If a partner site is not using its allocation, this bandwidth is wasted.

In the reverse direction, we use a mechanism called DVB [2]. In this there is an extra wrapping of the data over DVB frames. This wrapping indicates for which earth station the data is intended. While DVB does incur an extra overhead, of around 16%, it has two advantages. First it allows data to be allocated dynamically to the different remote sites depending on the traffic demands; this could be achieved by other level-2 mechanisms, but its use for TV broadcasting has allowed comparatively cheap hardware to be available. Second, it is particularly well-suited to the VSAT architecture, because of the economics of the usage.

In a satellite system, the transmitters are usually much more costly, for a given functionality, than the receivers. For this reason, one attempts to have simpler transmitters at the remote sites than at the hub. In a VSAT architecture, there is a powerful transmitter at only one site, the hub. For this reason, the extra complexity of the DVB transmitter at the hub is incurred only the once; the DVB receivers are low cost, and can be afforded at each earth station.

For the normal IPv4 operation, a schematic of the equipment used at the hub is shown in Fig. 3.

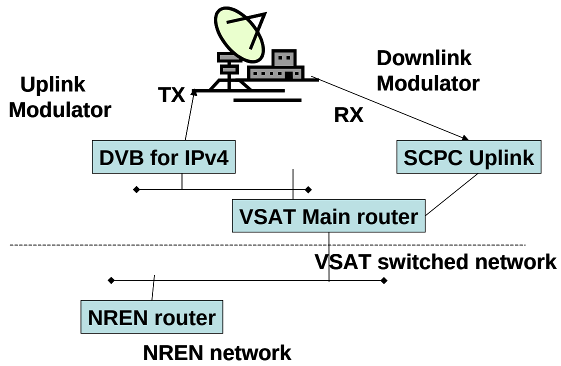

Figure 3. Schematic of Equipment used for IPv4 at the hub site

The partner site equivalent is shown in Fig. 4.

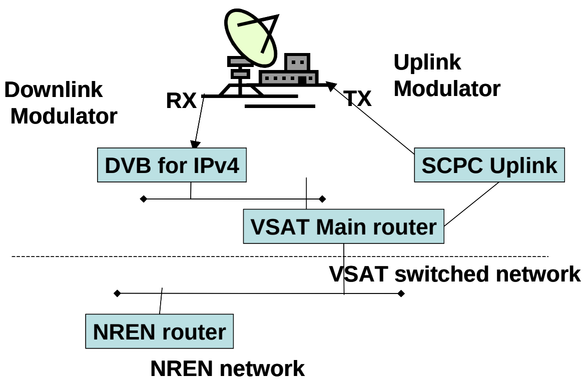

Figure 4. Schematic of Equipment used for IPv4 at the partner sites

The DVB encapsulator in Fig. 3 has to do the DVB encapsulation. The DVB decapsulator of Fig. 4 has both to strip off the DVB encapsulation, and forward the packet to the LAN.

4. Extending the Silk Network to Native IPv6

Unfortunately, the Harmonic DVB cards used by the Silk network in the hub are capable of operating only with IPv4 packets. Moreover, the most recent encapsulation encoding for IPv6 defined in the IETF is significantly different from that used on those cards. GCS, under a contract with the European Space Agency (ESA), has developed cards that do the more modern ULE[3] encapsulation for IPv6/DVB. Moreover, their cards have been designed to operate with both IPv4 and IPv6 packet formats. For this reason the Silk Project agreed to experiment with these new cards.
ESA/IABG agreed to upgrade the configurations of Figs 3 and 4 to those of Figs 5 and 6.

