A Perspective on how ATM Lost Control

Simon Crosby*, Sean Rooney†, Rebecca Isaacs‡, Herbert Bos§

Abstract

Contrary to the initial high expectations, ATM failed to become the universal network technology covering all services and running from the desktop to the backbone. This paper tries to identify the technological problems that contributed to this failure.

1 Introduction

From a service provider perspective, the Internet is a nightmare. Although enterprises are reaping tremendous benefit from lower communications costs and new IP-based applications, the anticipated profits to service providers from the Internet economy have not materialised, and record numbers of carriers and their suppliers are going out of business. The roots of the problem lie in the architectural principles of IP which delivers just one service: unreliable, insecure, best effort connectivity. Although tremendous benefit is gained from IP at the edge of the network due to its inherent simplicity and because it allows the application provider to be entirely freed from network constraints, there is an assumption that the providers will continue to increase capacity so that IP based applications will deliver high quality. Thus in the IP model the provider has to deliver constantly growing amounts of bandwidth, but is totally excluded from the value chain. Effectively the provider takes the role of a bit pipe, with no way to link their revenues to the value their customers derive from the network[1]. Furthermore, the best effort service offered by IP is inadequate for enterprise customers who have relied for years on the dependable alternatives of leased lines, Frame Relay and ATM.

Notwithstanding the ubiquity of IP at the edges of the network, ATM still has very healthy growth numbers. All DSL traffic crosses ATM, it comprises the core of most IP networks, ATM VPNs are commonplace and there is even a considerable market for circuit emulation services, including voice over ATM. The lack of any kind of service guarantees together with the difficulty for service providers to charge for value added to an IP network, suggests that ATM is a missed opportunity. So what went wrong?

2 The ATM value proposition

Borrowing from earlier control planes, specifically the narrowband Q.931 signalling standard, the ATM standardization community developed a User-Network Interface (UNI), and subsequent Network-Network Interface (NNI), that would allow users to dynamically signal for connectivity with the appropriate bandwidth and QoS guarantees that applications required. The notion of a User-Network signalling interface was seen as a crucial part of the value proposition for ATM.

From the point of view of the service provider, IP is flawed because there is no way to link revenues to the value customers derive from the network. In this respect, ATM had the potential to gain advantage over IP, provided that a suitable UNI could be developed. However, the UNI standardised by the ATM Forum was defective in at least three ways. Firstly, it was flawed because the UNI was tied to a specific

*simon@cplane.com, CPLANE Inc, USA
†sro@zurich.ibm.com, IBM Research, Switzerland
‡risaacs@microsoft.com, Microsoft Research, UK
§herbertb@liacs.nl, LIACS, The Netherlands

[1]Data has grown to roughly 60% of the traffic volume in some carriers but accounts for a tiny fraction of the profits so providers are being forced to put more assets in place to support their lowest margin business.

technology (ATM) requiring the signalling layer to be implemented ‘on the box’. This means that every access interface, in particular every operating system, edge router and device that might be attached to the network in the future, must implement a complex, heavyweight signalling stack. It doesn’t take much insight to see that to require that every OS implement the signalling stack was an act of insanity which doomed ATM to failure. Secondly, the model provided no accountability for signalling. That is, when the user exercises the UNI, and signals for an ATM connection, there is no hook into the billing system that allows the provider to account for and therefore bill for the value delivered over the interface. Finally, the ATM UNI was extraordinarily heavyweight and complex. Because it tried to define all possible service models for all possible applications, present and future, it became bulky and difficult to understand. At the same time, its closed nature made it impossible to provide new services for applications with specific requirements that were not catered to by the standards.

Thus the ATM UNI signalling protocol was crippled from the outset, imposing impossible burdens on both service providers and the users of the network. Perhaps it is time to re-evaluate the need for signalling across the User-Network interface, and this time to get it right. The key requirements are clear: separate the UNI from the underlying network delivery technology, make it easy for application users at the edge of the network, and allow providers to bill for services delivered using it.

Let’s step back for a moment to Signalling System Number 7 (SS7), a network of signalling channels used in the Public Switch Telecommunication Network. Why was it valuable? The answer is remarkably simple: It separated the signalling from the equipment. The SS7 Service Control Point allowed the provider to implement value propositions (albeit limited ones) independent of the underlying switched network. The provider owned the application (eg 800 number or ‘follow me’) and on the basis of performing some function in response to a user interaction via the ‘UNI’, the provider could charge, and make money. Although designed on this model, ATM control is too complex and presents the wrong level of abstraction to application developers and users. Small wonder then that almost all ATM connections are set up using the hopelessly impoverished pseudo signalling protocol SNMP: Providers ‘provision’ connections (usually sPVCs or PVCs), turning the dynamic nature of ATM into a sham, and its capabilities simply into a flexible virtual topology on top of fixed TDM capacity. This allows providers to traffic engineer for IP, deliver basic DSL ‘pipes’ and some voice traffic, with QoS, but the oft-touted B-ISDN dynamic service interface has never been deployed.

