Always On/Dynamic ISDN
A. Kuzma
Intel Corporation
Category: Draft White Paper
November 1996

# Always On/Dynamic ISDN Network Architecture

## 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 discusses issues related to AO/DI packet traffic network engineering between central office and the packet service provider. The issues are introduced and discussed in general terms; switch-specific information (provided by the vendors) provides and overview of the network engineering options for the respective switches.

This document serves as an introduction to AO/DI network engineering, and as an overview of the various options to provide satisfactory packet network performance in real-life deployments. With this knowledge, it is hoped that Local Exchange Carriers (LECs) can determine which particular mixture of packet network technologies to apply to a specific AO/DI deployment, in conjunction with their telecommunications equipment providers.

## 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.

Please see the "Vendors ISDN Association Always On/Dynamic ISDN, Draft RFC 001," at http://www.via-isdn.org for additional information on AO/DI operation between the CPE and the central office switch.

This paper focuses on the operations and engineering required to transport the X.25 packets between the central office and the packet-based data service provider. I.e., packet concentration, packet trunking and networking, central office packet engineering, and central office packet support.

Specific issues we consider are

- the use of Permanent Virtual Circuits (PVCs) and Switched Virtual Circuits (SVCs),
- the packet throughput of a packet handler,
- reasonable expectations of packet traffic,
- scenarios that create constrained packet traffic, and
- desirable methods to control packet traffic under constrained conditions,

# Brief Description of AO/DI Operation

AO/DI assumes modern central office (CO) switches capable of supplying National ISDN-1 (NI-1) or European Telecommunications Standards Institute (ETSI) equivalent. The importance of this assumption is that these central office switches are already configured with packet handlers. The packet handlers are used for signaling between the CO and the CPE, and have the ability to handle X.25 traffic as well.

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[1]. 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.921 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 BACP-based multilink protocol; we use the term “primary link” to refer to the link over which the control packets are sent.

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 impact on the telephony network has become a growing concern as Internet access has become increasingly popular[2].

## Motivation for Using AO/DI

The most common method used today is to connect to a packet service provider using an analog dial-up modem. ISDN dial-up is growing in popularity, especially among those who want to access Internet services. As far as the telephony network is concerned, the effect is the same: circuit-switched resources are used to carry traffic between the user and the service provider. With ISDN, it’s been relatively straightforward to aggregate both bearer channels for greater throughput; with Microsoft’s “ISDN Accelerator Pack 1.1”, it is now relatively simple to aggregate multiple modems, or a modem with an ISDN call.

It is well known that typical packet traffic is very bursty – a fact that shared media, such as Ethernet Local Area Networks, use to provide sufficient bandwidth to many users by statistically multiplexing the available bandwidth among the users in a subnet; Internet traffic follows this pattern. In contrast, the circuit-switched connection provides a maximum bandwidth at all times; hence, the connection has idle time.

By providing a packet network access as a starting point, and adding short-lived circuit-switched connections for throughput, AO/DI can help alleviate the problem of switched-circuit congestion through central offices and trunk lines.

## Network Architecture Evolution

It should be understood that the network architecture proposed herein is a step along an evolutionary path towards providing wide-area networks which are better suited to providing ubiquitous packet network access, in addition to the ubiquitous circuit-switched voice calls that the networks support today. As such, AO/DI represents a hybrid network that has both packet and circuit-switch characteristics, and, as much as possible, uses the existing strengths of each in a collaborative manner.

---

[1] The provider can be a public network service and/or a private network service. In either case IP is used to define the transport layer upon which applications are built. It is expected that business support of telecommuters represents an important use of both private networks (corporate-owned) and public networks (for access to the private networks).
[2] The “The War of the Wires,” The Economist, May 11, 1996, pp. 59-60. (http://www.economist.com/issue/11-05-96/wb1.html)

AO/DI uses ISDN because ISDN contains the necessary technology and network architecture components:

- a packet capability,
- a relatively high-speed circuit-switched voice and data.
- standardized and interoperable on a global scale, and
- becoming readily available.

