Always On/Dynamic ISDN
Request For Comments: 001
Category: Informational

A. Kuzma
Intel Corporation
October 1996

Always On/Dynamic ISDN

Status of this memo

This memo provides information for members of the Vendors ISDN Association. This memo does not specify a standard of any kind. Distribution of this memo is unlimited.

Abstract

This document provides motivations for a data network access service termed “Always On/Dynamic ISDN,” and outlines its operation.

Contents.

Introduction

Always On/Dynamic ISDN (AO/DI) is a networking service that provides an always-available connection to packet-based data services through the wide area connection. This service provides several advantages over current practices for dial-up access to packet services.

- For the end-user, there is no need to dial-up the service each time access is desired.
- For the packet service provider, it is possible to give the end-user a notification, such as the arrival of new mail.
- For the Local Exchange Carrier, the switched circuit trunk utilization is more efficient.

For a good overview of the current situation, please see the “The War of the Wires,” The Economist, May 11, 1996, pp. 59-60. (http://www.economist.com/issue/11-05-96/wb1.html)

Description of AO/DI Operation

AO/DI is based on using existing infrastructure of modern central office switches and using existing, or newly emerging, multilink protocols.

- Modern central offices are capable of supplying National ISDN-1 (NI-1), meaning these central office switches are configured with X.25 packet handlers.
- Examples of multilink protocols under consideration are the MLPPP, MP+, and BACP defined in Internet RFCs.

The basic idea of AO/DI is that an ISDN D-Channel X.25 call is placed from the subscriber to the packet data service provider. The multilink protocol and TCP/IP protocols are encapsulated within the X.25 logical circuit carried by the D-Channel. The Bearer Channels are invoked as additional bandwidth is needed. The Bearer Channels use the multilink protocol without the Q.922 and X.25 encapsulation used on the D-Channel.

Using the X.25 over the D-Channel, while admittedly not the most efficient protocol stack, allows AO/DI to take advantage of the existing packet handlers at the central offices. The link associated to the D-Channel X.25 packet connection is used as the primary link of the multilink protocol.

Because the D-Channel is an always-available connectionless packet-oriented link between the CPE and the central office, it is possible to offer an always-available service based on it. Further, because the D-Channel X.25 packets are handled at the central office by the X.25 packet handler, it is possible to route these packets without first crossing the time-division circuit-switched fabric of the switch, which reduces the impact to the telephony network.

The network diagram below provides an overview of how the components are linked together. For further discussion of the networking, please see the VIA AO/DI white paper “AO/DI Network Architecture, Revision 1.0” October 1996.

