

TCP & UDP with Network-independent Endpoints (TUNE)

Status of this Memo

This document is an Internet Draft. Internet Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its Working Groups. Note that other groups may also distribute working documents as Internet Drafts.

Internet Drafts are draft documents valid for a maximum of six months. Internet Drafts may be updated, replaced, or rendered obsolete by other documents at any time. It is not appropriate to use Internet Drafts as reference material or to cite them other than as a ``working draft'' or ``work in progress.''

Please check the `1id-abstracts.txt` listing contained in the internet-drafts Shadow Directories on `nic.ddn.mil`, `nisc.nsf.net`, `nic.nordu.net`, `ftp.nisc.sri.com`, or `munari.oz.au` to learn the current status of any Internet Draft.

Abstract

This memo proposes a strategy for providing TCP and UDP services over multiple network services in a manner compatible with existing TCP/IP applications. By introducing network-independent endpoint identifiers for transport entities, TUNE makes it possible to change network addresses (and potentially network protocols) without impact to the application layer. While this paper describes the use of Connection-Less Network Protocol [1] as a possible long-term network service for the Internet, the overall approach differs from the "TCP & UDP with Bigger Addresses" [2] initiative due to the addition of network-independent endpoint identifiers. TUNE also provides network layer programming interfaces which are unchanged from those of the current Internet Protocol (IP) suite and hence alleviate the need for application code changes during initial deployment.

Comments should be sent to the author (jcurran@nic.near.net).

Acknowledgments

Sincere thanks to the members of the TUBA mailing list who had to endure these ideas at very early stage, and in particular, thanks to Ross Callon, Dave Piscitello, and Paul Tsuchiya for bearing with my endless questions.

Table of Contents

1	Introduction
2	TUNE Architecture
2.1	Endpoint Identifiers (EIDs)
2.2	EID Resolution
2.3	TUNE Header Format
3	TUNE Administration
3.1	Autoconfiguration
3.2	Portability
3.3	Mobility
4	Appendixes
5	Security Considerations
6	References
7	Author's Address

1.0 Introduction

Recent growth in the Internet has revealed several significant scaling problems in the current Internet routing and addressing architecture. It is generally acknowledged that continued deployment of non-hierarchical addresses will result in unacceptable growth to transit routing tables in the Internet. While the "Classless Inter-Domain Routing" [[3](#)] approach to address assignment and route aggregation should minimize routing table explosion for the immediate future, it is not positioned as a long-term solution to the routing, addressing, and scaling problems.

From a purely technical viewpoint, strict hierarchical addressing could be used to create a highly-scalable Internet routing architecture. Topologically-assigned addresses for network attachment points would minimize the need for routing information exchange. It would be necessary to use addresses of sufficient size to fully represent the network topology, but large hierarchical addressing plans have already been proposed for some protocol suites [4].

There are two minor administrative problems in using topologically-assigned addresses as a basis for long-term solution to the routing problem:

Topologically-assigned addresses cannot provide organizations with geographic and service provider mobility. As a result, organizations using topologically-assigned addresses are forced into address migrations when relocating. While autoconfiguration and directory services may eventually minimize the use of network addresses in configurations, these services do not currently provide relief for network address changes.

Furthermore, the existence of larger hierarchically-assigned addresses does not insure their usage by the thousands of applications which currently work over today's Internet. Several of the proposed routing and addressing solutions call for changing the programming interfaces for TCP and UDP services to support larger addresses. As a consequence, every existing network application will have to be revised, and organizations will have to insure that they have the "upgraded" versions of all their network applications before moving to the new architecture. It will be necessary to phase in the usage of larger addresses in a controlled manner in order to bring about their acceptance by the Internet community.