Further, given ISDN's maturity and availability, ISDN is a natural network platform to understand the business and technological impacts of a mass-market packet network.

# Overview and Definitions

## *Networking Overview*

The figure below provides an overview of the end-to-end network. For this paper, we focus on the X.25 packet networking considerations. I.e., the flow of packets from the packet handlers, through packet aggregation, through the packet connection to the ISP. We assume that the ISP already has networking that provides access to the Internet or equivalent packet network.

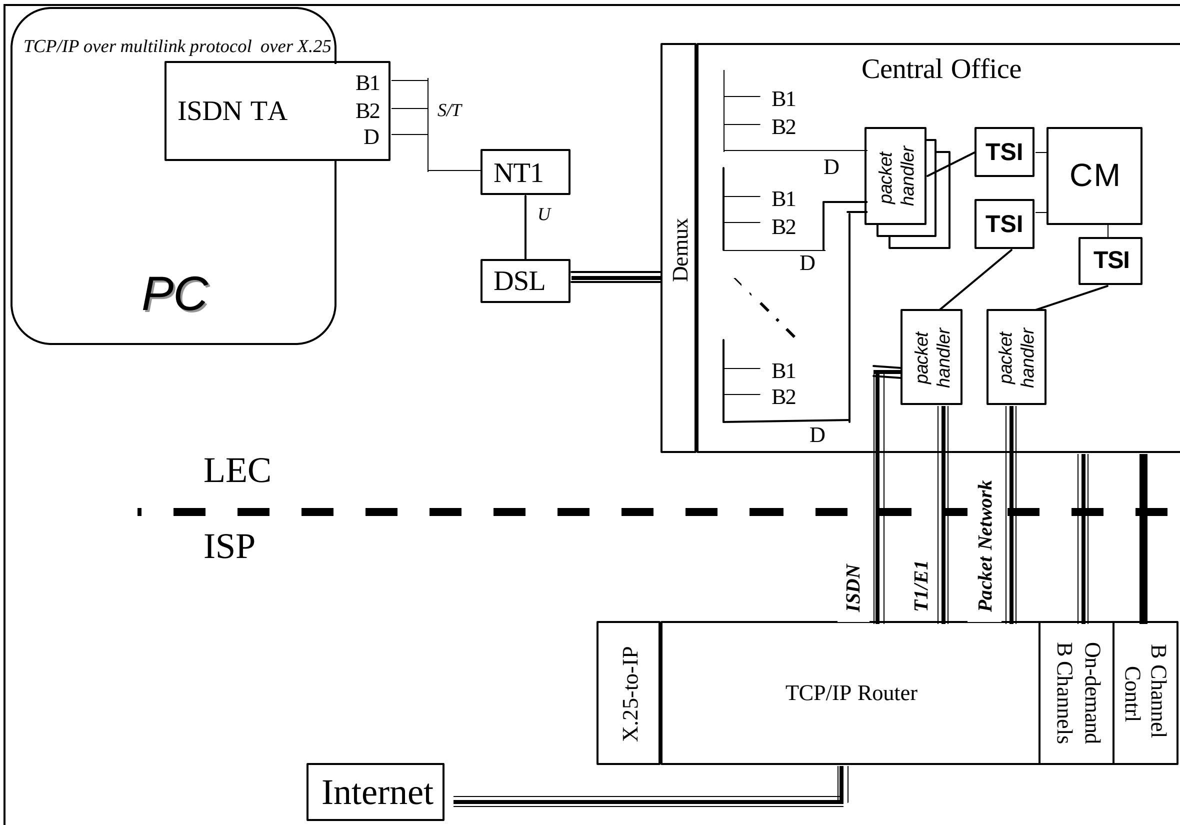


*This is a very simplistic diagram to illustrate the basic operation. This diagram does not mean to imply a restriction of one router per central office, or vice versa. As discussed in the "Wide Area Packet Networking Options," the router(s) can be made to have a virtual appearance across many central offices.*

