The ATM Forum

Technical Committee

ATM Forum Addressing: User Guide version 1.0

AF-RA-0105.000

January, 1999

The ATM Forum
Worldwide Headquarters
2570 West El Camino Real, Suite 304
Mountain View, CA 94040-1313
Tel: +1-650-949-6700
Fax:+1-650-949-6705

Acknowledgments

The Routing and Addressing working group was chaired by Mickey Spiegel. Greg Wetzel was the editor for the *ATM Forum Addressing: User Guide* specification, throughout most of its development. Tim Dwight took over editing responsibilities for Greg in the final stages of document preparation. The minutes at related working group meetings were recorded by Tim Dwight.

The following people made significant technical contributions to the *ATM Forum Addressing: User Guide* specification:

David Cypher
Tim Dwight
Bob Epley
Joel Halpern
Feiling Jia
Demos Kostas
Tom Shaver
Ethan Spiegel
Katie Taaffe
David Wells
Greg Wetzel

This page is intentionally blank.

1 Scope

This document on basic addressing capabilities for ATM Forum specifications is informational. It assumes a basic knowledge of the ATM Forum User-Network Interface (UNI) and Private Network – Network Interface (PNNI) specifications, but can be understood without extensive knowledge of these specifications. A companion document entitled "ATM Forum Addressing: Reference Guide" [ATMF ARG] contains further details on ATM addressing.

The User Guide provides an introduction to basic addressing concepts and terms, and examples of basic network addressing plans and operation. The Reference Guide provides a consolidated view of the addressing related aspects of the various ATM Forum specifications (e.g., PNNI 1.0, UNI Signaling 4.0, ILMI 4.0) and as such provides more detailed information.

The intended audience for this document is private network engineers and administrators, i.e., the people who design and run enterprise ATM networks. It may also prove valuable to ATM end-users and ATM Service Provider personnel.

The goal of the document is to allow ATM network administrators to understand:

- where to obtain ATM End System Addresses (AESAs);
- how they are used in private networks;
- what the ramifications are of attaching to ATM Service Providers (ASPs);
- how calls are established end-to-end in the situations where the private network is "stand-alone" and where it is attached to an ASP.

Where helpful, this document will draw analogies between ATM Forum addressing and IP addressing. In many respects they are similar, even if their use is different (e.g., IP datagrams contain source and destination addresses in each datagram, while only the connection establishment messages contain addresses, the Called and Calling Party Numbers, in ATM Forum specifications).

2 Concepts and Definitions

The text in this section is informative and does not attempt to define concepts and terms completely in all their potential variation. Its objective is to provide definitions of the typical uses.

2.1 ATM Addresses: What are They?

An address is associated with an interface, in particular, a UNI or its virtual equivalent, such as a Soft PVC owner. While it is possible that an address might be construed to identify other entities such as connection termination points or end systems, this is not the case in ATM Forum specifications. Consider, for example, a Public UNI. It is often the case that the device attached at the interface is an ATM switch in a private network. Thus the connection does not terminate there, yet the ATM Service Provider (ASP) network may assign an address to this interface. From the point of view of the ASP network, the connection "terminates", but in fact it does not. Consider also the establishment of an SPVC. The address used to designate the "destination" of the SPVC call setup is not associated with an end system. Therefore, an address designates the location at which the connection is to be delivered without regard to whether the connection continues beyond that interface by concatenation of further connection segments or whether the interface is attached to an end system (and thus truly terminates the connection).

An address may have two functions: location and identification. To ensure a manageable routing system, an address should have significance within the network topology, that is, it should be a *locator* that indicates *where* in the network topology (not geography) the interface can be found. As such, addresses should not be portable from one ATM Service Provider to another, or from one network to another, or even necessarily from one part of a network to another, as this defeats their ability to function as locators. There are, of course, exceptions, but every exception causes the routing system to have to deal with routes that cannot be aggregated and so increases routing table sizes throughout the ATM network (obviously an impediment to creating a network on a global scale).

The term "address" may sometimes be used as an *identifier* that indicates *what* the function is (rather than *where* the interface is). For example, Ethernets have addresses that are identifiers, but that are not locators. In another example, a telephone number has traditionally been a locator and an identifier. However, freephone (e.g., 800) numbers are only identifiers. Local number portability requirements remove the location function and leave only the identification function to a telephone number.

The ATM Forum uses the term address as a locator only (an exception is made for group addresses, which are discussed in [ATMF ARG]).

We distinguish between an address and an address prefix. An address prefix is an abstraction of the block of addresses (locators) that the prefix summarizes. Address prefixes occur in a variety of circumstances:

- call routing tables, in which the prefix stands for any address beginning with that prefix and that is not more specifically mentioned
- call screening, in which an address prefix is used as a pattern against which addresses may, or may not, match and then be screened (e.g., calling party and calling party prefix validation at a public UNI)

2.2 Categorization of Addresses by Structure

The ATM Forum specifications define the use of two major categories of addresses: E.164 and ATM End System Address (AESA, a.k.a. NSAP). The E.164 address space is the same as the one used by the telephone network. It is defined by the International Telecommunications Union – Telecommunications Sector, ITU-T [ITU-T E.164]. AESAs are derived from the International Standards Organization's (ISO) Network Service Access Point, NSAP [ISO NSAP]. There are several types of NSAP, and therefore several types of AESA (an AESA is simply an NSAP that is used to address ATM end systems).

2.2.1 E.164 Address

The ITU-T recommendation E.164 specifies the use of numbers in establishing ISDN calls by public networks (public here means recognized by the country's law and regulations as a network). The ITU-T assigns a country code which is generally associated with a specific country, except in cases of integrated numbering plans like the North American region in which eighteen countries including Canada and United States share a code. Although ITU-T Recommendation E.164 specifies a county code length of from one to three digits, only three digit codes are assigned on a going forward basis. E.164 also defines the maximum length of a number to be 15 digits (specific regions may define fewer than 15, but not more). Within a given region, a Numbering Plan Authority decides how the digits after the region code are to be structured and assigned, which differ from region to region, country to country.

In North America, the address structure is defined in the North American Numbering Plan (NANP). NANP numbers consist of a Numbering Plan Area (NPA, a.k.a. "area code") of three digits, followed by a Central Office Code (NXX) of three digits, followed by a Station Number (XXXX) of four digits. This results in a 10 digit numbering structure (plus "1" for the "North American country code"). The structure is commonly referred to as "+1 NPA NXX XXXX"[1]; a nomenclature which admittedly mixes field names and coding.

Note that the ITU-T uses the term "number" in E.164. This is because an E.164 "number" can be used for purposes other than to "address" a specific subscriber line (telephone). E.164 "numbers" are also assigned for the purposes of service identification. For example, the "800" numbers used in the NANP and other countries. These are not addresses (locators), per se, because they do not locate a subscriber line, rather they are numbers that the telephone network uses to identify a service. For example, in the case of "800" service, the called party pays for the call, not the caller, and the 800 number is translated into an address (non-800 number locator) to which the call is actually to be delivered.

Note also that there is a distinction between the "dialing plan" and the "numbering plan". Dialing plans are needed because the user of a telephone has only the keypad or dial to work with. A dialing plan is network specific and allows the user to escape from the normal mode of dialing a telephone number, to a mode in which the user instructs the network to do something. For example, by convention among local exchange networks in the USA, dialing "1" means that the user wishes to enter a 10 digit number, rather than the usual 7 digit number; dialing "011" means that

---

[1] Characters 'N' and 'X' in the Central Office Code and Station Number are used to denote different ranges of decimal digits. The digits valid in a particular portion of the address are defined by the numbering plan.

the user wishes to enter an international number; dialing "*69" means that the user wishes to return a missed call. Dialing plans should not be confused with numbering or addressing plans.

In ATM Forum specifications, only E.164 numbers that are addresses are recognized. Further, they are sometimes referred to as *native* E.164 addresses to distinguish them from an embedded E.164 AESA, discussed below.

2.2.2 ATM End System Address (AESA)

ATM End System Addresses (AESAs) are defined by the ATM Forum to have 20 octets and come in many types. The 3 types used in specifications UNI 3.1, UNI Signaling 4.0, ILMI 4.0, PNNI version 1.0, etc., are shown in Figure 1. Also shown is the Local AFI used by VTOA to the Desktop (see [ATMF ARG]). The various types of AESA differ mostly in what authority assigns them.