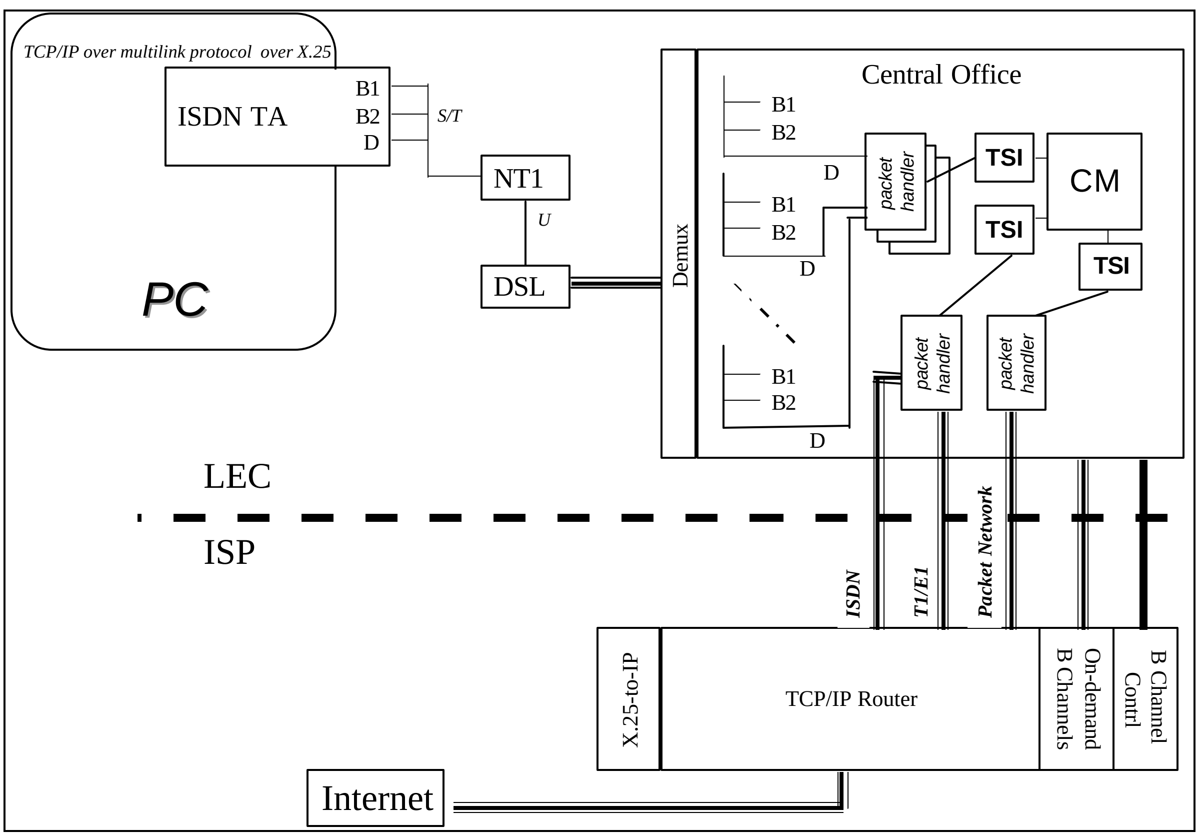

This is a very simplistic diagram to illustrate the basic operation. This diagram does not mean to imply a restriction of one router per central office, or vice versa.

Switched Virtual Circuits (SVCs) and Permanent Virtual Circuits (PVCs) between the subscriber (client) and ISP.

The binding of a subscriber's X.25 packet traffic to a specific aggregation channel depends on the type of connection made. For the PVC
this binding is permanent, whereas for the SVC the binding lasts as long as the circuit as active.

Considerations are:

- Not all the worlds Packet Handler implementations can be guaranteed to support PVCs. (We need to survey the major switch vendors to determine their Packet Handler capabilities.)
- Some service providers that own the ISDN infrastructure may not be an ISP in their own right and may be providing ISPs with a standard X.31/X.75 delivery of D-Channel traffic. If this is the case, there is a need to use (and change) X.121 addresses in order for a user (of the CPE) to be able to change ISPs easily.
- One European service provider will be delivering an ETSI packet handler protocol link to the ISP (or value add service supplier). This will allow the ISP to terminate any protocol within the LAPD.
- An SVC can be treated as a "permanent" connection. Once the call is established it does need to be cleared and can remain in the data state in a similar manner to a PVC.
- The success of X.25 networks was due in part to the use of SVCs and the ease of provisioning. Frame Relay, although successful, is extremely complex to provision because of it's PVC implementation and the same would apply to a managed service provider solution.

Given all these considerations, AO/DI uses SVCs.

Response to the Loss of an SVC

Under certain conditions, the ISP or the LEC may need to disconnect the X.25 SVC. Conditions under which this might happen include:

- Inability to contact the subscriber. This could be due to the PC being turned off, an equipment problem due to hardware or software in the network or the PC, etc.
- A need to redistribute the X.25 SVC across other LEC facilities due to traffic congestion.

When the X.25 SVC is disconnected, the PC should attempt to re-establish the SVC at the earliest convenient time.

Network Connection Sequence

An example of the calling sequence is shown below:

- The subscriber places an X.25 SVC call to the packet data service. This call is meant to be redialed only when the subscriber's PC is rebooted. Otherwise, once the X.25 SVC call is connected, it is intended to remain available.

- When additional bandwidth is needed, the appropriate phone numbers are exchanged between the subscriber's equipment and the packet data service provider's equipment to allow additional Bearer channels to be dialed. The Bearer Channels are routed through the switched fabric directly to the packet data service provider without the use of the packet handlers in the central office. Subsequent to successful connection, the multilink protocols are resolved to aggregate the additional bandwidth into a single transport connection.

Underlying Multilink Protocol Behaviors

AO/DI is built upon the underlying multilink protocol to negotiate for bandwidth, manage phone number exchange, and to aggregate the bandwidth of subsequent connections.

Currently, multilink protocols are not symmetric. A symmetric multilink protocol allows either end of the packet service to place a phone call to the other so that additional bandwidth can be added to the connection. This is a highly desirable flexibility, and one which this group will study as an option.

In today's multilink protocols, the call originator is required to dial the call, thus incurring the additional charges. This is an acceptable model to many people, but not universally so. To mitigate the asymmetry of multilink protocols, packet service providers can give their subscribers toll-free numbers. Please note that the availability of toll-free local data call services are subject to availability from the Local Exchange Carrier.

In any case, it may be desirable to have a user interface that confirms with the user the request for additional bandwidth, should the users be sensitive to these charges.