***Please see the vendors' contributions for a better explanation of how packets are handled in their respective switch architectures.***

Strictly speaking, the service can be to any packet network, but this paper focuses primarily on two business applications: the public Internet, and private Intranets. Both are based on the use TCP/IP protocols and provide TCP/IP services such as name resolution. Our reason for focusing on these two segments is:

1. these are the most popular packet networks, especially email and Web services,
2. Internet/Intranet requirements represents a stress-test of the solutions developed.

### *Function of the Packet Handler*

X.25 packets arrive at the Packet Handler over the D-Channel. The function of the Packet Handler is to route the non-switch control packets to their correct destination. The Packet Handler behaves as a packet forwarding agent for X.25 packets -- the equivalent of a bridge in a LAN. (Note: Not all Packet Handlers function this way. Please see "Switch-specific Considerations" for details.)

Before packets can be forwarded, a packet connection must be initialized so that the packets can have someplace to be sent. The packets will typically be aggregated for networking efficiency. At the time of initialization, the packets are "bound" to a specific aggregation channel for the duration of the connection; in the case of the Permanent Virtual Circuit, this would be forever (or a close approximation thereof), whereas for a Switched Virtual Circuit the binding would last as long as the virtual circuit is considered connected and available to carry traffic.

## AO/DI Subscriber Growth Assumptions

We need to have some reasonable estimates of AO/DI subscriber population in order to estimate the network impact. The impact is measured in two variables:

1. number of AO/DI subscribers per Packet Handler, and
2. number of affected Packet Handlers.

These are two important factors which can be controlled by the LEC to lessen the chances for resource-constrained packet networking.

In Year 1:

the number of users per Packet Handler is 10.
the percentage of packet handlers with AO/DI traffic is 10%.

In Year 2:

the number of users per Packet Handler is 50.
the percentage of packet handlers with AO/DI traffic is 25%.

In Year 3:

the number of users per Packet Handler is 100.
the percentage of packet handlers with AO/DI traffic is 50%.

In Year 4:

the number of users per Packet Handler is 100.
the percentage of packet handlers with AO/DI traffic is 75%.

## AO/DI Traffic Assumptions

Let me first be completely honest by admitting that these are only guesses based on some personal experience. I need to get some corroboration from ISPs about the size of the data files.

1. Packet Handler traffic is symmetric. Whereas my discussion covers sending of email, it also implies reception of email. Further, the Packet Handler resources are independent for transmit and receive. This means twice as many packets are processed through the Packet Handler, divided equally between transmit and receive.
2. Packet Handlers are configured handle up to 128 ISDN users.
3. The LEC has some control, through provisioning, so that AO/DI subscribers can initially be "sparse"; i.e., only a few AO/DI subscribers per Packet Handler.

4. IP packet size is 1024 bytes, equivalent to eight X.25 packets.
5. After switch-CPE signaling traffic, 14 kbps is available for X.25 traffic, on average. In some offerings, the D-Channel X.25 traffic is limited to 9600 bps.
6. The HDLC bit-stuffing accounts for 5% of available throughput. (13342 bps for X.25 traffic)
7. X.25 overhead accounts for 5% of throughput. (12674 bps for IP traffic, or 1584 bytes/second aggregate IP traffic, not inclusive of TCP overhead.)
8. TCP overhead accounts for 5% of IP throughput. (Approximately 1500 bytes/sec for TCP traffic.)
9. 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.
10. In Year 1, very few (less than 5%) of the subscribers will generate significant email with attachments (data files or multimedia) that will exceed 7500 bytes (per email note) in day-to-day use. (There may be transients as people start up their service, but these will settle quickly as people access their email much more frequently.)
11. By Year 2, 50% of the subscribers will be creating email larger than 7500 bytes (per email note), and sending it to more than one recipient.