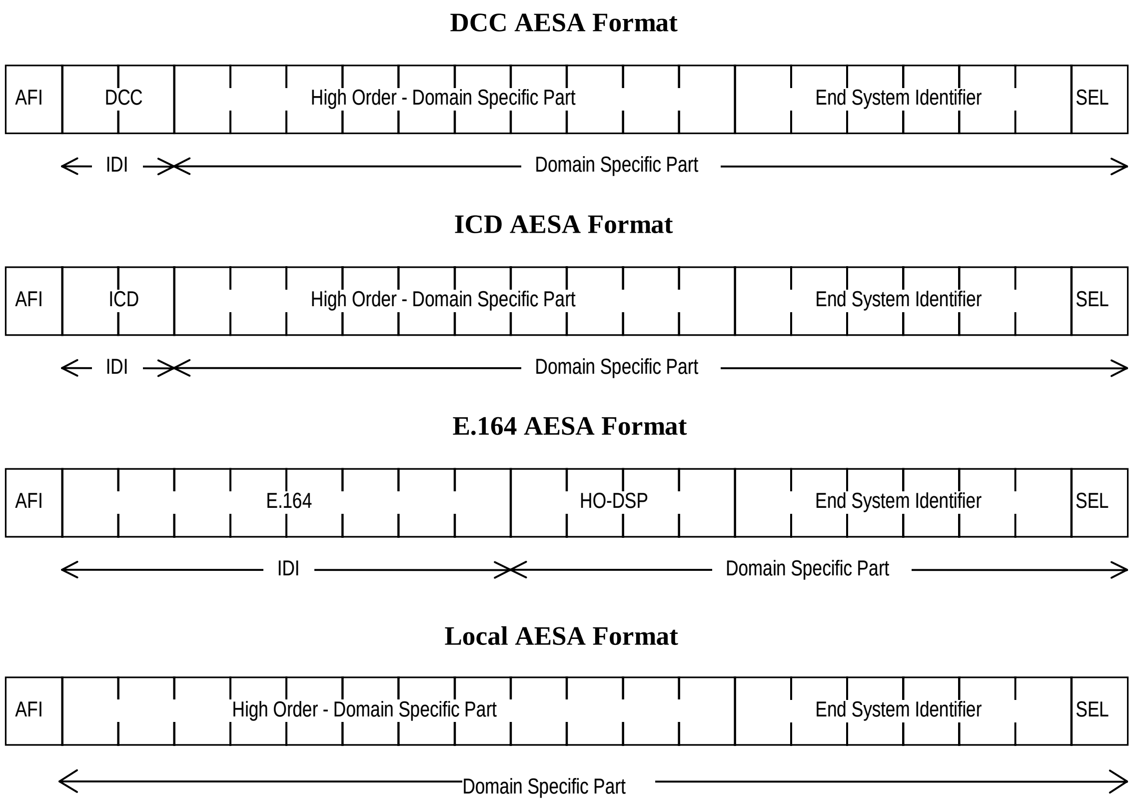

Figure 1 ATM Forum Defined ATM End System Address Formats

The fields defined by the ATM Forum specifications are:

- AFI: Authority and Format Indicator. The value of this field determines the type of the AESA (e.g., DCC, ICD, E.164, Local) and also indicates what authority can assign codes (the structure and values of fields) in the rest of the AESA. The AFI also indicates what encoding (binary or packed decimal) is used in the Domain Specific Part (DSP). The ATM Forum only uses AFIs that indicate binary encoding. The following four types appear in ATM Forum specifications:
  - DCC: Data Country Code. This type of AESA is assigned by ISO to national authorities (each country has a unique DCC code value). Each country is free to decide the structure and rules for assignment of the Domain Specific Part. The AFI for a DCC is 39 and the value of the DCC field indicates the country.
  - ICD: International Code Designator. An ICD is intended for use in the construction of internationally recognized codes (that is coded systems such as the merchandise bar code system, Dewey Decimal System, etc., not addresses). Each ICD code point assignee is free to decide the structure and rules used for assignment of the Domain Specific Part. The AFI for an ICD is 47 and the value of the ICD field indicates to which code set or organization that particular ICD is assigned.

- E.164: An E.164 AESA can be constructed by populating the Initial Domain Identifier of the AESA with a valid E.164 address (note that an *address* is explicitly required; numbers whose function is to identify services, are not acceptable). This AESA has a special case: when the HO-DSP, ESI, and SEL are all encoded as zero, the ATM Forum specifications recognize this as equivalent to a "native" E.164 address, and refer to it as an *embedded* E.164 AESA. The AFI for an E.164 AESA is 45.
- Local: The Local AFI, designated by AFI = 49, defines a structure that can be used by anyone within a private network. There is no Initial Domain Identifier, the octets following the AFI are all in the Domain Specific Part and can be structured by the user. In this case, private means just that: not interconnected with any other network, ASP or private.

- IDI: Initial Domain Identifier. The contents of this field vary depending on the value of the AFI. For example, with a DCC AESA (AFI=`39`), the IDI value of `0x840F` identifies the United States[2].
- HO-DSP: High-Order Domain Specific Part. This field has meaning as defined by the address authority controlling the AESA and its delegates (see discussion below). This component (together with the AFI and IDI) is typically used within the network to route a call to the appropriate switch.
- ESI: End-System Identifier. The ESI is usually an IEEE 802.2 Media Access Control (MAC) address. In this case the value of this field is filled in by each end system with codes derived from IEEE assignments and that (more or less) uniquely identify each interface. The code is typically burned into a ROM on the interface, by its manufacturer. [See www.ieee.org for assignments and structure.] The ESI is typically used by the terminating switch to select the interface to which the end system is attached.
- SEL: Selector. The selector is not used for ATM routing, but may be used by end systems. This field can be used by the end system for internal purposes (typically to identify the particular internal module that is to handle incoming calls, much like TCP/IP port numbers or an ISO/OSI network service access point). It can also be used to differentiate multiple addresses associated with the same interface.

E.164 numbers are generally used to identify an interface between an ASP network and a private ATM network. Different forms of E.164 AESAs use the HO-DSP, ESI and SEL components differently, however.

- In embedded E.164 AESAs, the HO-DSP, ESI and SEL components are not used (the corresponding fields must be set to zero).
- In E.164 AESAs other than embedded E.164 AESAs, the HO-DSP, ESI, and SEL may be administered by the private organization for use within the private ATM network. In this case, the E.164 number identifies the private organization as the authority for the assignment of the HO-DSP, ESI, and SEL.

The encoding of information into the AFI and Initial Domain Identifier fields of an AESA are determined by rules in ISO 8348, Addendum 2 [ISO NSAP]. See the *ATM Forum Addressing: Reference Guide* [ATMF ARG] for more details.

Assignments of AESAs generally follow a series of authorities, the top level assigning some part of the prefix and delegating authority to a lower level authority for the rest. ISO and the ITU-T jointly assign the AFIs, the DCC country codes, and the ICD values, and then assume that the particular country, organization, or region has a mechanism to handle requests within its scope of authority. Thus the American National Standards Institute, ANSI, handles assignment of codes within the US DCC address space, the Federation of the Electronics Industry, FEI, handles requests within the United Kingdom, Deutsche Industrie-Normen, DIN, handles requests within Germany, etc., These authorities, in turn, may assign a sub-authority. Each sub-assignment may, or may not, define some further structure in the HO-DSP of the AESA. These definitions are done from left to right to form a "prefix" of the AESA.

Thus there is a hierarchy of registration authorities involved in assigning an AESA. At each level in the hierarchy there is a *registration point* beyond which the sub-authority may define the structure of the remaining bits/octets. The root of the hierarchy is the joint ISO/ITU-T registrar, and the registration point is defined immediately after the

---

[2] The notation "`0x`" indicates a hexadecimal (base 16) number.

Initial Domain Part (AFI and Initial Domain Identifier). Subordinate authorities can further define sub-authorities and sub-registration points.

For example, ANSI is the authority for the DCC with country code = 840 (USA) with registration point after octet 3. [ANSI NSAP] defines a structure (see Figure 2) with sub-authorities as registered organizations with the sub-registration point after octet 7. The first octet after the AFI and DCC code for the USA is assigned as a Domain Specific Part Format Indicator, DFI, a field that allows subsequent fields to be formatted differently depending on the value of the DFI. The DFI currently has a single value of 128 (0x80). The three octets after the DFI are allocated to a field called the “Organization name”, a 24-bit value which is assigned by ANSI to organizations. The remaining fields of the HO-DSP (that are defined in [ANSI NSAP]) are not interpreted with ATM addressing, that is, these fields called Routing Domain, Area, and Reserved, are used within the ATM address without regard for the definitions in [ANSI NSAP].