TCP & UDP with Network-independent Endpoints (TUNE) is one possible strategy for resolving the routing and addressing problems in the current Internet while accommodating existing applications. The TUNE architecture provides for the delivery of TCP and UDP messages in a manner similar to IP, but goes further by formally separating endpoint identification from network attachment information to improve mobility and scaling. By allowing the use of existing IP addresses as one class of endpoint identifiers, TUNE provides for an initial transition to TUNE services without application code changes.

This document describes the overall TUNE architecture and a migration plan for transition to CLNP network services. While TUNE will function over almost any network service, CLNP's existing hierarchical network addressing will require minimal routing information exchange once TUNE is deployed. In addition, the

Curran

Expires April, 1993

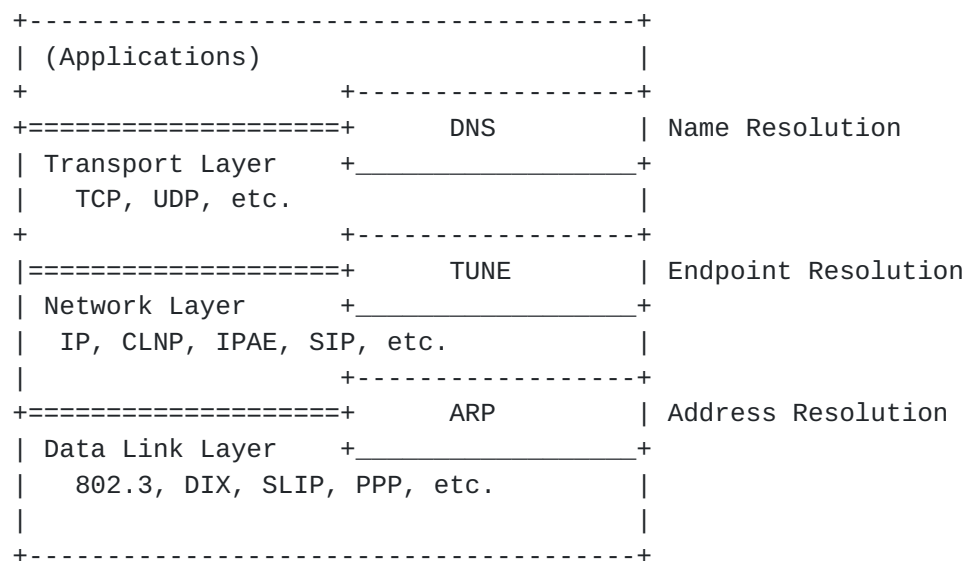
[Page 3]

availability of CLNP services in the current Internet will accelerate the deployment process and eliminate the risks associated with the development of a new network layer service.

2.0 TUNE Architecture

TCP & UDP with Network-independent Endpoints (TUNE) details a convergence layer between the Internet transport and underlying network datagram delivery services. The purpose of this layer is to allow identification of transport endpoints independent of the underlying network addressing scheme. This is analogous to the role that the Address Resolution Protocol [5] plays in allowing network address selection unconstrained by physical interface address.

TUNE's relationship to other network services can be depicted as follows:



TUNE provides two functions: a mapping function from endpoint identifiers to network protocols and addresses, and the labeling of messages with transport endpoint identifiers.

While TUNE is not required for successful implementation of TCP and UDP over new network protocols, the benefits of using TUNE include:

- o Globally unique endpoint identifiers which are independent of network attachment point.
- o Endpoint identifiers translated to network addresses in a manner transparent to the applications.
- o Underlying transport provided via multiple network services

with diverse address formats.

- o Hiding of network layer information from higher layers.
- o Backward-compatible programming interfaces for TCP and UDP services, and application compatibility during transition.
- o Full utilization of the 32-bit Internet Protocol address space and its existing allocation infrastructure.
- o Interoperability between IP and TUNE hosts during transition.