### *Year 1*

#### Reasonable Expectations of Packet Traffic

Email size will be approximately 700 bytes per note, and users will typically retrieve two notes at the same time. This means about 1500 bytes will be sent over X.25, equivalent to twelve X.25 packets. These twelve packets will be sent in 1 second.

#### Packet Handler Throughput Considerations

Were all users to request this service simultaneously, the Packet Handler would need to service 120 packets/second. This is highly unlikely. Even if all the users decided to access the two email notes within 15 minutes of each other, there is be an available time of 900 seconds, of which 10 is needed for transfer.

Under this scenario, estimate the packet load as 1 user, maximum.

### *Year 2*

#### Reasonable Expectations of Packet Traffic

Assume we stop sending data over the X.25 connection when the Bearer channel is established and negotiated with the ISP. When the amount of data is larger than 7500 bytes, we invoke a B-channel; further, this B-channel establishment, negotiation, and use 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. The number of X.25 packets is approximately 36.

#### Packet Handler Throughput Considerations

Again, using the 15 minute activity window, but now with 50 users, we have 1800 packets in 900 seconds. There is high probability that the Packet Handler will need to handle more than two subscribers. Under this scenario, estimate the packet load as 3 simultaneous users, maximum.

### *Year 3 and beyond*

#### Reasonable Expectations of Packet Traffic

The X.25 traffic generated per user does not change from the estimates for Year 2. The number of users per Packet Handler doubles, however.

## Packet Handler Throughput Considerations

Again, using the 15 minute activity window, but now with 100 users, we have 3600 packets in 900 seconds. There is high probability that the Packet Handler will need to handle more than four subscriber. Under this scenario, estimate the packet load as 5 simultaneous users, maximum

This is likely to be a high-end estimate as the 15 minute window is artificially small over this many people. If we linearly distributed the traffic over a 60 minute period, the total number of packets is still 3600, but the average load would be much reduced.

## *Summary of Packet Traffic Under the Assumptions*

The packet traffic estimates are tabulated below for each year according to the assumptions. At this time, please take this table with a large grain of salt – I'm not a statistician, so my calculations are rather simplistic.

| | maximum # of subscribers per PH | # of active subscribers per PH | # of packets/second/user | Elapsed time to get packets (seconds) | Packet exchange duration window per hour (seconds) | Total # of packets/second | Estimated maximum # of simultaneous users per PH per second | Estimated average # of simultaneous users per PH per second | Estimated maximum # packets/second at PH | Estimated average # packets/second at PH | Estimated maximum # packets/second/user at PH | Estimated average # packets/second/user at PH |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| Year 1 | 128 | 10 | 12 | 1 | 900 | 120 | 1 | 0.011 | 12.0 | 0.133 | 1.200 | 0.013 |
| Year 2 | 128 | 50 | 36 | 3 | 900 | 1800 | 3 | 0.056 | 36.0 | 0.667 | 0.720 | 0.013 |
| Year 3 | 128 | 100 | 36 | 3 | 900 | 3600 | 5 | 0.111 | 60.0 | 1.333 | 0.600 | 0.013 |
| Year 4 | 128 | 100 | 36 | 3 | 3600 | 3600 | 1.25 | 0.028 | 15.0 | 0.333 | 0.150 | 0.003 |

## *Alternative Traffic Estimation*

The primary application, at least initially, driving D-channel traffic is E-mail. A "typical" user might send and receive 20 messages per day. In addition, many E-mail clients poll for E-mail and, depending on the polling interval, will generate additional traffic for that process even when there are no messages to move. It's likely that other applications would also use the link, such as stock-quote or news headline services (such as a lower-bandwidth version of Pointcast). One could imagine that a user might move 100K bytes per day in one direction, most of it during the 10 hours between 8 AM and 6 PM, for an average of 10K bytes/hour.