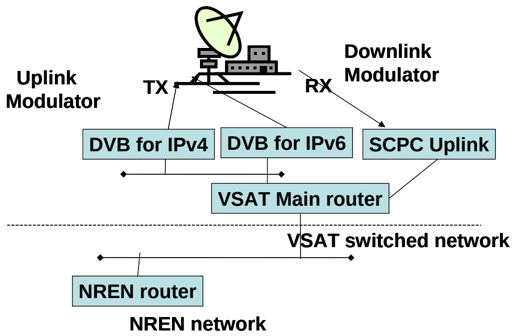

Figure 5. Schematic of Equipment used for IPv4/6 at the hub site

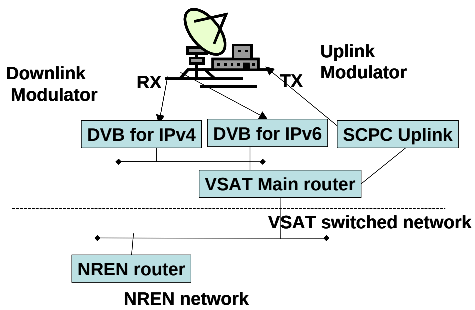

Figure 6. Schematic of Equipment used for IPv4/6 at the partner site sites

This extra equipment caters for the absolute requirement that we did not endanger the IPv4 operations; it allows the IPv6 experiments to be carried out without any change to the IPv4 services.
At the hub, there is an extra DVB gateway with routing functionality, which contains a card to do IPv6/DVB encapsulation based on the new ULE specification. It is mainly an IP router (IPv4 and IPv6) and as such plays an active role in the IP world (e.g. routing, filtering). It also has some satellite link management capabilities. Any native IPv6 traffic to the Eastern sites is routed through this new DVB gateway from GCS rather than the normal one from Harmonic. On the receiving side there is another PC router, with an extra card doing IPv6/DVB decoding. In the East->West direction, the normal SCPC channel can be run directly with the dual-stack traffic. It is necessary to ensure, of course, that the appropriate traffic goes through the different DVB modules.. ESA/IABG have provided a DVB gateway and 6 DVB receiver cards. The gateway and one of the receivers is sited in Hamburg (which has a test rig like those at the partner sites); the other five cards are sited in the partner sites of Almaty, Ashgabat, Baku, Tashkent and Tblisi. This has been done under the auspices of the contract IABG has from ESA.
The partner sites have provided a PC to house the new cards, and made available some PCs to run IPv6 services.

5. Bandwidth for Native IPv6 Over Silk

In order to maintain the normal IPv4 operations, ESA and Eurasiasat provided extra satellite bandwidth for testing the new equipment for three months. This bandwidth is used in two ways. On the West-> East DVB path, a separate DVB channel has been set up for IPv6 operation. On the East -> West SCPC path, the packets can run dual mode; here we have just made a small addition to their normal bandwidth allocation.

6. Applications Testing with IPv6

We have had only limited testing in the native IPv6 environment. This is mainly because of lack of time, and the need to train people at the partner sites to configure and use the new facilities. We have deliberately used some demanding applications – namely the Mbone tools [4] VIC/RAT for audio-video conferencing. Both VIC and RAT have long been IPv6-enabled, mainly under the 6NET project.
VIC and RAT are also multicast-enabled, and we have run these extensively in that mode under the 6NET project. However, we did not consider it worthwhile to install also a multicast infrastructure in the partner sites at this time. Instead we ran a Reflector at DESY, which acted as a Multiplexing Control Unit (MCU), and ran the audio and video streams as Unicast ones. With this configuration, we have run successful conferences with a substantial number of sites operating simultaneously.
The Silk network is not generously provisioned with bandwidth. It runs typically at -700 Kbps on the SCPC paths, anda shared 20 Mbps on the DVB one. In order to provide adequate quality for the conferencing, we implemented Quality of Service (QoS) on the Cisco

routers. This QoS is based on the IP address of the Sources, and gives priority to the streaming traffic. Clearly this is not a scalable situation, but it verified that satisfactory performance was achievable.

7. Requirements for IPv6/DVB Boards

While the above experiments were successful, they revealed a severe economic problem with these boards. We mentioned already that DVB gave a 16% overhead on the normal IPv4 operation. However, there we used fairly complex coding (8PSK) of the analogue frequency in the hub transmitter. As a result we achieved a raw data rate of 1.92 bits/Hz, and even after allowing for the DVB overhead,a useful data rate of 1.66 bits/Hz.

