CITIBANK'S GLOBAL DATA NETWORK

Malcolm Hamer
Regional Telecomms Head, Asia Pacific, Citibank N.A.

Introduction

Citibank's global data network was initially put in place in the years 1979-1983 to support electronic banking for Citibank's corporate clients - enabling them to obtain up-to-the-minute account information on all their accounts with Citibank around the world, and to electronically initiate a variety of transactions. Although electronic banking is still one of the primary applications of the network, the period 1983 to the present time saw a rapid growth in internal Citibank traffic, supporting the globalization of banking products and activities by more closely coupling activities within the bank across 75 countries.

Citibank's approach to technology has been based on a philosophy of decentralized development and rapid response to new market opportunities. As a result there is a wide variety of hardware and software platforms in use around the world. This contrasts with other large corporations who have stuck to a single hardware vendor and highly centralized software development. Citibank's approach has meant that the global network has had to be highly flexible with regard to host computer interfaces, terminal interfaces, and communications protocols. This has led to the evolution of a multi-layered network using simple and proven techniques wherever possible.

This paper describes some of the lessons learned in building this network over the last 15 years.

A brief history

The original form of the network was a point-to-point network based on Infotron statistical multiplexers. Only one electronic banking service, for US dollar accounts, was offered to customers around the world. Customers would use dial-up data connections via the local telephone network to access the nearest Citibank branch, and the central-site modems there would connect them directly to a channel on the statmux that took them to the host computer in New York. As further services were added it became necessary to have a different group of telephone lines and central-site modems in each country for each service. This soon became unmanageable and so the statmuxes were replaced with switching statmuxes from CASE. These made it possible for a customer to get into the network and select any service by means of a short code (e.g. :44 for electronic banking in London, :41 for Zurich, :65 for Singapore, etc).

Thus, a customer with accounts in several countries could establish a single network session and then, within that session, set up a series of host sessions and get a snapshot of his position with Citibank worldwide in a few minutes. This process could be automated on PCs, which were starting to be used by many customers.

In some respects this approach was crude, because it pushed the task of integrating information out to the customer's PC, instead of requiring that Citibank provide some sort of global database. However, it enabled electronic banking to be available to customers very early on, without the need for major development efforts on mainframe software.

Since 1977, Citibank had a completely separate message switched network, using telegraph-speed circuits (50bps or 75bps in most cases) for handling international customer transactions. Once the data network (or GTN as it came known - Global Telecommunications Network) became established, the branch-to-switch and switch-to-switch channels of the message network were moved onto the GTN. This greatly reduced the cost of the message network, since the bandwidth unit costs in the GTN were a lot lower than the rentals of the exclusive telegraph circuits.

About the same time (1982) an electronic mail system was connected to the network on an experimental basis. Staff needing to communicate across borders and timezones quickly seized on this as an effective tool for managing global customer relationships, global projects, information sharing, and more rapid execution of administrative actions. Within three years almost all staff who handled international business, including all senior management, were using electronic mail regularly. Today, electronic mail represents about 40% of all session-based communication on the network and about 30,000 employees are on the system.

All the communication described so far was full-duplex, asynchronous, with host echo of typed characters. This was initially by necessity. However, it was soon recognized that this method of communication, although it had some drawbacks (such as delayed echoing of typed characters onto the user's screen when going over satellite circuits, and inefficient use of bandwidth in the terminal-to-host direction), it provided a great deal of flexibility in terms of vendor independence. A "dumb" terminal, or PC running terminal emulation software, could be used to communicate equally well with a DEC mainframe or an IBM mainframe, since asynchronous communication represented a "lowest common denominator" of protocols.

In the period 1984-1989 a number of requirements emerged for point-to-point synchronous channels to link mainframes to one another. The CASE statmuxes could not initially support synchronous connections, and even when synchronous facilities were introduced they did not perform well. What was done, therefore, was to build a new network "layer" underneath the statmuxes, using GDC time-division multiplexers (TDMs). These broke down the leased circuit bandwidths into useful chunks that could be applied to synchronous communication, voice communication (digital), or fed into the statmux layer.

In 1989 the capacity limitations of the CASE statmuxes were reached as network usage grew, and it became clear that the asynchronous host interfaces (over 100 of them on some hosts) were difficult to manage in terms of fault isolation. It was therefore decided to replace the statmux layer with packet switches from Telenet (subsequently acquired by Sprint, and more recently partnered with Alcatel). This process took three years to complete, starting with the first installations in 1990. The packet switches made it possible to convert to X.25 host interfaces, which had many advantages, and to greatly expand the capacity of the network.

In parallel with this the increased use of DEC VAX computers to support trading room operations and electronic mail created a DECnet "layer" in the network. This used the packet layer and the TDM layer to provide raw channels to link the VAX hosts. Also, a number of international SNA networks were rolled out, also taking raw bandwidth from the TDM layer.