The packet handler is indifferent to who generates the packets (that is, it shouldn't matter whether the packets are generated by one or many users from a capacity standpoint), and flow of packets will be throttled by the network if bursts arrive simultaneously, so the traffic will be naturally smoothed. If there are 128 users on a given PH, the total average traffic through that PH will be 1.28Mbytes/hour, or 10K packets per hour. If the PH is linked to the X.25 cloud via a 56Kb circuit, there's a capacity of roughly double that, and presumably additional links or higher-bandwidth circuits could be added if traffic needs dictate. (Other experts would have to confirm that a typical PH can operate comfortably at these levels.)

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

As mentioned earlier, 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 world's 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 not 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.

# General Networking Options for Aggregated Packet Traffic

Every major switch offers several options for the aggregation. Examples are X.25 aggregated into B-Channel or a Frame Relay circuit. In the vendor-specific sections, the various methods for aggregation are discussed further.

## *Wide Area Packet Networking Options*

It may be desirable for the ISP to have a presence at multiple switches to serve a larger number of customers. Using the packet networks, it may be possible to create a virtual presence for the ISP. This is done by aggregating the X.25 traffic onto a larger packet network (larger in both bandwidth and geography). This Wide Area Packet Network is connected to the switch from which the ISP has a physical presence.

A wide variety of telecommunications equipment is available for the many packet networks currently offered. WAN packet network options include X.75 and X.75'. Depending on the needs of the ISP and relative costs, the appropriate network can be engineered.

The choice of packet networking equipment is best made in conjunction with the facilities engineers and the switch vendor(s) and other telecommunications equipment suppliers.

## *Constrained Packet Traffic Scenarios*

If there are too many subscribers, and/or subscribers with heavy packet traffic, it is possible for:

1. the Packet Handler to become swamped, and/or
2. the packet aggregation channel to have more demand than capability.

Hence we can find the network in a constrained packet network condition. In order to prevent this several options are available:

1. add more packet handling capability, and
2. improve the overall capability of the packet aggregation networks.

While these are ideal solutions, in the short term they are contrary to the philosophy of an evolutionary approach to improving the network's packet capabilities. In the short term several options are available:

1. let the packets drop. This is the brute force method of flow control. Unfortunately, it places extra burden on the ISP equipment, provides poor service to the customer, and favors a bandwidth-grabbing behavior (the first few packet users could grab an unfair share of the bandwidth). In practice, given the bursty nature of Internet/Intranet use, this may well be a workable solution for a year.
2. put a flow control at the client (subscriber) side that prevents anyone user from monopolizing the X.25 packet network resource. This could be done by putting in a flow control that throttled the total X.25 throughput with a random sequence. This would allow others access to the packet network.
3. While strictly outside the scope of VIA, proper tariffs will help. By establishing a generous, but not infinite limits to the amount of packet traffic, along with low cost Bearer channel rates, users and ISPs will be encouraged to behave in economically rational ways and hence balance their use of the packet network and the circuit-switched network resources.

## CPE Impact: Desirable Methods to Lessen Constrained Packet Traffic Impact

We have identified the following behaviors which the CPE (client) can help lessen the probability of packet congestion.

1. Put in a medium-term limiter to allow others access to D-Channel resources provided by the network.

## Switch-specific Considerations

The following information has been provided by the switch vendors. VIA is grateful to have their contributions and thanks them in advance for their input and cooperation with this effort.

### *Lucent Technologies, Inc.*

The following information was submitted by Ted Kraft of Lucent Technologies. It has been edited to fit within this white paper, but not substantively changed.

#### Function of the Packet Handler (PH)

The 5ESS PH does indeed terminate X.25 layers 2 and 3. All D-channel messages go to the PH. The PH examines the layer 2 SAPI and decides if the message is call control or X.25. If the message is X.25, it terminates layer 3. (Call control is sent to the Switch Module Processor.)

As regards relationship between the TSI and the PH shown in the "Networking Overview" diagram, the actual operation needs to be clarified. The TSI provides the path for packets that need to go from Switch Module to Switch Module through the Communications Module (CM), but there is not any real switching; the TSI simply gets packets on the inter-module time slots.