2.1 Endpoint Identifiers (EIDs)

Datagram-oriented transport protocols have traditionally used network layer addresses during the specification of communication endpoints. It is common for such network addresses to contain has some embedded information regarding the underlying network topology. IP addresses, for instance, implicitly identify the particular network which serves the end system. In the case of CLNP, addresses that are compatible with OSI Intra-domain IS-IS routing protocol [6] include explicit topological information in the "Area Address" (formed from the Initial Domain Part (IDP) and the High-order portion of the Domain Specific Part (HO-DSP) of a given NSAP.) Neither IP nor CLNP provide a endpoint identifier which is independent of network topology. As a result, topological information which is specific to the network must be made available to the applications for inclusion in their transport service requests. Likewise, applications using transport services become dependent upon the specific network attachment points at the time of configuration.

TUNE transport services do not require specification of topology information (implied or otherwise) by applications. Communication is provided between globally unique "endpoint identifiers". These endpoint identifiers (EIDs) are assigned hierarchically (for administrative convenience and directory manageability) but do not embody any network topology information. A system with two interfaces (on the same or different networks) will use a single EID for normal communications. Additionally, when a system moves from one network to another, the endpoint identifier does not change. This provides TUNE with endpoint identifiers that are independent of network topology and inherently mobile.

The success of a globally unique identification system depends ultimately upon the ease of obtaining identifiers. If the process for acquiring such an identifier is burdensome, then the probability of duplicate identifiers being used increases dramatically. For this

reason, TUNE make use of existing allocation schemes for endpoint identifiers whenever possible. [Appendix A](#) contains a preliminary specification of the format and allocation process for EIDs.

2.2 EID Resolution

In order to transport a TCP or UDP message to the appropriate endpoint, it will be necessary to determine the corresponding network service and address for any given EID. Given the current state of directory services in the Internet, the Domain Name System [[7](#)] is a likely choice for providing the EID to network address mapping function. One consequence of using DNS is that each TUNE system will require access to a DNS nameserver via a supported network protocol. Initially, it is anticipated that the existing IP nameservers will be used (with CLNP-only systems using transition services), but deployment of TUNE-supporting nameservers will occur as TUNE systems become more common. [Appendix B](#) specifies a DNS services plan for TUNE using the existing DNS software. This plan supports any NSAP format and provides flexibility in the selection of NSEL values for the TUNE service.

Once the destination network service is determined, the TUNE message is passed to the appropriate network service along with the destination network address. Delivery of the message is handled as is normal for the particular network service selected; TUNE does not affect the datagram forwarding process in any manner.

2.3 TUNE Header Format

Both TCP and UDP currently include a "pseudo-header" during checksum calculation. The pseudo-header includes the network-level source and destination along with the protocol identifier and transport message length. The purpose of this checksum is to insure that the message was delivered to the appropriate endpoint.

Since an endpoint identifier is no longer related to the network layer addresses, the use of these pseudo-headers by TCP or UDP must be revised with new endpoint semantics. Under TUNE, the destination of TCP or UDP datagrams can be uniquely identified by the EID and the port number alone. It is not necessary for pseudo-headers to contain network addresses at all, since EIDs are globally unique. In addition, maintaining network addresses in these headers can actually prevent communications in some instances because the messages arrived via a different network protocol or interface than the sender intended. The elimination of network address information in these headers provides for consistent endpoint specification during

mobility, and allows the use of multiple network services (with potentially different characteristics) by a single endpoint.

The revised TCP and UDP pseudo-header for TUNE operations is:

```

      0          7 8          15 16          23 24          31
+-----+-----+-----+-----+
|   source endpoint id (octets 1-4)   |
+-----+-----+-----+-----+
|   source endpoint id (octets 5-8)   |
+-----+-----+-----+-----+
| destination endpoint id (octets 1-4) |
+-----+-----+-----+-----+
| destination endpoint id (octets 5-8) |
+-----+-----+-----+-----+
| flags | protocol|   data length   |
+-----+-----+-----+-----+

```

Since TUNE decouples the one-to-one relationship between network attachment points and transport endpoints, the specific destination EID must be available in the network or transport message so that messages for the destination system may be recognized and processed. Source EIDs may also be desired in transport messages for logging, accounting, or authorization purposes.