AFI	DCC	DFI	Organization Name	Specified by Organization	ESI	SEL
0x39	0x840F	0x80				

Figure 2 Structure of the HO-DSP within the US DCC AESA

A prefix of an AESA is any bit-string that is from 0 to 152 bits (19 octets) in length. The Selector field cannot be part of a prefix.

2.3 ILMI Address Registration

ILMI Address Registration is a protocol exchange between a switch and an end system that allows for auto-configuration of the end system and switch.

As shown in Figure 1 and Figure 2, all of the formats have a 13 octet structure followed by a 6 octet End System Identifier (ESI) and a 1 octet Selector (SEL). Integrated Local Management Interface (ILMI) address registration uses this 13-6-1 octet structure.

A switch uses the ILMI address registration protocol to provide a 13-octet prefix to the end system [ATMF ILMI]. Conceptually, the end system obtains its End System Identifier (ESI) from a ROM on its network interface card (NIC) and concatenates that, and a zero octet for SEL, with the prefix provided by the switch. The resulting 20 octet address is then returned to the switch using the ILMI address registration protocol. The switch uses this address (usually after validating the prefix of the returned address) to identify the UNI to deliver calls to that address.

2.4 Multiple Addresses and Addressable Entities

It is entirely possible to associate multiple addresses or address prefixes with a given interface. This association may actually locate the interface, in the case of an address, or a route toward a destination/location, in the case of a prefix. In this sense an end system may be capable of receiving calls at a number of different addresses. The situation where multiple addresses at a single interface is most common, however, is at a public UNI. When the public UNI is used to interconnect an ASP network with a private network, it is, in fact, being used as an NNI. In such circumstances, there may be many addressable interfaces on systems, and hence addresses, reachable via that UNI. The set of addresses may be represented by an enumeration of addresses, an address prefix, a set of address prefixes, or some combination of addresses and address prefixes.

When used as an NNI, the addresses and address prefixes associated with a UNI should be thought of as routes toward the destinations/locations summarized and not as identifiers that distinguish interfaces.

When a single address or address prefix appears at multiple interfaces in the global ATM network, it usually designates a route toward the destination(s) summarized by the address or prefix. If a full 20 octet address appears multiple times, then in one instance it is used as an address to identify the interface, and in all the others it is used as a route toward that interface.

2.5 Categorization of Addresses by Ownership

AESA prefixes are differentiated in this document by who has obtained a particular prefix, and for what purpose.[3]

- An ***ATM Service Provider prefix*** is a prefix allocated to an ATM service provider by a national or world registration authority. Prefixes that are extensions of such prefixes are also ATM service provider prefixes. Addresses derived from such prefixes are ATM service provider addresses. An ASP may sub-allocate part of its address space to its customers.
- A ***Customer-Owned ATM prefix*** is a prefix allocated directly to a private network by a national or world registration authority. Prefixes that are extensions of such prefixes are also customer-owned prefixes. Addresses derived from such prefixes are customer-owned addresses.
- An ***Unregistered ATM prefix*** is a prefix that is not obtained from a national or world registration authority, or an extension of such a prefix. Private ATM networks may use unregistered prefixes to derive unregistered addresses. These may only be used within that private network. See Section 4.3.

  If a network wishes to use unregistered addresses, it is recommended that they be formed using the local AFI (`49`), which is defined explicitly for this purpose.

  Unregistered ATM Prefixes can also be formed through the use of other address formats which are not registered to the network in which they are deployed. While this may work (for a time), it presents a number of serious issues and is strongly discouraged.

Implications of using customer owned vs. service provider addresses, are discussed in sections 4 and 5.

2.6 Service Provider Addressing Support

The use of various address types (e.g., E.164, the AESA based on Data Country Code) is allowed by ATM Forum specifications. This means that the protocols provide the means for an ATM service provider to construct a network that uses any or all of them. This does not imply, however, that an ATM service provider is required to support all of them.

An ATM service provider that offers SVC services to its subscribers will typically also offer addressing services in conjunction with it. Addressing services are expected to include the assignment of addresses from an address space administered by the ASP and *may* include the capability for the subscriber to provide to the ASP addresses and/or address prefixes that the subscriber has been assigned by a recognized address registration authority.

All aspects of call establishment may be subject to restriction by an ASP with the concurrence of the calling and called parties. For example, call requests to destinations on the same ASP as the originating customer, but to addresses other than those assigned to that customer may be blocked or not. Some customers may request that other customers not be allowed to call them. Since not all ASPs are interconnected, a customer of one ASP may or may not be able to request a call setup to a customer of another ASP (customers should investigate the status of reachable parties during subscription).

2.7 An Analogy with IP Internetworks

IP addresses are assigned by the Internet Assigned Numbers Authority (IANA). Each IP version 4 address is defined to be 32 bits. With the implementation of Classless InterDomain Routing (CIDR) in the Internet [RFC 1519], prefixes of IP addresses are assigned to organizations. The prefix is denoted in the usual dotted-decimal form with the assigned part filled in, the unassigned part filled with zeros, and a "/" followed by the prefix length, e.g., `12.0.0.0/8`.

The IANA today assigns address space to IP Network Service Providers (NSPs) who assign some part of that space to downstream Internet Service Providers (ISPs), who, in turn, assign part of that to organizations that attach to them. The intent is to ensure that addresses are assigned so that their location function is maximized, that is, tied to

[3] "Ownership" in this usage means the organization is allowed to use the prefix or addresses, not that it is free to reassign it to someone else, or dispose of it in any way not agreed to by the controlling registration authority. The organization does not *own* it, but is assigned it by the true owner: ISO or ITU-T and their delegates.

the backbone topology. In today's Internet, only the very largest organizations may obtain addresses directly from the IANA.

When an organization receives an assignment, it becomes the assigning authority of the remainder. This uses a hierarchy of assignment authorities. AESA assignments followed this tradition.

Most Internet addresses are restricted today to Service Provider assignments. The reason for this is partially historical and partially due to the phenomenal growth rate of the Internet. The historical reasons have more to do with the fact that the Internet did not have an address assignment plan initially. This resulted in what is known today as the "toxic waste dump" (addresses in the `192.0.0.0/24` range). These prefixes were allocated without regard to where in the Internet topology the networks using them fell (and indeed, the Internet backbone topology has changed as well). As a consequence these address prefixes are more identifier than locator, and each one of them appear individually in the routing tables of every Internet backbone router—they cannot be summarized. There are some `6400` routes, or approximately 12% of the Internet routing table in backbone routers at the major exchange points (as of March 1998) that are in the "toxic waste dump".

2.8 Address Aggregation and Routing Systems

Address summarization/ aggregation is one mechanism that is used to allow networks to scale. If all the addresses assigned in a locality of the network share a common prefix, then that prefix can be used in other parts of the network instead of enumerating each local address, much as area codes provide some level of aggregation of telephone numbers.

An address aggregate is a set of addresses that are derived from a common prefix **and** a common locality within the network (e.g., the interfaces on a switch or cluster of switches). If addresses from a common prefix **do not** have a common locality, then they cannot be aggregated. The reason is that the prefix representing the aggregate is used as a common route toward all the addresses in that aggregate (there are ways to provide for exceptions). If they don't share a common locality within the network, then it doesn't make much sense for them to share a common route.

Thus one goal of address plan design is to consistently assign addresses in a locality of the network topology from a common address prefix. If an address or a smaller aggregate (longer prefix) is assigned outside the locality of the larger aggregate, then it is called an *exception* and requires a separate route entry in the routing table to identify it throughout the network.

2.8.1 Address Aggregation Examples

2.8.1.1 Campus Network

Assume the network administrator has assigned a prefix of `39.840F.80.809134.0001` to a specific campus/site network. The campus network administrator can refine that prefix by extending it by two fields, each of length one octet[4], one to enumerate each building and one to enumerate the floors of each building (see Figure 3).

The addresses within the campus aggregate nicely into the single site prefix.

---

[4] Each "digit" in the address, is encoded using 4 bits. Thus the 2-digit fields in the example each occupy one octet of the 20-octet address.