Strictly speaking, client (CPE) behavior is left to vendor-specific implementation. A vendor may choose to provide differentiation in the features, behaviors, and look-and-feel of the dialer (the software that controls the addition/subtraction of B-Channels) to meet customer requirements for a specific business environment. For example:

- for a voice call, the dialer would need to grant exclusive access of one of the B-Channels.
- for requests to add B-Channels, the dialer may query the user to grant permission depending on the requester; a remote corporate access request might be granted automatically, whereas an unknown source would require manual intervention.

Invoking Additional Bandwidth

Given the relatively low net bandwidth of the packet service due to the low bandwidth of the D-Channel (16 kbps total) and the encapsulation, this method has limited applications. Application domains where the low-bandwidth, always-available are useful are: basic ASCII email services, news feeds, and automated data collection.

To improve the low-bandwidth of the D-Channel X.25 circuit (referred to as the "primary link") is used to add Bearer Channels on demand. The B-Channels are invoked to temporarily boost data throughput, then the B-Channels are disconnected. This mode of operation statistically multiplexes the switch fabric and inter-office trunk lines across more users, thus reducing the traffic impact to the wide area network. Using the B-Channels for bandwidth-on-demand is good for both the Local Exchange Carrier and the packet service provider as compared to having users "camp on" a Bearer Channel.

Traffic Estimates

Traffic estimates and assumptions are presented in the VIA AO/DI white paper "AO/DI Network Architecture." Based on these estimates, the following triggers can be used as a stimulus for requesting additional bandwidth.

Triggers for Requesting Additional Bandwidth

Bearer channels are added if the traffic will take more that 5 seconds to transmit through the D-Channel X.25, or if the pending data is larger than 7500 bytes.

When the amount of data is larger than 7500 bytes, we invoke a B-channel; further, this B-channel establishment, negotiation, and initialization for data takes 3 seconds, meaning the D-Channel X.25 is active for only 3 seconds, or approximately 4500 bytes, before data is no longer sent across it. Thereafter, as long a the B-Channel(s) is active, the X.25 is used to carry BACP traffic, only.

A Example Heuristic for Adding Bearer Channels

One method for determining when additional bandwidth needs to be added is described below. The heuristic is based on an implementation shown in the figure below.