#### Capacity/Engineering limits:

The 5ESS switch comprises numerous Switching Modules (SMs). In principle, a switch can accomodate 192 SMs, but few metro switches are that large. More than 3000 BRIs can be terminated on an SM depending on engineering considerations.

Each SM contains up to 80 Protocol Handlers, which as noted above, terminate Q.921 and X.25. Of these 80 PHs, 75 active (5 are stand-by, one on each of five shelves).

## Packet Handler constraints:

| | |
|---|---|
| Max no. of logical channels per B-channel | 127 (permanent or switched) |
| Max no. of logical channels per D-channel | 15 (permanent or switched) |
| Max no. of ports per PH<br>*each packet B-channels uses 1 port*<br>*each packet D-channels use 1/4 port* | 32<br>*(so max B/PH = 32)*<br>*(so max D/PH = 128)* |
| Max no. of logical channels per PH | 255 (summed over all pkt chans) |

The PH has a software resource called Logical Channel Control Blocks. Each switched or permanent virtual circuit uses one LCCB. There are 255 LCCBs available on a PH.

The maximum throughput on a X.25 PH is about 200 packets per second summed over all ports, and is relatively insensitive to the packet size. The maximum thoughput is not an absolute number, but above about 200 pps, the delay per packet begins to increase more rapidly. PHs handling trunk connections are significantly faster, and should not constitute a bottleneck for this proposal.

## Wide Area Packet Networking Options

If the packet call must be carried out of the switch to another X.25 network, an X.75 or X.75' trunk is needed. One trunk equals one 64 kbps time slot on a T1, and:

| | |
|---|---|
| Max no. of X.75 and X.75' trunks per PH | 4 (combined X.75 and X.75') |

### *Nortel*

Nortel's DMS-100 and Packet Handler are both designed to quickly and efficiently handle packet email assumptions presented in an earlier section of your document.

The DMS-100 supports both PVC and SVC; however, it would appear that SVC would be the most appropriate implementation scenario for AO/DI.

Our high-level view does not indicate any DMS architectural limitations for deployment of AO/DI. We are currently conducting more in-depth modeling to determine if there are design implications for the future. We anticipate the results of our modeling could yield recommendations for the most effective deployment scenarios based on the assumptions outlined in the Year-1 through Year-3 examples.

Figure 1 shows the high-level design of the DMS 100 architecture including the DMS packet handler known as the LPP. This diagram indicates a data network access example for AO/DI.

