**Bionic Buffalo Tech Note #22:**

# Asynchronous Transfer Mode (ATM) Addresses

*last revised Monday 14 September 1998*



## OVERVIEW

Asynchronous Transfer Mode (ATM) cells each contain addresses, used to route the cells through the network. This paper introduces the two parts of the address (VPI and CPI), and provides a general description of how they are used in the context of the overall architecture.

## LOCAL ADDRESSES: VPI & CPI

There are two kinds of interfaces between nodes in an ATM network: the User-Network Interface (UNI) and the Network-Network Interface (NNI). (Each of these has variations, so not every interface of one type is compatible with every interface of the same type.) These are two kinds of software/protocol interfaces; hardware is a separate issue.

NNI is used between switches within networks, usually large-scale public or private networks. Some switches use UNI, rather than NNI. Other switches use UNI for connections to the outside of the network, and NNI for connections to other switches within the network.

UNI is used everywhere else. Most ATM equipment uses some form of UNI, including workstations, telephones, faxes, PBXs, and other devices not part of the backbone or major switching systems. Almost all customer equipment uses UNI, including use for connections to the phone company or other external carrier.

ATM messages are encapsulated into packets, called *cells*. Each cell contains a *virtual channel identifier* (VCI) and a *virtual path identifier* (VPI), which is used to route the cell (if outgoing) or identify the cell's source (if incoming). At the NNI, 12 bits are available for the VPI. At the UNI, 8 bits are available for the VPI. For either NNI or UNI, 16 bits are available for the VCI. Although the packet format allows VPIs and VCIs of these sizes, most equipment does not allow all bits to be used, and sets the unused (high-order) bits to zero.

A virtual path, identified by the VPI, is a route between two nodes in the network. A path may pass through an intermediate node, without terminating. (As an example, see the path from *A1* to *C1* in the illustration, below.) Each virtual path may contain one or more virtual channels, each identified by a VCI. (For example, the path from *A1* to *C1* contains virtual channels *U* and *V*.) Virtual channels are one-way. Two-way communication requires a pair of virtual channels.

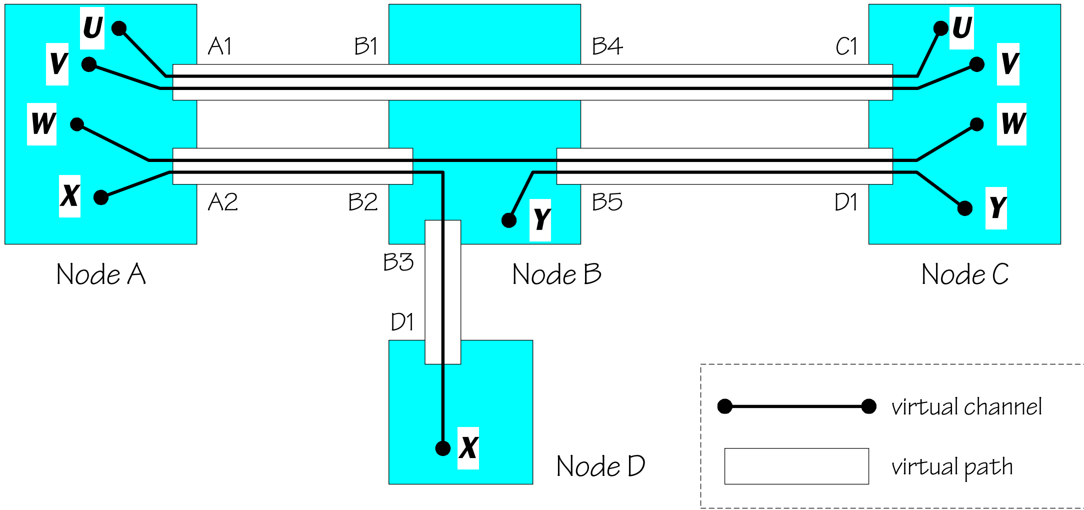


A virtual channel may exit one virtual path, which terminates within a node, and then enter a different virtual path terminating at the node, to proceed thereby to another node. (Channel *X* illustrates this.)