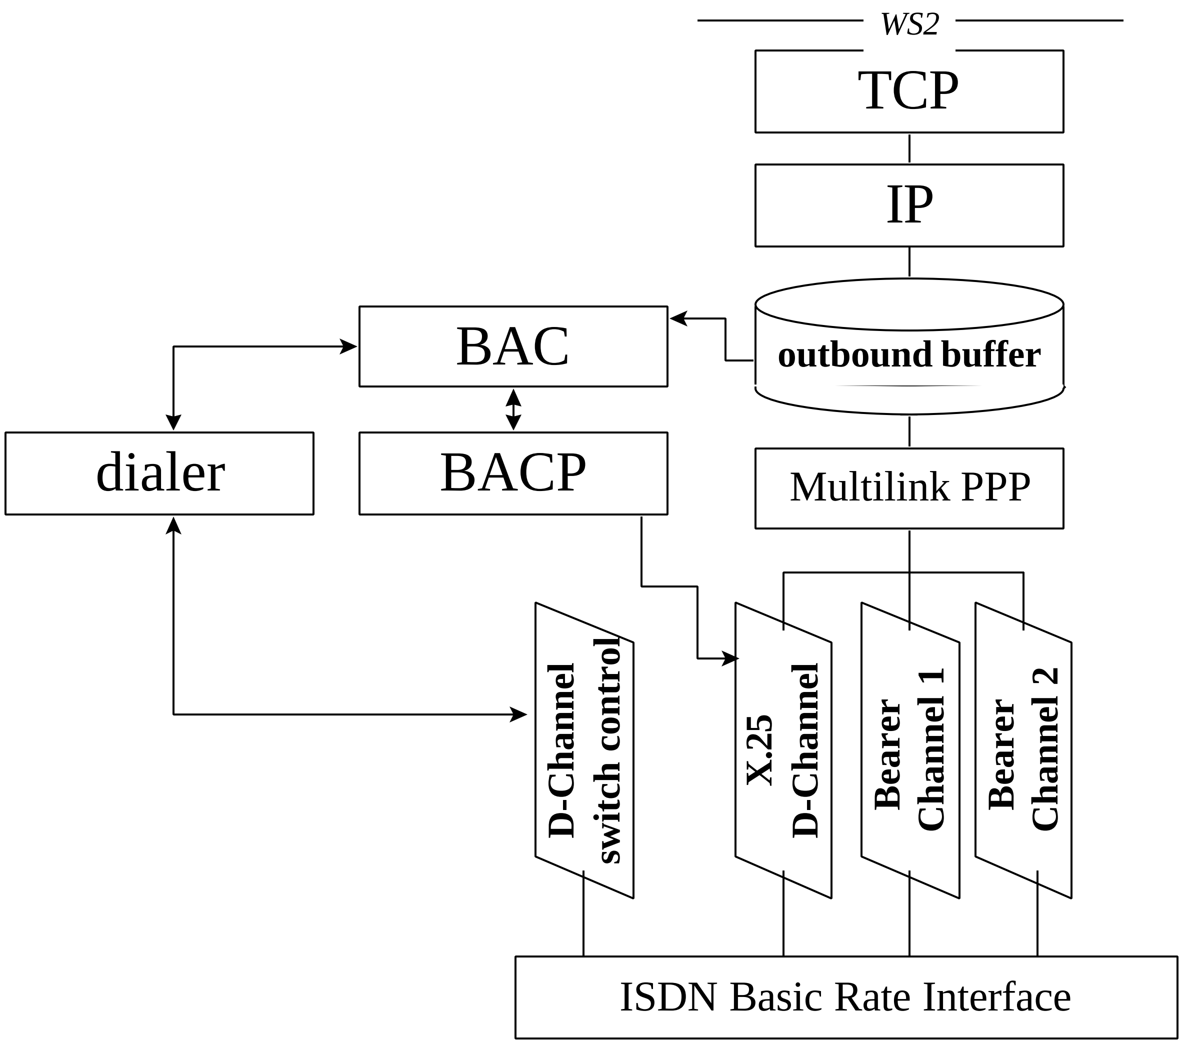

This is a simplistic diagram to illustrate an example of how the determination of when bandwidth should be added. This example shows a method which looks at the total IP traffic. Improvements to this structure would allow applications to reserve bandwidth across one or more of the Bearer channels.

- Is the packet service outbound queue getting full? Where full means that at current throughput, will it take longer than 5 seconds to empty? Will it take longer than 10 seconds?
- If the time to empty the queue is less than 5 seconds, use only the primary link (D-Channel X.25) without invoking a B-Channel.

- If the time to empty the queue is more than 5 second, but less than 10, use only the primary link (D-Channel X.25) of the multilink protocol to request one B-Channel.
  - If a B-Channel is available, use the multilink protocol to augment the packet service connection.
  - If a B-Channel is not available, continue to empty the queue and monitor for queue fullness and B-Channel availability.
- If the time to empty the queue is more than 5 second, but less than 10, use primary link (D-Channel X.25) of the multilink protocol to request two B-Channels.
  - If two B-Channels are available, use the multilink protocol to augment the packet service connection.
  - If only one B-Channel is available, augment the connection with the single B-Channel. Continue to empty the queue and monitor for queue fullness and B-Channel availability.
  - If a B-Channel is not available, continue to empty the queue and monitor for queue fullness and B-Channel(s) availability.

Following this heuristic allows the user the freedom to use the ISDN resources in multiple methods without affecting the ability to augment bandwidth when available. For example, the user may be having an ISDN-voice call simultaneously with the use of the AO/DI.

Connections to the Packet Data Service Provider

Routing D-Channel X.25 packet calls to the packet service provider can be done more efficiently without over-loading the switch trunk lines and the switching fabric. For efficiencies, the X.25 can be concentrated into standard WAN connections (e.g., T1 or PRI) between the central office and the packet service provider; several central office-to-packet service provider options are available and can be decided on their own merits between the Local Exchange Carrier and the packet service provider.

References

[1] K. Sklower, B. Lloyd, G. McGregor, D. Carr, “The PPP Multilink Protocol (MP)”.

[2] Simpson, W., Editor, “The Point-to-Point Protocol (PPP)”, STD 51, RFC 1661, Daydreamer, July 1994.

[3] K. Smith, “Ascend’s Multilink Plus (MP+),” RFC 1934, April 1996.

[4] C. Richards, K. Smith, “The PPP Bandwidth Allocation Protocol (BAP), The PPP Bandwidth Allocation Control Protocol (BACP)”, July, 1996,

[5] A. Kuzma, et al, “AO/DI Network Architecture, Revision 1” Vendors ISDN Association white paper, October 1996.

Author’s Address

Andrew Kuzma
Intel Corporation
MS: JF2-007
5200 NE Elam Young Parkway
Hillsboro, OR 97124, USA

Email: Andy_Kuzma@ccm.jf.intel.com
Phone: (503) 264-8520
Fax: (503) 254-6157