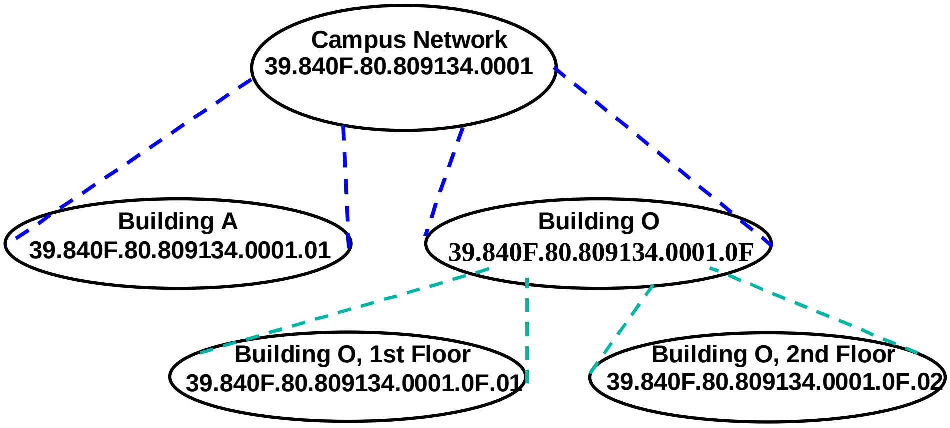

Figure 3 Campus Network with Address Aggregation

2.8.1.2 Wide Area Enterprise Network Using Customer-Owned Addresses

Suppose that a USA-based enterprise obtained an address prefix from ANSI in the US DCC space and proceeded to extend that prefix by one octet, creating a field it used to enumerate its major sites. Suppose it has major sites in San Francisco, Chicago, and Dallas. Each campus constructs its numbering plan as in the previous section. The enterprise has constructed its address plan very well, and it allows aggregation of its address space within its own network.

Now suppose the company connects the ATM networks at each site into an ATM Service Provider network. Since the enterprise did not obtain addresses from the ASP, its addresses bear no relationship to the topology of the ASP network and the ASP's addressing plan. For each interface in the ASP connected to the enterprise, the ASP must support at least one exception route everywhere in its network. If all the ASP's customers do this, the number of routes grows as the number of interfaces in its ATM network (and thus among all ASPs willing to carry that route within the global ATM network).

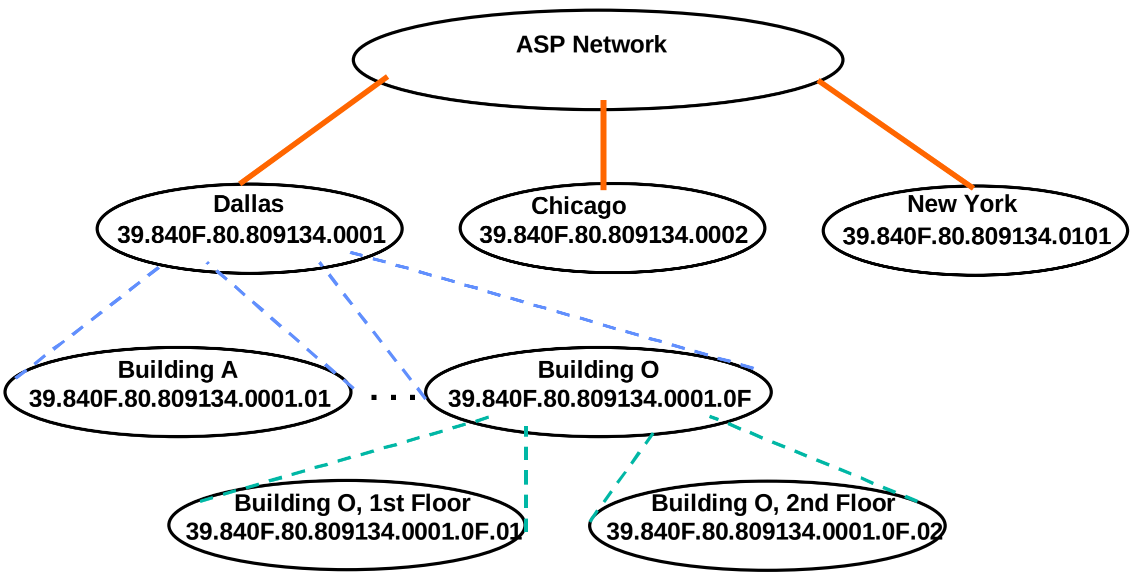

Figure 4 Hierarchical Network Structure without Address Aggregation

The Internet used to assign addresses like this, and Internet service providers used to have to install routes in their routers based on addresses their customers obtained independently. The Internet has undergone a significant change in the past few years with respect to IP address assignment. This change has been driven by a scarcity of IP addresses (eased by CIDR) and the exploding size of the global Internet routing table (it grew as described above: one or more entries per attachment to ISPs across the globe). Internet service providers have now implemented rules that dictate when a customer's address space can be routed within the global Internet in an attempt to lower the rate of routing table growth. This depends heavily on the consistent assignment of IP address prefixes by ISPs to their customers and on the aggregation of these address prefixes within the global Internet.

2.8.1.3 Wide Area Enterprise Network Using ATM Service Provider Addresses

The example given in Figure 5 shows an ASP network with a corporate customer with three major sites: London Main, London South, and Berlin. The ASP network-wide prefix is "`47.0079`". The corporation has obtained 3 prefixes from its ASP, each of length 72 bits (9 octets), which are: "`47.0079.8000.1122.0001`",

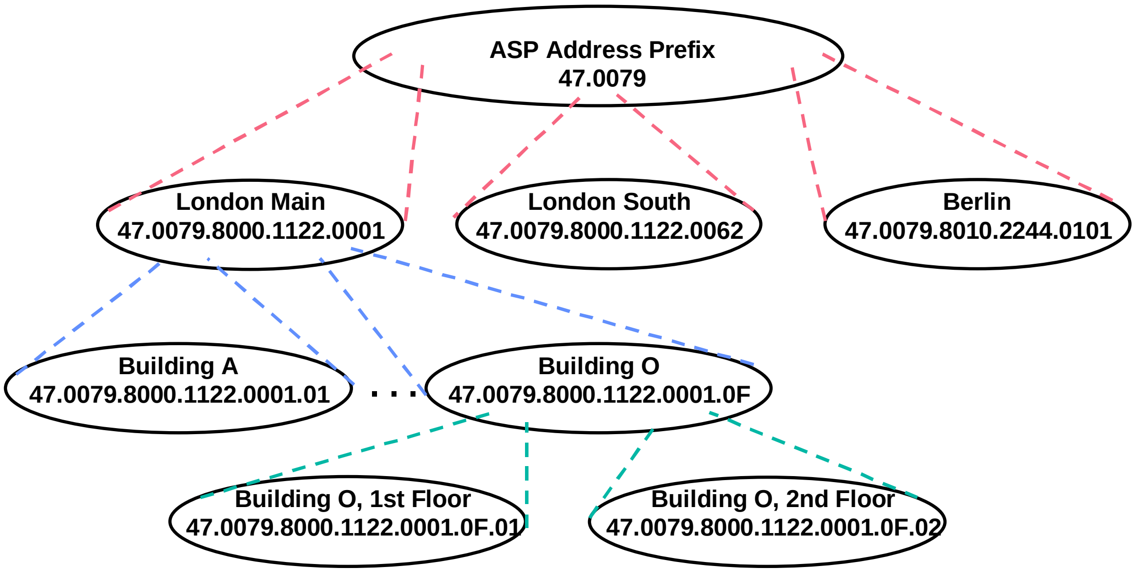

Figure 5 Hierarchical Network Structure with Address Aggregation

"`47.0079.8000.1122.0062`", and "`47.0079.8010.2244.0101`", respectively. It has decided to use the next 8 bits to designate buildings within the site, and 8 additional bits to designate the floor within that building. So that in London Main, Building A is represented by the extension "`01`" and Building O by extension "`0F`".

Within this scheme, particular switches would be numbered by extending the floor prefix.

2.8.2 AESA Prefixes and PNNI

AESA prefixes are most commonly used in the context of the ATM forum Private Network – Network Interface (PNNI). PNNI specifies that switches can be arranged in a hierarchical topology and that each switch is assigned a node prefix based on its place in the hierarchy. This is somewhat analogous to the use of subnets within an IP address block, or the use of area or city codes and exchange codes in telephone numbers. Switches at one level in the hierarchy derive their addresses by extending the prefix of their parent node by some number of bits sufficient to allow a numbering of all the nodes at that level.

This numbering assumes that the addresses of all end systems will be derived from the node prefix of the switch to which that end system is attached, and that the higher level peer group prefixes aggregate all the lower level ones.

2.8.3 ATM Name System

The ATM Forum has specified a protocol, the ATM Name System (ANS), that can be used to translate ATM system names to AESAs [ATMF NS]. The ATM Name System is based on the Internet's Domain Name System protocol and derives its name structure similarly. The specification provides the details of the information kept for each ATM end system, how the end system contacts the ATM Name System server, and the form of the queries and replies. ATM names use the same names as in the Domain Name System, but new resource records are defined to contain the ATM addresses associated with those names.