A VPI and VCI are meaningful only in the context of a specific connection and node. As a cell passes through the network, the network may give the cell a different VPI and VCI as it passes through each node. However, the VCI and VPI do not change *between* nodes, since (of course) there is no way to renumber them in transit.

If a node has more than one connection or link, then each link may number VPIs independently. In other words, <u>two different VPs with the same VPI may found in a single node</u>. These diagrams do *not* emphasize this point.

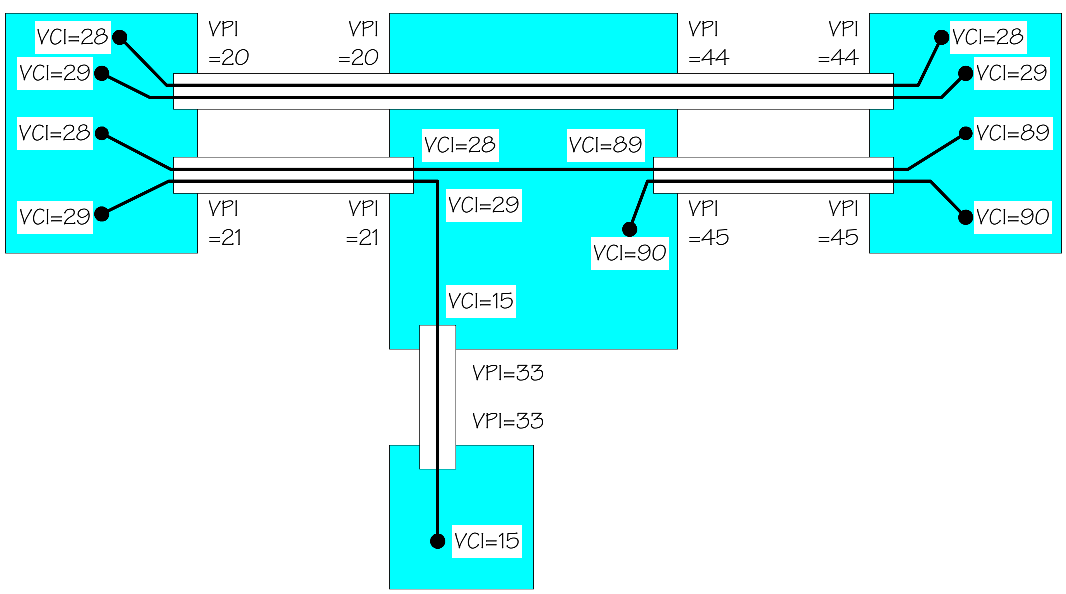


 14 September 1998
http://www.tatanka.com

A virtual path, which passes through an intermediate node without termination, may be given two different VPIs on the two different links. However, the VCI of a cell within a path will not change within a node unless the containing path terminates at that node.

## ALLOCATING VPI & VCI NUMBERS: PERMANENT ASSIGNMENT

Certain VCI & VPI numbers are assigned in advance, either by specification or by optional local configuration, while others are assigned dynamically as required.

Various ATM specifications allocate VCI and VPI numbers for specific purposes. For example, an endpoint will use VPI=0, VCI=5 to ask the network to set up a call to another endpoint. There are also specific addresses reserved by non-ATM specifications for applications on ATM networks. For instance, the DAVIC specification calls for VPI=0, VCI=33 as part of the MAC layer initialization procedure for set-top boxes.

An important permanent assignment is VPI=0, VCI=16, which belongs to the *Interim Local Management Interface* (ILMI). ILMI defines a *Management Information Base* (MIB) accessed using the *Simple Network Management Protocol* (SNMP). The database allows end-points to register their addresses for retrieval by other nodes. If those addresses are assigned systematically, then a simple directory service is available to the network. (For example, the ESI can be assigned conventionally as part of the complete address, so that addresses containing specific ESI values can have special significance.)

System designers sometimes assign specific VPI and VCI numbers for permanent or semi-permanent channels to specific applications or for specific connections. Uses include not only application-level communication (such as video transport) but also management-level or system level communication (for initialization, diagnostics, or control).