The cards used by IABG were designed to prove the concept of new ULE encapsulation for IPv6/DVB, not as a commercial product. As the development of the ULE functionality has been done on a Linux platform, and as there are no DVB cards available for Linux with 8PSK support, a DVB card with QPSK coding had to be used. Hence we could attain only 0.92 bits/Hz. Even with IPv6/IPv4 tunnels over our standard equipment, we are able to do considerably better than the data rates with the IPv6 boards. Moreover, the use of a different DVB channel for IPv6 is not very efficient; when we do not have IPv6 traffic, that DVB bandwidth is wasted. All these problems would have been overcome if the GCS board had support for 8PSK modulation, or the Harmonics board had support for dual-mode operation.

The ULE encapsulation adopted by the new boards has now been ratified as a standard. Thus we may expect more suppliers to provide such boards; hopefully some will include the above missing features.

8. Our Future Plans and Native IPv6

The European Commission (EC), under the aegis of the 6NET project, has agreed to provide some limited bandwidth to help in the dissemination of information on the use of IPv6 to the Silk community. They have not stipulated whether we use native or tunnelled IPv6 – which does not really impinge on the users or even partner NREN operators. For this reason, in most of our future work over the next few months, we expect to continue with tunnelled variant.

With the current boards we expect to continue for a short time to compare the overheads introduced by the different encapsulation schemes versus the gains from modulation and Forward Error Correction. All IPv6 service transmission over the satellite path will be done, however, with our normal boards and tunnelled IPv6/IPv4.

As part of the dissemination activity, we have already provided an IPv6 workshop for the Silk community. This has clearly motivated the partner countries to learn more about the technology. We now plan to provide two further sets of services in the IPv6 environment. First, we will introduce the media services based on IPv6. Here we will include the three sets of applications developed under 6NET: H.323, VoIP and VIC/RAT. How large a deployment will be made depends on the demands of the partner sites; we expect it will be modest. Second, we will provide some transition mechanisms, and will ensure that they use the IPv6 tunnelled path. In the process, we expect to monitor both the relative performance of the IPv4 versus the IPv6 services, and to develop mechanisms for ensuring that certain services use IPv4 and others IPv6. In the process of this activity, we will investigate whether the finer granularity of QoS signalling in IPv6 may be used with advantage in this context.

9. Conclusions

We have shown that it is possible to extend an IPv4/DVB satellite service to dual-stack operation fairly simply. This extension can be done gradually, with minimal impact on those users not wishing to take advantage of IPv6 functionality; the effort required of those who do desire to acquire IPv6 experience is minimal. While some deficiencies were found in the simple hardware used, these can be overcome in a straightforward manner. In the process we have extended the transmission environment and the geographic communities that have acquired experience of IPv6 and dual-mode operation.

10. Acknowledgements

We wish to acknowledge gratefully the support ofthe following: NATO under the Silk Project; the EC under the 6NETand SPONGE projects; ESA/IABG/GCS under their “Preparation for Introduction of IPv6 in Satellite Communications” and “Standardisation Support of Enhanced IETF IP Encapsulation Techniques” projects; Cisco for their support under their donationprogramme; DESY for hosting the Silk hub; Eurasiasat under their testing programme; and ISOC under their training programme. It is the fact that all these agencies agreed to support this activity, that has made it possible to bring IPv6 to this emerging Internet community. Finally, it would not have been possible without the dedicated activity of the technicians in the Silk Partner sites who participated in this work.

11. References

[1] www.silkproject.org/.

[2] ETSI EN 301 192 v1.3.1: “Digital Video Broadcast (DVB); DVB specification for data broadcasting”, ETSI Norm, January 2003.

[3] G. Fairhurst, B. Collini-Nocker: draft-fair-ipdvb-ule-02.txt “Ultra Lightweight Encapsulation (ULE) for transmission of IP datagrams over MPEG-2/DVB networks”, October 2004.

[4] www-**mice**.cs.**ucl**.ac.uk/multimedia/software/.