Network administrators who maintain a large ATM enterprise network may wish to consider the use of the ATM name system to reduce the amount of work (and inaccuracies) involved in manually typing in 20 octet AESAs. In some cases, the AESA may not be portable among ASPs (as noted above, portability of addresses has serious consequences for the global ATM network routing system). However, names are almost always portable. For this reason, ATM end system names should be the items exchanged among users, not AESAs.

The ATM Name System is not widely deployed. While it is based on the DNS, the new resource record types containing the ATM End System Addresses are not recognized by most DNS server software, nor is there easily available client software that performs a query for that information. There are experimental implementations of the ANS [EPFL ANS].

2.9 Routing and Call Delivery

Routing in this document is used to mean either:

1. The procedures whereby a (source or intermediate) node, using the call set-up destination address (called party number), determines where the call set-up should be forwarded next. In PNNI networks this process determines not only the next node, but all nodes from the source to the destination; or

2. The PNNI routing information distribution protocol which allows a node to advertise the prefixes to which it has direct reachability. This process also advertises the status (and the resource allocation) of all directly attached links. This information is distributed to all nodes in the network.

PNNI routing information distribution allows each node to construct a routing table that is used to determine the path from a source node to the egress node serving a given destination (as defined by an AESA prefix). Prefixes of varying lengths, depending on how much is necessary to distinguish an appropriate route, are kept in the routing table.

The route from the source to the destination is based on the called party number's 19 octet prefix in conjunction with the routing table.

2.9.1 Longest Prefix Matching

Routing within the AESA address space is based on matching the called party number against the entries (prefixes) in the routing table. There may be several entries that begin with the same few octets. For example, all US DCC AESAs begin with the same 4 octets; all US DCC AESA prefixes within the same organization begin with the same 7 octets (where the 8th octet might be used to distinguish one site from another). Thus each entry in the routing table has a prefix and its length (in bits) associated with a source route (if using PNNI) or an out-bound trunk (possibly more than one with a metric allowing selection of the "best").

So the 19 octet prefix of the called party address might match, to some number of bits in length, many of the prefixes in the routing table. The best match is the prefix in the routing table that matches the 19 octet prefix of the called party to the largest number of bits; the longest prefix. The source route or out-bound trunk of the best match determines where the call should be routed.

2.9.2 Address Prefix Aggregation

Longest prefix matching allows for *address aggregation*. Because an ASP prefix is usually structured geographically in its high order parts (by country, region within country, etc.), routing tables for nodes in New York do not need to know all the address prefixes of nodes in California. They need to know only that those nodes can be reached by traversing one or more of the trunks directly attached. Thus all the 13 octet prefixes (104 bits) of UNIs and nodes in California can be represented as a single, shorter, prefix (say 72 bits long, for example). This means that routing tables can be quite small and that scaling properties are good. Note that the aggregation of addresses typically is not automatic, but done by administration during the network planning phase.

When customer *networks* are attached, however, there are problems. Customers use addresses that are not necessarily part of the ASP's numbering plan (for a variety of reasons). These addresses are not necessarily distributed across the ASP network in any fashion that might allow address prefix aggregation (i.e., the customer's high order address might be structured by organization within the enterprise, e.g., finance, engineering, sales). The numbering plan must align with network topology to allow aggregation to occur, and we have no guarantee that a particular customer has numbered their network to align with ASP's topology. So aggregation may not be possible.

In fact most customers with locations in many parts of the US (or the world) will likely have numbering plans that use geography, at least roughly, in its structure. ASPs may be able to aggregate to some degree using these addresses. However, each point at which the customer attaches to the ASP network typically results in one or more exception routes. The ASP's routing table size grows linearly with the number of attachments using customer-owned addressing. Such growth inhibits scaling to a global network.

2.9.3 Distribution of Addressing Information

Address prefixes are distributed using a routing protocol, e.g., PNNI, that allows a node to pass to each of its neighbors' information about what it has directly connected to it and what it has learned from its other neighbors. This is called "flooding". In a connected network, all nodes are reachable, via some concatenation of neighbors, from all others. This information from one node is passed, hop by hop, to all the others. Aggregation occurs along the way, so not all information is passed; as information is aggregated, detail is lost, but reachability is not.

3 Example Network Topologies

This document discusses only the scenarios in which the Private Network is based on PNNI (or other network based on ATM End System Addresses) and their interworking with ATM Service Provider networks based on ATM End System Addresses or E.164 addresses. The assumption is that if there are multiple ASPs involved, they all support connection setup based on the same address categories (i.e., all AESA or all E.164, but not a mix of AESAs and E.164 unless all ASPs support both categories).

This document specifically does not cover topologies in which multiple interconnected ASPs do not support the same addressing and routing capabilities (e.g., mixtures of E.164-only and AESA-only ASPs).

It is a goal of this section to describe end-to-end examples. That is, each example will begin by assuming an end system, given an ATM name supported by the ATM Name System server, looks up that name, retrieves one or more addresses, and uses those addresses to form a connection setup request. The network (private or combination of private and ASP) receives the request and routes the call request to the destination end system. The request is then accepted and the acknowledgement is returned to the calling party.

3.1 Scenario 1: Stand Alone Private ATM Network

This scenario assumes only a stand-alone private network with no ASP network involved. The ATM Forum specifications limit the addresses used at private UNIs to be AESAs (that is, native E.164 addresses are not allowed).

Assume that all end systems are assigned a name and that they obtain their addresses using the ILMI protocol (and thus are derived from the prefix of the switch to which they are attached). The end system also learns which ATM Name System server it should use via the ILMI protocol (this is supported in ILMI 4.0—in earlier versions of ILMI, this information must be manually configured).

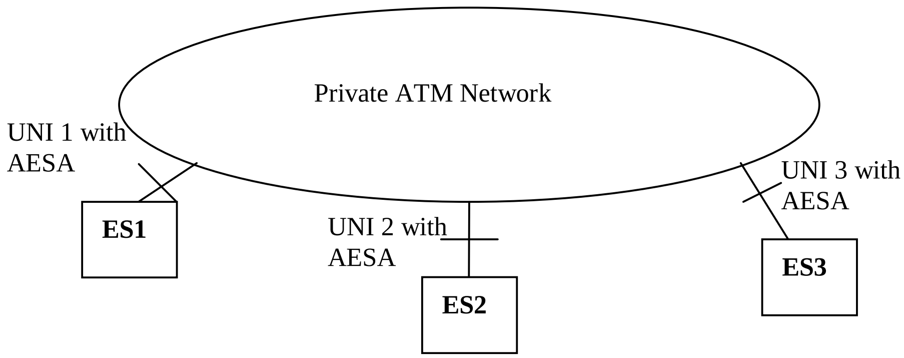

Figure 6 Private Network

3.1.1 Calling Party Initiation

The end system receives a request to establish a call to an ATM name. This request could be generated by the user of an application program running on the end system or generated internally by an automated process.

Assume that the end system has already established a connection to its designated ATM name server. Using that connection, it sends a name query and receives the ATM address corresponding to that name (or an error).

The end system then forms a call SETUP[5] message using the address of the named destination in the Called Party Number and its own address in the Calling Party Number. The original request also should have contained the traffic parameters of the call (traffic category, etc.). The end system uses this information to complete the mandatory and optional elements of the SETUP message.

The end system transmits the SETUP to its attached switch on the UNI signaling channel, which is used for all call control messages between the end system and the switch.

The network routes the call to its destination. Meanwhile, the network may respond with a CALL PROCEEDING message to let the end system know that the SETUP was received (this is optional).

3.1.2 Network Routing

The network routes the call setup message based on the called party number. Each end system communicates its connection control messages to a switch over the signaling channel. Similarly, when two switches are interconnected using PNNI, they have at least one signaling channel on the PNNI interface and use that to communicate call control information (there is also a routing control channel used to exchange routing information).

The set of signaling channels interconnecting the switches and end systems comprises the signaling network. The signal processing units in the switch serve multiple functions, one of which is to route messages to the proper destination. In the case of a call setup, the destination is identified by the Called Party Number.