These first two assignment mechanisms can create *permanent virtual circuits* (PVCs). Such PVCs require a simple ATM architecture:

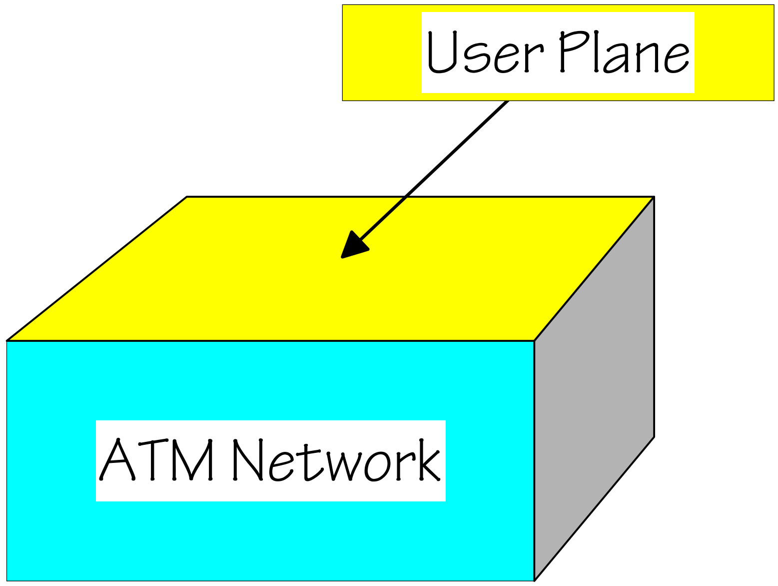


The top surface of the block is called the *user plane*, because it represents the network as seen from the point of view of the user. For example, viewing a simple PVC from the user plane, it may appear as follows:

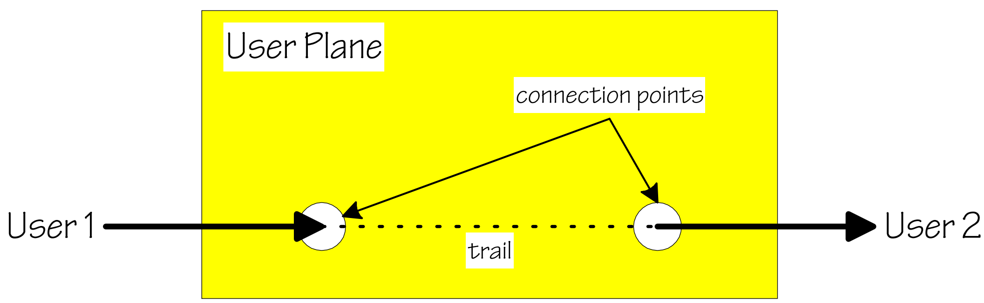


From the user plane, the network simply associates two connection points, allowing the user to move messages between them. The details of how the connection is made are hidden from the user. From the ATM network's point of view, however, there are protocol stacks involved. If the same connection is viewed from the side of the block, it looks like this:

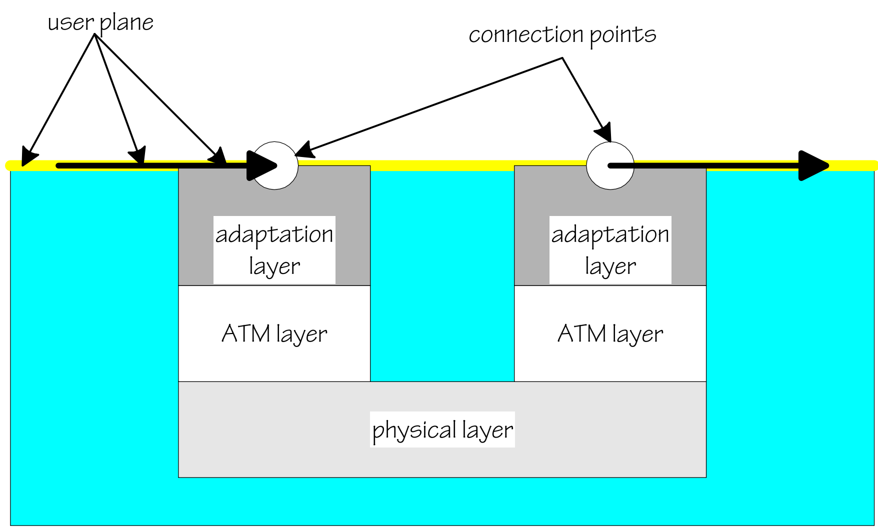


The ATM layer of these protocol stacks translates the VPI and VCI numbers, as the cell moves through the network. Other functions of the ATM layer include cell header process, cell multiplexing and demultiplexing, and some flow control. If a cell passes through a single intermediate node while going from one connection point to another, the associated protocol stack might have the following configuration:

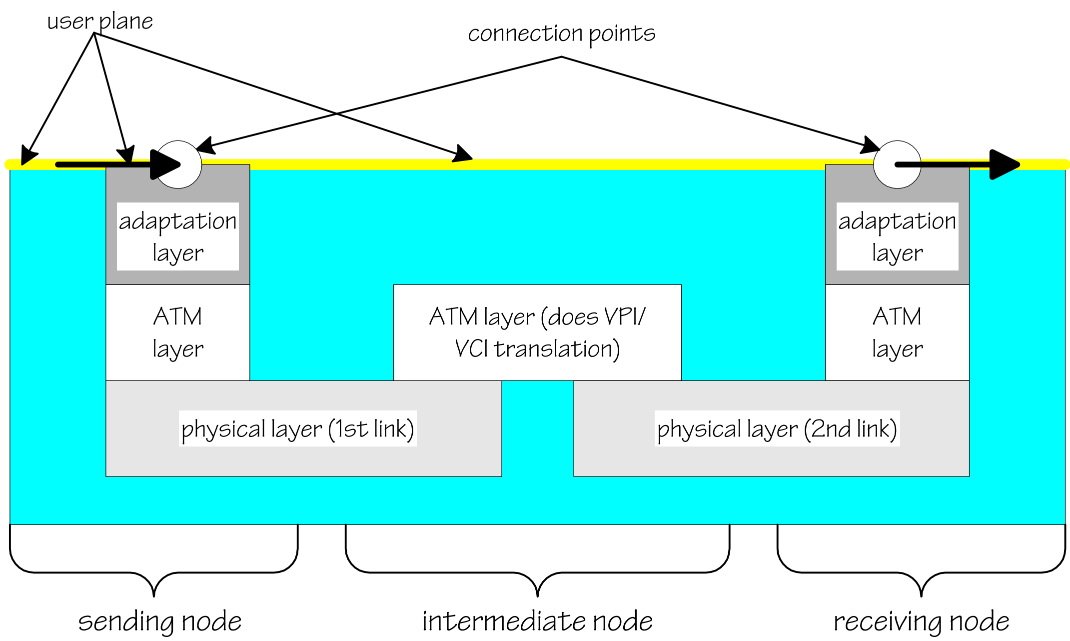


In the most complex and interesting cases, however, the VPI and VCI numbers are assigned dynamically as part of the routing and switching process, as connections are established and broken down. Adding these functions requires increasing the complexity of the architecture, as explained in the next section.

## The Control Plane

To allow applications and system software to control the behaviour of the connections and protocol stacks, ATM adds several more planes to the model. After the user plane, the next plane is the *control plane*.

The control plane is a parallel entity to the user plane. Control messages use the same underlying transport structure as the user connections. Viewed from the user's perspective, the two planes together exist side-by-side. The difference is that, while the user plane connects one user to another, the control plane connects users to the switch or local exchange.