3 Open Signalling

SS7 on the PSTN was successful because it offered service providers a way to add value to the network and to charge accordingly, but the system is closed to third parties, inhibiting the development and deployment of new and innovative services. The need for an open network was not recognised by the ATM Forum, who, by not specifying the interface between the control plane and the physical network, tied the control software to the physical devices. Although architecturally distinct, cell forwarding and connection establishment functions tend to be integrated, and this results in switch-specific control systems which can only be maintained and evolved by vendors.

In academia such concerns led to the establishment of the OPENSIG forum to develop a more flexible control plane for ATM. Open signalling systems require that there be a clear divide between the control plane and the data plane, with an open, switch-independent control API. Open signalling ideas were heavily influenced by the industrially developed IP Switching[4] which discarded entirely the standard ATM control plane. An IP Router built over an ATM switch made a local decision as to whether the packets of a flow should be forwarded at the software IP layer or, if the flow was sufficiently long-lived, to cut through onto the ATM hardware layer. IP Switching was ultimately not successful because hardware IP routers became commonplace, but its General Switch Management Protocol for communication between the IP control layer and the ATM forwarding layer is an example of a switch-independent control API which has since been standardized by the IETF[3].

The large number of proposals to replace the ATMF control plane seriously questions the basic assumption of the standardization bodies that there is a single control interface which will provide for the needs of all applications and services. In response to this, the open signalling community addressed the need to support a plurality of ATM control systems. Work done at Cambridge, using virtual resource partitioning, showed how this could be achieved[7] and was extended to support novel control planes[5, 1] and a virtual network service[8, 6, 2]. These academic efforts have in turn been adopted by industry through the Multiservice Switching Forum and the IEEE P1520 proposed standard for network APIs.

4 Lessons learnt with ATM

IP is the basic unifier of the enterprise application suite, and is the primary vehicle for applications of value. This indicates that the control procedures for interacting with the network should occur at a much higher layer, specifically the application layer, to allow applications to directly interact with the network to dynamically request the security, bandwidth QoS and other attributes that are required.

Pipe based services (whether ATM or IP based) relegate the service provider to the role of simple bit shifter, but the value proposition has moved outside the network. The right way to solve the control problem is to ensure that it is *out of band* so it allows customers to dynamically request network services to meet application specific needs (bandwidth, QoS, etc). It should also be technology independent, so that the *service* layer concepts of value can be delivered independent of the underlying delivery technology (ATM, GigE, MPLS, IP or whatever comes next). This motivates for development of 'application layer signalling' which enables direct interaction between application components and the network, without requiring detailed knowledge of the network infrastructure, and encompassing both servers and networks.

Why is application layer signalling important? Primarily because the only successful service provider models are those that allow customer-network interaction. This is possible in the PSTN, but it is impossible in the 'best effort pipes' based IP network. Providers need to be able to interact with their customers, to implement services for which their customers are prepared to pay a premium. If they are unable to do so, they will become just utility bits-per-second providers.

5 Conclusion

Ultimately open signalling failed; ATM's control plane has been abandoned and ATM is used as one of many layer-2 protocols in the public internet. Retrospectively it can be seen that simplicity of deployment was the killer argument for IP. However, ease of deployment does not equate with simplicity of management and IP's lack of precise control over the data path makes it hard for network operators and service providers to differentiate their offerings in a meaningful way. ATM's relegation as nothing more than a flexible way to compensate for the most blatant failings of IP, as far as reliability and predictability of service are concerned, is a result of over-complex protocols without an adequate understanding of the business driver.

6 Acknowledgments

Thanks are due to Kobus van der Merwe for his valuable insights and discussions.

References

[1] Bos, H. Open extensible network control. *Journal of Network and Systems Management 8*, 1 (Mar. 2000).

[2] Isaacs, R., and Leslie, I. Support for resource-assured and dynamic virtual private networks. *IEEE Journal on Selected Areas in Communications 19*, 3 (Mar. 2001), 460–472. Special Issue on Active and Programmable Networks.

[3] Newman, P., Edwards, W., Hinden, R., Hoffman, E., Ching, F., Lyon, T., and Minshall, G. Ipsilon's General Switch Management Protocol specification version 2.0. RFC 2297, Mar. 1998.

[4] Newman, P., Minshall, G., and Lyon, T. IP switching: ATM under IP. *IEEE/ACM Transactions on Networking 6*, 2 (Apr. 1998), 117–129.

[5] Rooney, S. The Hollowman: An innovative ATM control architecture. In *Integrated Network Management V* (San Diego, USA, May 1997), A. Lazar, R. Saracco, and R. Stadler, Eds., IFIP & IEEE, Chapman & Hall, pp. 369–380.

[6] Rooney, S., van der Merwe, J. E., Crosby, S. A., and Leslie, I. M. The Tempest: A framework for safe, resource-assured programmable networks. *IEEE Communications Magazine 36*, 10 (Oct. 1998), 42–53.

[7] van der Merwe, J. E., and Leslie, I. Switchlets and dynamic virtual ATM networks. In *Integrated Network Management V* (San Diego, USA, May 1997), A. Lazar, R. Saracco, and R. Stadler, Eds., IFIP & IEEE, Chapman & Hall, pp. 355–368.

[8] van der Merwe, J. E., Rooney, S., Leslie, I., and Crosby, S. The Tempest—a practical framework for network programmability. *IEEE Network Magazine 12*, 3 (May/June 1998), 20–28.