The switch maintains a routing table (based on PNNI routing exchanges) that allow it to decide which of several possible out-going paths to use to progress the call toward the destination. The switch chooses a complete node/link path from itself to the "nearest" PNNI node advertising the best match for the Called Party address (directly or through an external interface). Here "best" means the node advertising the longest prefix that matches the called party (e.g., if the called party is "`12345`", node A advertises "`123`", node B advertises "`1235`", and node C advertises "`1234`", then node C has the longest prefix match). Once the "egress" node is identified, the "shortest path" from the ingress to the egress which meets the requested Quality of Service (QoS) constraints, is selected.

---

[5] Words in all upper case are names of control messages in the ATM signaling protocol.

As the call progresses through the network, *call reference values* are created. These values allow the signal processing units to later distinguish this call from others between the same calling and called end systems and also to be able to reverse the path from the called end system to the calling end system.

Once the SETUP message is routed through the network, it is transmitted across the UNI (again on the signaling channel) and so *offered/delivered* to the called party.

3.1.3 Called Party Acceptance and Connect

The called end system receives the SETUP message and can determine in any way it likes whether or not to accept the offer of connection establishment. If the called end system decides it will accept, it sends a CONNECT message back using the call reference value received in the SETUP (which also contained the VPI/VCI to be used for the connection, if accepted).

If the called end system does not want to accept the call, it responds with a RELEASE COMPLETE message that may contain a cause code indicating why the call was rejected.

3.1.4 Calling Party Connect

After the destination has indicated acceptance of the call by returning a CONNECT message, the network reverses the path of the connection using the call reference values. Eventually it responds to the originating end system with a CONNECT message containing the VPI/VCI of the newly established connection along with its traffic parameters (which may be different than originally requested due to negotiation). If the network could not complete the call or the destination rejected the call, then a RELEASE message is returned with cause codes.

The end system responds to the CONNECT message with a CONNECT ACKNOWLEDGE (and to a RELEASE message with RELEASE COMPLETE). Once the connection is acknowledged, the end system may begin sending data on the connection.

The end system can subsequently initiate connection clearing by sending a RELEASE message on the signaling channel.

3.2 Scenario 2: Private ATM Network Interconnected by an ATM Service Provider Network with AESAs

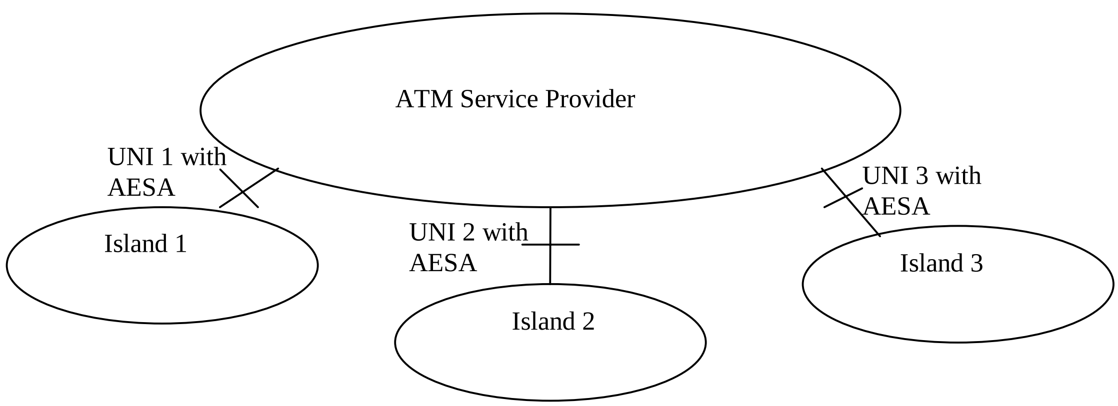

Figure 7 ASP with AESA Addressing

In this scenario we describe two (or more) PNNI islands of the same enterprise connected across an ASP that provides call setup using AESAs, see Figure 7. Each PNNI island is attached to the ASP using a public UNI. Therefore, PNNI is used within each island and within the ASP, and does not cross the UNI. Routing information is not passed dynamically across the UNI as in PNNI. There are 4 separate PNNI networks in Figure 7.

In this example we assume that the addressing used within the “islands” is based on prefixes obtained from the ASP. We also assume that the call setup is requested from an end system in one island to an end system in another island.

Call setup is the same as in scenario 1 with the exception that it proceeds on two levels: one within the private network and one within the ATM Service Provider. Within each island, a call setup is routed between the end system and the public UNI attaching the island to the ASP.

Within the ASP, the call setup is handled as in scenario 1. The major difference is that ILMI is not used to configure addresses at the public UNI. Instead, the prefix (address aggregate) of the particular island is configured as a static external route on the ASP side of the UNI. PNNI passes this route to other ATM nodes in the ASP network so that a call to called party numbers within the set summarized by a configured prefix, can be routed.

Within each of the islands, a “default route” (prefix with zero length), or specific prefixes for all the other islands, are configured as external static routes on the private network side of the public UNI, pointing toward the ASP. PNNI passes these routes to other ATM nodes within that island.

In this scenario, because ATM service provider addresses were used (see the definition of this term in section 2.5) the called and calling party addresses remain constant across both of the private network islands and the ASP network.

3.3 Scenario 3: Private ATM Network Using E.164 AESA Interconnected by an ATM Service Provider Network

In this scenario, we presume two (or more) PNNI islands of the same enterprise connected across an ASP. The PNNI islands are addressed using E.164 AESAs. The ASP can provide call establishment using E.164 AESAs in the called party. The E.164 AESA was returned by the ANS.

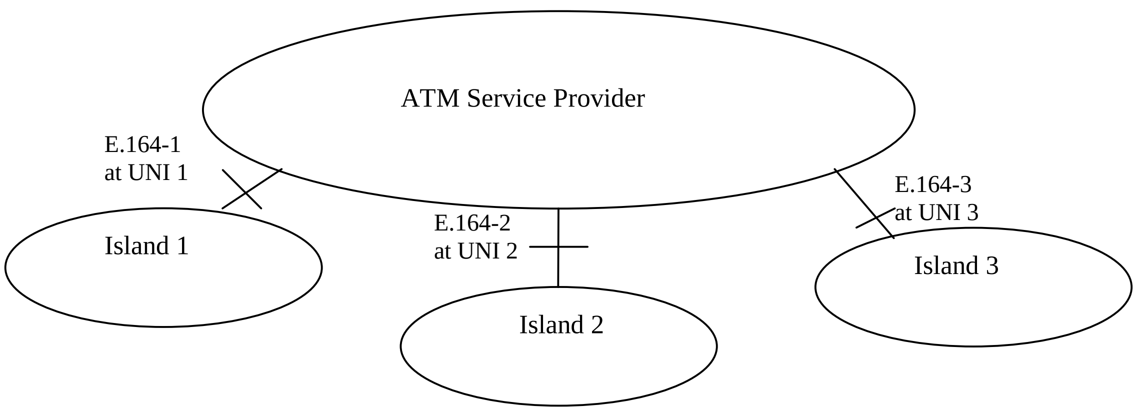

Figure 8 ASP with E.164 Addressing Private Network with E.164 AESA

Within each island, the end system is addressed using an E.164 AESA. The E.164 AESA is derived from the E.164 address the ASP assigns to the Public UNI between the ASP and each island. The ASP assigns E.164 numbers to its UNIs, and the E.164 numbers are to be used in the format of the E.164 AESA.

Within each island, node prefixes are obtained by structuring the HO-DSP octets of the E.164 AESA format shown in Figure 1.

3.3.1 Calling Party Initiation

When the end system queries the ATM name server to look up the destination address, it receives an E.164 AESA in the response. It places the E.164 AESA in the Called Party Number Information Element (IE) as in scenarios 1 and 2 and passes the SETUP message to its attached switch.

The end system also places its E.164 AESA in the Calling Party Number IE.

3.3.2 Private Network Routing to ASP Interconnection

The attached switch recognizes the E.164 AESA as it does any other AESA, by longest prefix match. As in the first scenario, a default route or set of specific routes pointing toward the ASP are advertised by the private network side of the public UNI between the ASP and the island, and propagated via PNNI throughout the island.

So the call is delivered across the public UNI to the ASP.

3.3.3 ASP Routing to Private Network Interconnection

The ASP routes the call based on the Called Party Number. Routing within the ASP network is as it was in scenario two with the exception that the addresses involved are all E.164 AESAs.

3.3.4 Private Network Routing to Called Party

The call is delivered to the first switch in the destination island across the associated public UNI.

Call routing is based on the Called Party Number, an E.164 AESA, which is appropriate for use within an island using a numbering plan based on AESAs. Routing occurs as in scenario one.

3.3.5 Called Party Acceptance and Connect

The destination end system again may use any criteria it wishes on which to base a decision to accept or reject a call.