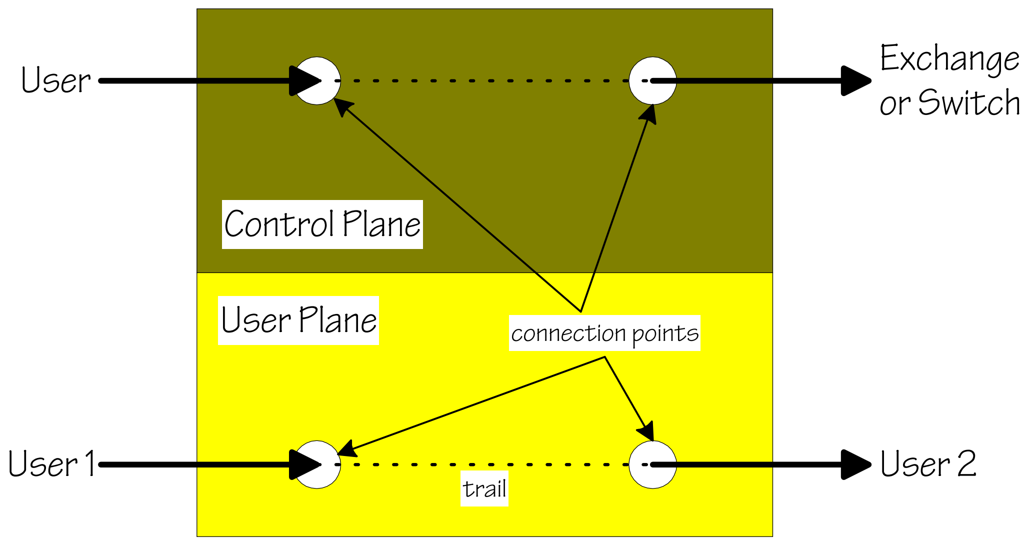


Signalling System 7, or SS7, also calls the user plane, the *information plane*. Because of the close relationship between the architectures of ISDN, B-ISDN, ATM, and SDH, the terminology often is interchanged.

Overall, the ATM network with the user and control planes can be viewed this way:

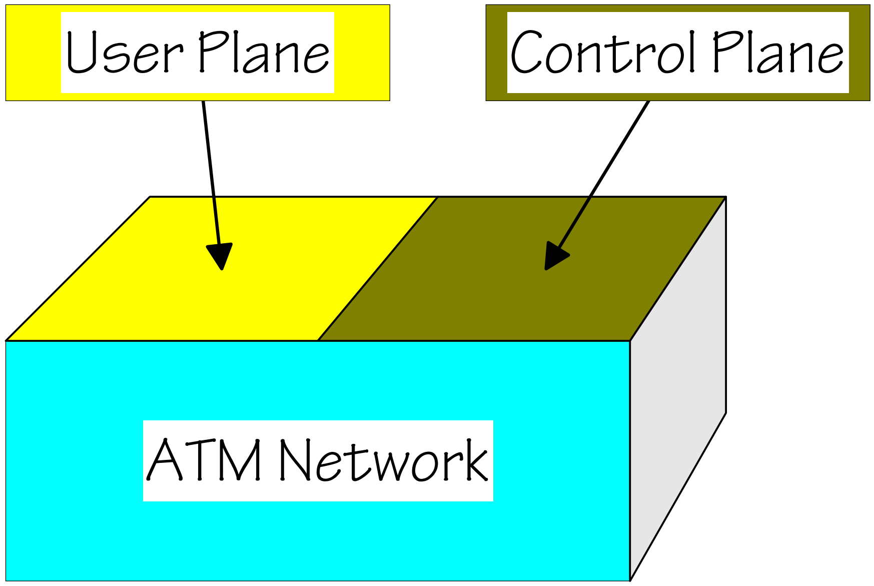


Underneath the user and control planes are the same protocol stacks, although there are differences at the adaptation layer level.

The connection between the user and the exchange or switch is used for signalling protocols, which handle call establishment and release, as well as some other control functions relating to switched services.

## Cross-Connects, Concentrators, and VPCIs

There are two basic types of ATM switch: VP switches (often, these are also simply known as cross-connects), and VC switches. A VP switch can terminate or re-route paths. A VC switch can terminate and re-route both paths and channels.

Cross-connects are (obviously) simpler than VC switches. One common use of a cross-connect is as a concentrator or multiplexor. A cross-connect can take a fast link containing more than one virtual path, and re-route those paths into separate, slower links. In such applications, the end-points of a network often are “aware” of the cross-connect and the way it is connected to the network.