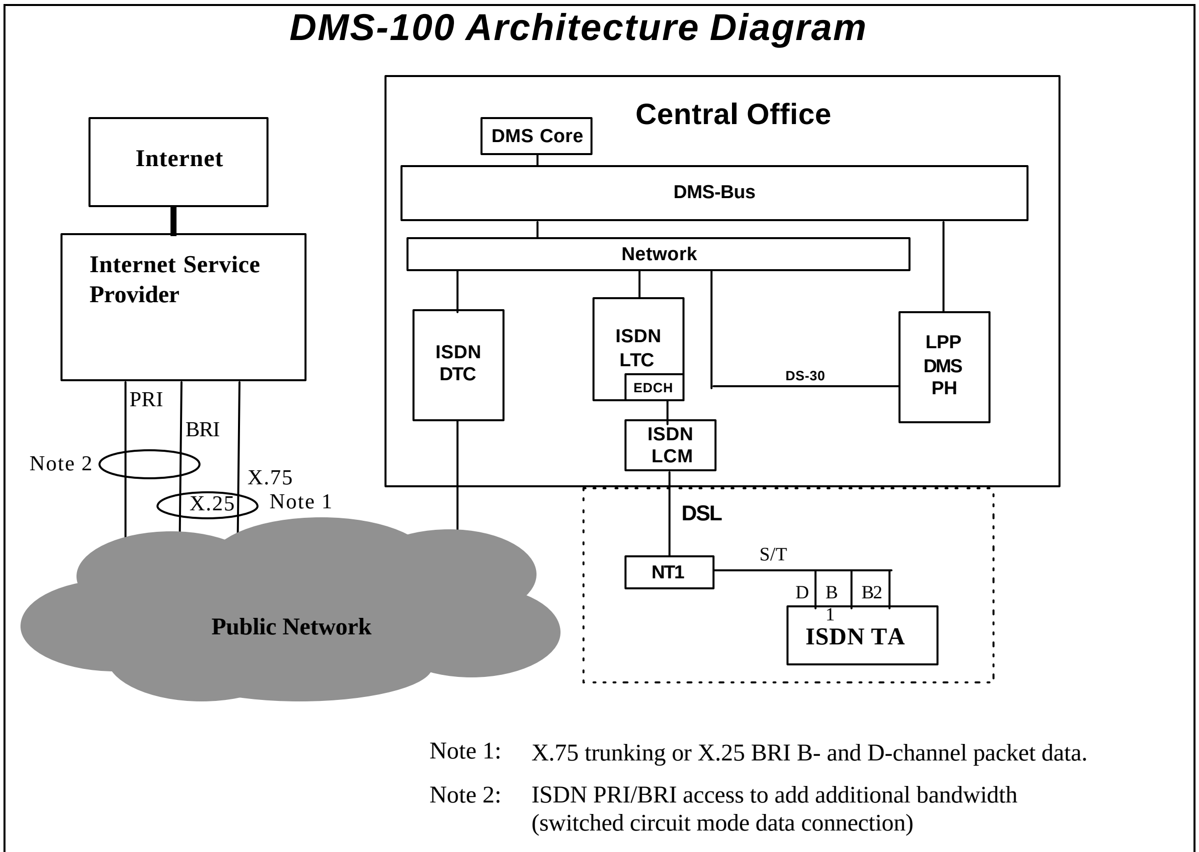


Figure 1

*(Editor's Note: when viewed on the computer monitor, some of the connection lines in this figure may not be visible at normal screen resolutions. If this appears to be the case, either print the document, or view the document at a higher magnification.)*

The following acronyms should help you identify components of the diagram.

| | |
|---|---|
| EDCH | Enhanced D-Channel Handler |
| ISDN DTC | ISDN Digital Trunk Controller |
| ISDN LCM | ISDN Line Concentrating Module |
| ISDN LGC | ISDN Line Group Controller |
| ISDN LTC | ISDN Line Trunk Controller |
| LPP | Link Peripheral Processor |
| PH | Packet Handler |

### *Siemens*

## 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, “AO/DI RFC 001,”  Vendors ISDN Association, October, 1996

# VIA Authors’ 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) 264-6157

**Rob Harrison**
Virtual Access (UK) Ltd
Woodside House
Woodside Road
Winkfield
Windsor
Berkshire
SL4 2DX
United Kingdom

Phone: +44 (0)990 364364
Fax: +44 (0)1344 893519
Email: RobHarr1@aol.com

**David Frankel**
Jetstream Communications
1054 S. DeAnza Blvd., Suite 110
San Jose, CA 95129, USA

Phone: (408) 777-4333
Fax: (408) 777-4343
Email: dfrankel@jetstream.com

**Peter Geier**
Eicon Technology, Inc.
14755 Preston Road
Suite 620
Dallas, TX 75240

Phone:  +1 (972) 239-3270 ext. 333
Fax: +1 (972) 239-3304
Email: peterg@eicon.com

# Switch Contributor Contact

## *Lucent Technologies*

Ted Kraft
Tel: (630) 713-7434

### *Northern Telecom*

Wayne Shaw
Tel: (919) 992-8441

Nancy Adcox
Tel: (919) 992-5280

### *Siemens Stromberg-Carlson*

Ray Massé
Tel: (561) 955-8141