Starting in 1993 the need for a network layer to interconnect LANs became apparent. In order to address this the implementation of a router network layer was started, using Cisco routers. This rides on top of the TDM layer and functions in parallel with the packet layer. The new LAN/LAN applications that have been put on the "router backbone" have required large amounts of bandwidth, so many leased circuits are in the process of being upspeeded.

What we have today is thus a multi-layered network. At the lowest layer is the TDM network, breaking out raw bandwidth (generally from 64kbps to 2Mbps) into useful chunks. Sitting on top of this, in parallel with one another, are the packet layer (used mainly for terminal-to-host sessions, and a certain amount of X.25 host-to-host communication), and a router network layer (used mainly for LAN/LAN and LAN/host communication). And on top of this there is a global DECnet layer and some regional SNA networks. With recent changes, the DECnet layer now utilizes the packet layer, the TDM layer, and the router network layer to provide its inter-area links. The TDM layer is also used to provide digital voice channels between PABXs. This is done on an "opportunistic" basis, using spare bandwidth that becomes available as leased circuit bandwidths are increased in standard steps, so the amount of voice traffic on a route increases or decreases depending on data usage: meeting demand for data communication is always the main purpose of the network.

Lessons learned

Nothing that was done at any stage in the evolution of the global network was very adventurous by industry standards, except perhaps in terms of scale. (The geographical coverage of 75 countries surpasses that of many of multinational corporations, and even some of the international data networking companies and consortia.) We were deliberately conservative, sticking to proven technologies that could be depended on in variable field conditions (poor air conditioning, unstable power, noisy leased circuits, etc). What we learned was that even our deliberate conservatism was sometimes optimistic.

These are examples of some of the lessons we learned, not a complete inventory:

- *Beware of the software components of tried-and-tested products when the product is used in large-scale situations*. Because of the geographical scope of the Citibank network, and the high traffic levels in certain parts of the network, we are constantly finding product bugs that other customers have not found yet. We have found weaknesses GDC's routing table updating software that cause it to be unable to cope with several circuit failures at once, weaknesses in Sprint's table-generation software which cause it to produce fatal errors when there are more than 500 mnemonic addresses stored per node, weaknesses in Sprint's X.25 layer-2 software that cause heavy load conditions to get trapped in an error loop, and weaknesses in Cisco's router software that cause it to crash under heavy loads - to name but a few deadly bugs.

- *Do not rely on IBM's asynchronous communication support*. In 1982 we needed to include IBM hosts in the general scheme of using asynchronous communication as the lowest common denominator. We quickly found that IBM did not understand async. What we did was to use protocol convertors to adapt IBM's screen-based, synchronous communications to the async world. This proved very successful because it was a stable solution, standing independent of complex IBM communication software, which could support any one of about 15 common "dumb" terminals.

- *Make the best of a product's strengths and work around its weak points*. We spent 18 months trying to get our SNA network layer to run through our packet layer using Sprint's SDLC-support software. In the end we gave up and resorted to plugging raw bandwidth from the TDM layer into the SNA components. Sprint know X.25 but they are useless when it comes to IBM protocols. We have lost the hoped-for optimization of bandwidth between SNA and X.25 traffic, but at least the SNA network is stable. To give an example the other way round: we found that the DECnet layer was unreliable when we ran it through the TDM layer because it was easily disrupted when a leased circuit failed and the TDMs started re-routing channels. We moved some of the DECnet channels onto "DLMs" (X.25 connections via the packet network) and performance improved, because DECnet expects DLMs to drop out from time to time and will immediately initiate a re-connect and settle down quickly.

- *Standards aren't standard*. Even though there has been much talk about the benefits of international standards, there are very few products which can be plugged into other products without quite a lot of fine tuning. When we started converting async network/host interfaces to X.25 interfaces we found that it takes about three months per host to work out all the problems of incompatibility between the packet switch X.25 software and the host X.25 software. In some cases it took as long as six months for a particular host.

Conclusion

Citibank's global data network, which covers 75 countries, consists of a number of layers that meet different needs: TDM (synchronous channels and digital voice), packet (terminal-to-host and some host-to-host communication), router (LAN/LAN and LAN/host communication), DECnet, and SNA. It has evolved over 15 years and we have deliberately avoided state-of-the-art technology, depending instead on tried-and-tested components that will perform under tough field conditions. The scale of the network has exposed flaws in even well-tested products. Overall, our experiences have moved us in the direction of greater caution and extreme skepticism about manufacturers' claims about new products and technologies which may work in the USA where bandwidth is cheap, reliable, and terrestrial, but are troublesome in the field of global networking.

Presented at:
AIC Conference
Hong Kong, February 3, 1994