3.4 Scenario 4: Private ATM Network Using Customer-Owned Addresses Interconnected by an ATM Service Provider Network with E.164 Addressing at the ASP Border

In this scenario, we presume two (or more) PNNI islands of the same enterprise connected across an ASP. The ASP can only provide call establishment using native E.164 addresses as the called party.

Within each island, the end system is addressed using an AESA. Because the ASP does not allow the use of AESAs through its network, but rather uses the ITU-T standard E.164 numbering plan at its UNIs, an interworking function is needed. Note that this behavior is not specified in any ATM Forum specification and that its description here is informational. It is one way to achieve call setup in this scenario. The interworking function can be implemented on a gateway switch on the enterprise side of the UNI or on the ASP side of the UNI. Here we describe the interworking function as being on the enterprise side. Other implementations are possible.

3.4.1 Calling Party Initiation

When the end system queries the ATM name server to look up the destination address, it receives the destination system's AESA. It places the AESA in the Called Party Number IE and passes the SETUP message to its attached switch.

The end system also places its AESA in the Calling Party Number.

3.4.2 Private Network Routing to ASP Interconnection

The attached switch routes the call using the Called Party Number. As in the second scenario, the private network side of the Public UNI between the ASP and the island is advertising either a default route or a set of specific routes, pointing toward the ASP.

The last switch in the island must perform an interworking function on the Called and Calling Party Numbers. In particular, it must move the Called and Calling Party Numbers, which are AESAs, into the Called and Calling Party Subaddress information elements and replace the Called and Calling Party Numbers with the E.164 addresses of the destination and local public UNIs, respectively. Thus the call is delivered across the public UNI to the ASP.

How the translation of Called Party Number to destination E.164 address is accomplished is outside the scope of ATM Forum specifications. Some switches allow the enterprise network administrator to build translation tables for this purpose. We note, however, that while this mechanism is straightforward when interconnecting the sites within a single enterprise, the maintenance of such translation tables for any-to-any call setup becomes infeasible. Thus this mechanism does not support an environment in which a user can place a call to an arbitrary end-point on a global public ATM network.

When the islands use E.164 AESA numbering plans based on the E.164 numbers assigned to the interfaces between the ASP and the islands (see section 5.1.3), rather than using customer-owned numbering plans (see section 5.1.2), the translation of Called Party Number to destination E.164 address may be performed algorithmically.

3.4.3 ASP Routing to Private Network Interconnection

The ASP routes the call based on the Called Party Number (and, in fact, carries the Called Party Subaddress transparently to its egress UNI). The Calling Party Number and Subaddress are also carried. Routing within the ASP PNNI network is as it was in scenario one with the exception that the addresses involved are all in E.164 format, rather than AESA format. The ASP neither examines nor acts upon the contents of the Called and Calling Party Subaddresses.

3.4.4 Private Network Routing to Called Party

The call is delivered to the first switch in the destination island across the associated public UNI. This switch recognizes that the Called Party number matches the E.164 address of its island and replaces it with the AESA in the Called Party Subaddress (and deletes the Called Party Subaddress).

Similarly, the Calling Party Subaddress is copied into the Calling Party Number and the Calling Party Subaddress is deleted. Routing the call is thereafter based on the Called Party Number, which is now an AESA, and appropriate for use within an island using a numbering plan based on AESAs. Routing occurs as in scenario one.

3.4.5 Called Party Acceptance and Connect

The destination end system again may use any criteria it wishes to determine whether to accept or reject a call.

4 Private Network Naming and Addressing

The decisions surrounding how to obtain addresses and what type of address to use are always difficult ones. If the private network will never be interconnected with an ATM Service Provider, then unregistered addresses can be used (e.g., much like using RFC1890 IP addresses). See Section 4.3 for details. If the private network will be connected to an ASP, either immediately or eventually, then registered addresses should be used.

Just as with IP networks, administrators may obtain registered addresses directly from one of the address registrars, or from their ATM Service Provider. The advantages and disadvantages of these courses of action are described here.

4.1 Using ATM Service Provider Addresses

As with Internet Service Providers, ATM Service Providers usually assign address space when a private network subscribes for service. The length of the assigned prefix may vary depending on the size of the private network, or it may be fixed. Either way, the prefix can be obtained and enables aggregation of addresses within the global ATM network (assuming the service provider is doing proper aggregation).

If the private network changes service providers, then the private network node prefixes must be re-administered, and the ATM Name System servers changed to reflect the change in addresses. Because the end systems learn their full addresses using ILMI address registration, the task of changing prefixes is easier than it was in IP networks before the advent of Dynamic Host Configuration Protocol (DHCP). Nonetheless, there may be many ATM switches in the private network, and each of them must be changed.

The advantage of using ASP-provided addresses is that there will be a much greater chance that the addresses will be supported throughout a global ATM network.

4.2 Using Customer-Owned Addresses

Private network administrators may also go directly to address registrars to obtain their own addresses. In this way, the private network need not change their prefixes as they change ASPs. Which type of address to use is typically not a technical issue, but rather one of address availability and the qualifications of the applicant.

- The British Standards Institute, BSI, controls the assignment of ICDs on behalf of ISO. BSI has assigned one ICD code for use with “Identifiers for Organisations for Telecommunications Addressing (IOTA)”. The IOTA scheme allows organizations to obtain identifiers that can be used to construct ATM addresses derived from this single ICD code. See Section 6.1 for more information.

  Before IOTA, entire ICD codes were assigned to organizations without strict adherence to the original purpose. Today BSI strictly limits the assignment of ICD codes. BSI has indicated that entire ICD codes will **not** be given to organizations to use as AESAs and has asked that such requests not be made. ICD codes that have already been assigned by BSI are valid for use and may be further subassigned by ATM Service Providers to their customers.

- Enterprise networks desiring their own address space can also obtain addresses within the DCC AESA space of their home country. Since each country uses different procedures to apply for address space, and the structure of the address space may vary from country to country, specific steps cannot be generally given. See Section 6 for more information about address registration in various countries and refer to [ATMF ARG].

Note that, as with Internet Service Providers, ASPs will not often accept a customer’s address prefix for routing within the ASP network (in the Internet, this depends mostly on how “large” the private network is; there is very little experience with this in ASP environments to date).

Customer-owned addresses are usually topologically significant within the customer network. When these networks are interconnected with ASP networks, however, the customer-owned addresses may not be topologically significant within the larger context. As discussed in Section 2.8, the use of customer-owned addresses in the ASP environment has a large impact on the global ATM network routing system.

4.3 Using Local (Unregistered) Addresses

If a private network wishes not to obtain registered addresses and not to interconnect with an ASP, there is a third way to obtain addresses. AFI 49 indicates a type of AESA that is always to be considered “local”. An address with this AFI will never be accepted by an ASP as a legitimate address, but it provides a way to obtain an address prefix without involving anyone else.

The AFI is encoded as with other AFIs in octet #1 (see Figure 1). The HO-DSP field extends from octet #2 through octet #13 (octets #14 through #20 are still used for the ESI and SEL fields). The private network may use this prefix in any way it sees fit. Beware, however, that the network will have to be renumbered in case it is interconnected with any other networks, private or ASP.

5 Interworking Private Networks over ATM Service Provider Networks

5.1 Customer-Owned Versus Provider Numbering Plans

In a combined ASP and private network environment there has always been tension over who “owns” the addresses numbering the networks. This dichotomy is usually characterized as customer versus provider numbering.

5.1.1 ATM Service Provider Numbering Plans

In an ATM Service Provider numbering plan, the service provider should assign the addresses to be used by the private network in a way that facilitates routing across the global interconnected network. This type of numbering best suits the ATM Service Providers since it allows the highest level of address summarization and hence minimizes the size of routing tables, overhead due to routing information distribution, etc.

The numbering plan within a single ATM service provider and all customers of that provider can be consistent and well organized. Inter-provider routing is based on the (few) address prefixes that are assigned to each ATM service provider. Routing is efficient. One well-known example of this type of address assignment is the telephone network (before the advent of local number portability).

The customer, on the other hand, may have problems with this. Since the ATM service provider assigns the addresses, if the customer decides to change service providers, this happens at the expense of renumbering all of the end systems and/or nodes in the customer's private network. Various auto-configuration protocols can be brought to bear on the end system reconfiguration issue, but internal node numbering must still be changed manually.