If an end-point uses a VPI to specify a path, the result might be considered ambiguous when cross-connects are used. VPIs are unique only on a given link, and cross-connects have multiple links. The VPI will not uniquely identify all of the paths entering the cross-connect. To allow unique identification of the paths traversing or terminating at a cross-connect, the *Virtual Path Connection Identifier* (VPCI) is used.

VPCIs are returned to the endpoints by the switching protocol when connections are established.

A VPCI is 16 bits long, longer than a VPI. This reflects the fact that, with cross-connects, the number of VPCIs may easily exceed the number of unique VPIs.

A VPCI identifies a *Virtual Path Connection* (VPC), which can be considered a sequence of individual *Virtual Path Links* (VPLs). Each of the VPLs of the VPC has its own VPI. Since the VPI changes upon entry to and exit from a switch (that is, the virtual path is re-numbered), a VPL generally extends only over the physical link from one switch to the next.

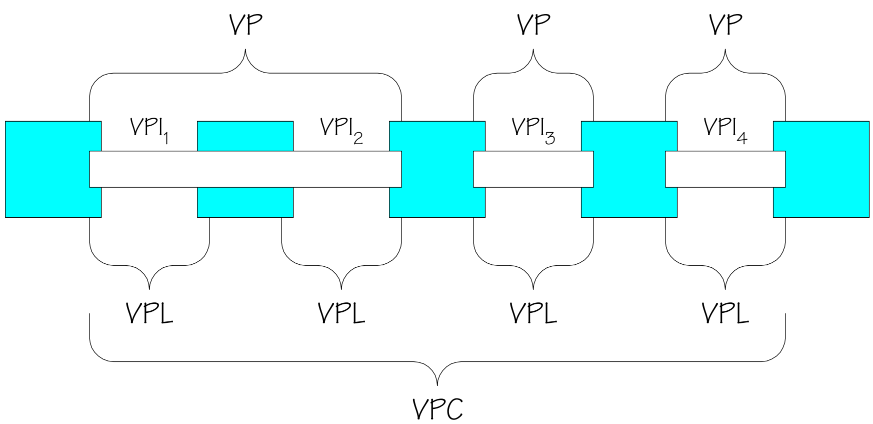


The corresponding entity for channels is a *Virtual Channel Link* (VCL), which has its own VCI. Accordingly, VCLs change when paths are terminated or the channel is switched from one path to another. A sequence of VCLs from one user to another is a *Virtual Channel Connection* (VCC), identified by a *Virtual Channel Connection Identifier* (VCCI).

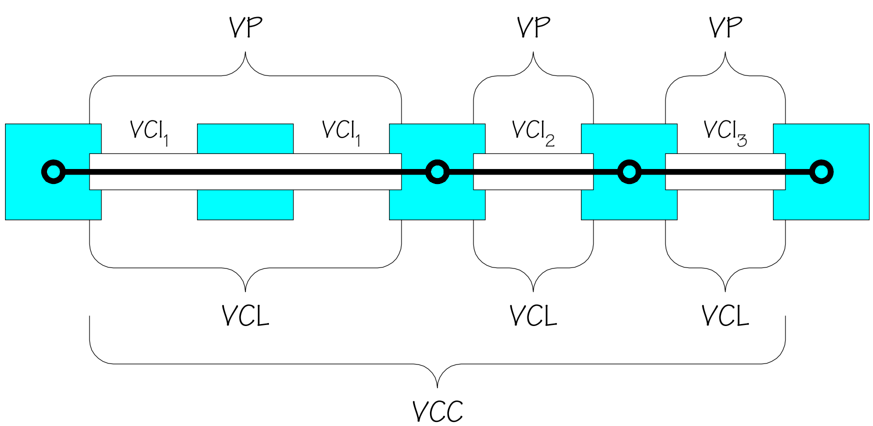

VPCs and VCCs, along with their identifier VPCIs and VCCIs, are used also in discussing configurations which do not include cross-connects. In general, they are used when discussing complex paths or connections that span multiple links and nodes.

## Switched Services and Control Plane Signalling

In addition to facilitating dynamic assignment of VPI and VCI, the control plane signalling protocols give the user a regional or global view of the network.