Another problem with basing the numbering of the private network on addresses obtained from the ATM service provider, occurs when a customer network simultaneously connects to multiple providers (e.g., for redundancy, or because provider A is preferred in area X and provider B is preferred in area Y). Which provider's numbering plan should be used as the basis for the customer's numbering plan? The customer will generally prefer to use one numbering plan within the private network. If only one of the service provider numbering plans is used, this typically means that calls will only be delivered via one of the providers (since routing and addressing are intertwined). If more than one provider numbering plan is used, the customer must logically partition the private network based on the different numbering plans. This will complicate management of the customer network (and likely cause other problems as well).

5.1.2 Customer Owned Numbering Plans

In customer-owned numbering plans, the customer obtains the addresses used to number the private network and informs the ATM service provider of the address prefixes used in the private network. To change providers requires no renumbering on the part of the customer.

The ATM service provider is faced with a number of problems. Foremost is that the prefixes assigned to customer's private networks cannot (in general) be summarized to reduce routing table size and routing information distribution traffic. Each customer has a prefix unrelated to another customer's prefix, so the ports on a given switch are numbered without regard to a common switch or topological reference. This means that routing tables at every switch in the network must maintain information on the address assignment of every port on every switch since address summarization cannot be done.

See Section 6 and [ATMF ARG] for information on obtaining addresses of a particular type.

5.1.3 Hybrid Numbering Plans

Hybrid numbering plans joining provider and customer-owned numbering have been formulated. The ATM E.164 AESA is one such structure. The AESA contains an E.164 number that is used to address the ASP-private network boundary and the rest of the AESA is used to provide private network numbering. In this sense, the ASP network number is used as the most-significant part of the private numbering plan prefix.

These schemes usually suffer from the same drawbacks as provider numbering plans in that they do not allow transparent migration from one provider to another and do not easily support connections to multiple providers. Any hybrid numbering plan uses parts of the provider and parts of the customer assigned address to create a new address type. If the customer assigned part is in the higher-order part of the prefix, then it really behaves like a customer-owned numbering plan. If the provider assigned part is in the higher-order part of the prefix, then it really behaves like an ATM Service Provider numbering plan.

6 Address Registration Information

6.1 International Code Designator

British Standards Institute, BSI, controls the assignment of ICDs on behalf of ISO. In general, entire ICDs (i.e., new values of the ICD field) are not assigned for the purposes of ATM addressing. See Section 4.2 for more details.

BSI has assigned one ICD code for use with "Identifiers for Organisations for Telecommunications Addressing (IOTA)". The IOTA scheme allows organizations to obtain identifiers that can be used to construct ATM addresses derived from this single ICD code. Figure 9 shows the structure of an ICD AESA constructed using an IOTA identifier (the IOTA identifier is called an "Organization ID" in the figure).

AFI	ICD	Organization ID	Specified by Organization	ESI	SEL
0x47	0x0124				

Figure 9 ICD AESA using IOTA Identifier

Information on the IOTA scheme is available from http://www.bsi.org.uk/disc/iota.html.

Applications for an IOTA must be made to:

IOTA
BSI-DISC
389 Chiswick High Road
LONDON W4 4AL
UK
Tel: +44 181 996 9000

An application form may be obtained from the BSI web page referenced above.

6.2 Data Country Code

6.2.1 Data Country Code Registrars

A useful list of national DCC authorities and contact information is maintained by the Federation of the Electronics Industry, FEI, the ISO country registrar for the UK. This information may be found at:

http://www.fei.org.uk/fei/dcc-nsap.htm

6.2.2 US Data Country Code

The DCC address space within the USA is administered as organization names, rather than addresses. The name is used to construct an address prefix as described later in this section. If an organization wishes to use US DCC address space, it must obtain an organization name from the American National Standards Institute, ANSI. This organization name is numeric and is actually a unique branch point (arc) within the OSI object-id tree subordinate to:

{joint-iso-ccitt(2) country(16) USA(840) organization(1)}

For those readers who are unfamiliar with ISO naming conventions, the important point is that the organization name is guaranteed to be unique within the global naming structure. This makes them useful for generating unique identifiers, such as X.500 distinguished names, X.400 Originator/Recipient names, and US DCC-based ATM addresses.

The assignment by ANSI of an organization name to an organization is permanent.

Once the organization has a numeric organization name, a US DCC AESA can be constructed by converting it to hexadecimal and inserting it in the Organization field. Figure 10 shows an example of a US DCC AESA when the numeric organization name is 113526[6]. The registered organization then becomes the registration authority for all addresses derived from further extension of that prefix.

AFI	DCC		DFI	Organization Name			Specified by Organization	ESI	SEL
39	84	0F	80	01	BB	76			

Figure 10 Organization Name used in US DCC

[6] `113526` decimal equals ((0 x 16^5) + (1 x 16^4) + (11 x 16^3) + (11 x 16^2) + (7 x 16^1) + (6 x 16^0)).

You may also request a specific alphanumeric name which can be associated with the numeric name (e.g., an X.500 organization name). There are rules concerning the assignment of the alphanumeric name (see [ANSI NSAP]). Note that alphanumeric names are not used to construct US DCC-based ATM addresses.

ANSI's telephone number is `+1 212 642 4900`. Request the "Procedures for Registering Organization Names in the United States of America Under the Joint-ISO-CCITT Arc". Do not ask for an AESA registration package as they will not understand what this is.

7 Summary

This document is intended for informational purposes only and does not amend, extend, or restrict any ATM Forum specification.

An overview of the current ATM addressing capabilities and practices are provided to aid private network operators in deciding how to design and operate their networks. This guide:

- Defines basic terms and concepts.
- Provides example scenarios of various addressing configurations in private and ATM Service Provider networks.
- Discusses the various addressing options available and how to obtain addresses to implement them.

8 References

[B-ICI v2.0] ATM Forum, "BISDN Inter Carrier Interface (B-ICI) Specification, version 2.0 (Integrated)," af-bici-0013.003, December 1995.

[B-ICI v2.1] ATM Forum, "Addendum to BISDN Inter Carrier Interface (B-ICI) Specification, v2.0 (B-ICI Specification v2.1)," af-bici-0068.000, November 1996.

[ILMI v4.0] ATM Forum, "Integrated Local Management Interface (ILMI) Specification," Version 4.0, af-ilmi-0065.000, July 1996.

[PNNI v1.0] ATM Forum, "Private Network-Network Interface Specification version 1.0 (PNNI 1.0)," af-pnni-0055.000, March 1996.

[SIG 4.0] ATM Forum, "ATM User-Network Interface (UNI) Signaling Specification Version 4.0," af-sig-0061.000, July 1996.

[ATMF NS] ATM Forum, "ATM Name System Specification Version 1.0," af-saa-0069.000, November 1996.

[ISO NSAP] ISO/IEC Information Processing Systems – Data Communications – Network Service Definition Addendum 2: Network Layer Addressing. International Standard 8348/Addendum 2, ISO/IEC JTC 1, Switzerland, 1998.

[ANSI NSAP] American National Standard for Information Processing Systems--, Data Communications--, "Structure and Semantics of the Domain Specific Part (DSP) of the OSI Network Service Access Point (NSAP) Address," ANSI X3.216-1992, July 6, 1992.

[RFC 1519] Classless Inter-Domain Routing (CIDR): an Address Assignment and Aggregation Strategy. V. Fuller, T. Li, J. Yu, & K. Varadhan. September 1993.

[RFC 1518] An Architecture for IP Address Allocation with CIDR. Y. Rekhter & T. Li. September 1993.

[RFC 1918] Address Allocation for Private Internets. Y. Rekhter, B. Moskowitz, D. Karrenberg, G. J. de Groot & E. Lear. February 1996.

[RFC 2008] Implications of Various Address Allocation Policies for Internet Routing. Y. Rekhter, T. Li. October 1996.

[RFC 2050] Internet Registry IP Allocation Guidelines. K. Hubbard, M. Kosters, D. Conrad, D. Karrenberg, J. Postel. November 1996. (BEST CURRENT PRACTICE)

[RFC 2101] IPv4 Address Behaviour Today. B. Carpenter, J. Crowcroft, Y. Rekhter. February 1997.

[EPFL ANS] Web site for the Laboratoire de Reseaux de Communication at the Ecole Polytechnique Federal de Lausanne,http://lrcwww.epfl.ch/linux-atm/. ATM Name System implementation for Linux.

[ATMF ARG] ATM Forum Addressing: Reference Guide

[ITU-T E.164] "Recommendation E.164: The international public telecommunication numbering plan", May 1997, ITU-T.

[ITU-T E.191] "Recommendation E.191: B-ISDN numbering and addressing", October 1996, ITU-T.