If ATM networks were limited to fixed VPI and VCI numbers, then the total address space would be no greater than 28 bits (or 32 bits for NNI nodes). By dynamically assigning VPI and VCI numbers, a much greater addressing range is possible.

The dynamic assignment of addresses at each node is based on a very different philosophy from that of common IP networks (using TCP/IP or UDP/IP). In IP networks, each node knows the global address (IP number) of the nodes to which it is connected. The total number of addresses (or nodes) in an IP network is limited to the largest IP number (32 bits), while there is no practical limit to the total number of nodes in an ATM network. The limitation at an ATM node is on the number of *simultaneous* channels and paths, and is not practically reached.

(Note that IPv6 increases the address space to 128 bits, and that there are various address translation schemes to ameliorate the problems of having a restricted address space in IP networks. We also emphatically state that this situation, out of any specific context, does not make ATM's addressing scheme "better" than IP's in any general sense. There are disadvantages to ATM's approach, as well as advantages.)

A user on an ATM network, who wishes to connect to, or call, another user, sends a *setup* message to the network. In this case, the "network" is the exchange or switch, which is reached using a pre-assigned VPI and VCI in the control plane. In response, the network sends a similar setup message to the destination node. The response to a setup message is a *connect* message. After the destination node connects, then the network connects to the caller.

The setup message contains the global or regional addresses of the calling and called nodes, as well as other parameters regarding the type of service and protocols to be used for the connection. The addresses used are not VPI/VPI numbers, and are described in the next section.

In addition to these messages, there are additional messages to indicate progress and to acknowledge successful connection. The overall procedure is:

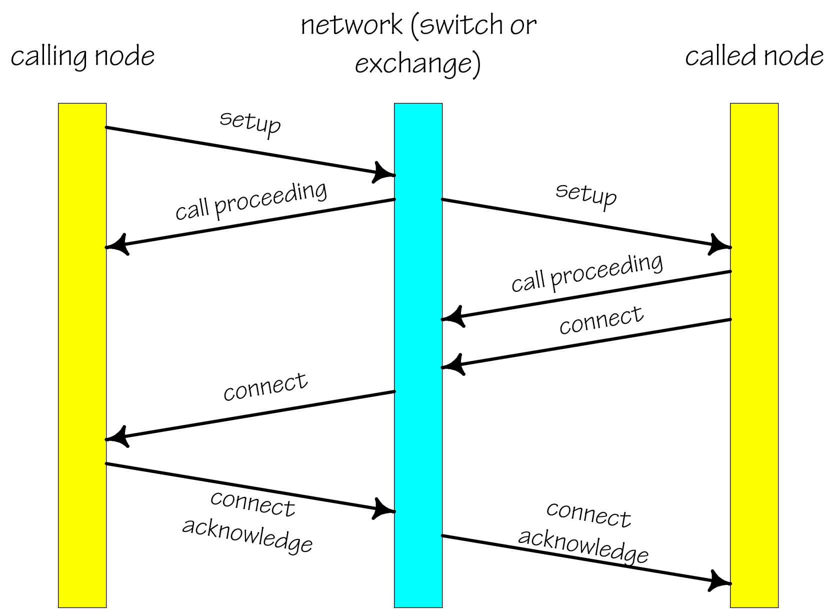


Setup messages, either from an end-point to the network, or from the network to an end-point, may suggest or require VCI/VPCI values. (VPCIs are used instead of VPIs because - as explained above - VPIs are ambiguous in the context of cross-connects.) If a setup message requires specific values which are not available, then the call may fail to complete. The values eventually used, when the call is successfully completed, are returned in the connect message.

## Address Formats for Call Establishment

The call setup message (as specified by Q.2931) allows six address types for the calling and called nodes, and three standard address formats. The six address types are:

- *network-specific numbers* use a format defined by the specific network, and are used for administrative purposes (such as "operator" or "information")
- *abbreviated numbers* use a format defined by the specific network, and are intended for short forms of other address types
- *subscriber numbers* are local to a specific area, PSTN, or network, and normally use one of the three standard formats (described below)
- *national numbers* include a subscriber number, along with routing information to a specific PSTN, region, network, trunk, or other subset of a country's networks, and normally use one of the three standard formats
- *international numbers* include the country code, and normally use one of the three standard formats

- *unknown types of numbers* may use one of the three standard formats, or may use another format

In the three standard address formats, an *Authority and Format Identifier* (AFI), consisting of two BCD digits in a single octet, indicates the format of address used. Not all networks support all of the formats.

- Data Country Code (DCC) format (AFI=39)
- International Code Designator (ICD) format (AFI=47)
- ISDN Telephone Number (E.164) format (AFI=44 for decimal numbers, AFI=45 for binary numbers)

Each of these address formats includes a single octet *selector* field, which is ignored by ATM itself, but which may be used by the end systems as needed. Together with the AFI and selector octets, the total available space for addresses in the setup message is 20 octets (160 bits) for each end point.

Some networks, by private agreement among the parties, also may support other forms of addressing. However, the available address space in all cases is limited to the same 20 octets used for the standard formats.

Architecturally, the standard addresses are equivalent to OSI Network Service Access Point (NSAP) addresses, and have the same formats. The selector field is equivalent to an NSAP selector.

## DCC Format

This format conforms to IEEE 802.x recommendations, as commonly used in LANs, but also found in other environments. The authority for the country specified in the Data Country Code (DCC) field of the address assigns or delegates assignment of the addresses. ISO 3166 gives the DCC values.

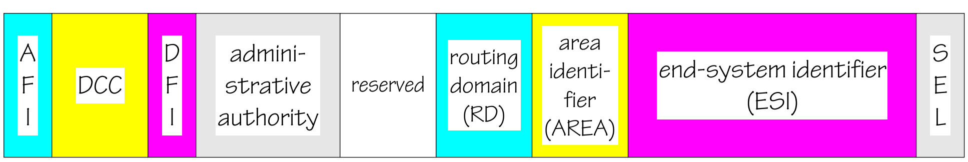


The third field of the address, after the AFI and DCC, is the *Domain Format Identifier* (DFI), which gives the format of the rest of the address. The *administrative authority* identifies the national organization, which allocates the subsequent fields of the address.

The *routing domain* (RD) and *area identifier* (AREA) are part of a hierarchical address space organization. The *end-system identifier* (ESI) gives the address of the specific end-point. To allow autoconfiguration of addresses, the network can base the ESI on the MAC address.

## ICD Format

This format is almost identical to the DCC format. International organizations (as opposed to individual countries) use this format for address allocation. The DCC field is replaced by an ICD field. The British Standards Institute maintains values for the ICDs.

## E.164 Format

This format is based on ISDN telephone numbers. Eight octets (64 bits) are used to specify the number, which can be up to 15 decimal digits long.

ISDN numbers consist of a *country code*, a *national destination code*, and a *subscriber number*. Country codes are from one to three digits. Usually, the subscriber number is the "local" telephone number, and the national destination code is the area code, city code, or other routing information within the country. Both the subscriber number and the national destination code may be of any length, as defined by national standards, so that the total length is no more than 15 digits.

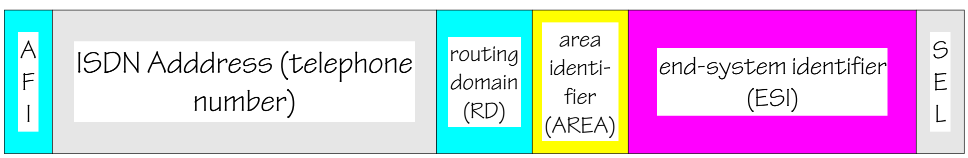


The other fields may be used for routing, or for specifying sub-addresses within a destination subscriber's number.

# Internetworking Addresses

When ATM networks are connected to non-ATM networks, there are sometimes formalized relationships among the different networks' addressing scheme. In some circumstances, it is possible to compute all or a portion of one network's address from the other network's address.

Because of the complexity of the issues, and the variety of situations, detailed discussion of these address relationships are beyond the scope of this document.