As part of the TUBA discussions, it has been suggested that these EIDs might be embedded in network addresses via various mappings. This restricts the choice of network addresses (to those conforming to the mapping algorithm), and such mapping may not be available over other network services. Another consequence of embedding EIDs into network addresses is that the resulting transport identifiers used by applications become dependent on network addressing structures and thus subject to change with network changes.

TUNE maintains independence from network layer addresses by prepending TCP and UDP messages with the revised pseudo-header during delivery. This revised header (referred to as the "TUNE header") precedes the TCP or UDP data in a TUNE datagram and contains the protocol field needed for demultiplexing at the destination TUNE system. This is a significant departure from current IP architecture, and will result in larger network layer datagrams in most cases. While the addition of 24 octets to the average Internet datagram is considered heresy by many, the header will result in a minimal performance impact when used with large datagram sizes. The impact on low-speed circuits would be much greater, but may be alleviated through the use of header compression techniques in a manner similar to IP [8].

Curran

Expires April, 1993

[Page 7]

3.0 TUNE Administration

By introducing organizational endpoint identifiers, TUNE creates basic infrastructure for several administrative features which have been previously viewed as network-layer services. In particular, EID identification of transport messages can simplify the deployment of autoconfiguration, portability, and mobility for end-systems.

3.1 Autoconfiguration

In order for a TUNE system to properly interact with other TUNE systems, it must:

- o Have an operational network service
- o Know its own EID
- o Have access to the EID resolution service

Autoconfiguration of network layer services has been successfully deployed for IP [[9](#), [10](#)], and is under development for CLNP services. As TUNE operates on top of network services, it does not interfere with the operation of any network-layer mechanisms.

TUNE does provide an additional key which autoconfiguration techniques could use to gain hardware independence. Rather than utilizing an IEEE identifier from one of its network interfaces as a key for system information, (such as name and bootfile selection) it becomes possible to index such information by the EID, and have systems store (in a non-volatile memory) their EID. Changing either the network level address or the hardware address (such as board swap) would not affect the boot process. In some environments, it may be useful to store a secret key in addition to the EID for use via ticket granting service authorizing use of the EID.

TUNE operation will require either non-volatile storage of the EID resolution service configuration or use of the autoconfiguration process to obtain such information. BOOTP is already capable of returning the necessary DNS information, but any additional information (such as network service preference) will have to be added as BOOTP extension or stored locally.

3.2 Portability

Portability, loosely defined as the ability to carry a system somewhere and have it operate at the destination, has been considered

a desirable characteristic of the Internet's next architecture. This is particularly true if portability is going to be used to maintain topologically-assigned network addresses.

Since EIDs are not dependent upon any particular physical network, TUNE systems which are moved to a new network attachment point only need to update the EID-to-network-address mapping in order to achieve portability. This works regardless of the cause of the address change (physical move, provider change, etc.) Such address changes are completely transparent to applications under TUNE, since applications use EIDs.

Due to the use of DNS services for EID->address mapping, TUNE cannot provide "instant" portability during its initial deployment. This stems from the inability to perform distributed updates to the DNS database from remote systems. Organizations will enjoy portability for all of their TUNE systems, but will have to manually update the EID mapping when systems are moved.

The addition of dynamic update facilities could be done via DNS extensions for dynamic database update, or could be done as a separate facility which interacts directly with the master nameserver as needed. The advantage of a separate facility is the ability to readily add site-defined validation for such updates, whereas a DNS-based solution can readily be found during autoconfiguration and would not require specific autoconfiguration support as a result. Additional work is needed before making a determination for this topic.

3.3 Mobility

Mobility, defined here as the ability to operate a system continuously while changing network location, has been expressed by many as a desirable characteristic of the next Internet architecture. Several mobile technologies including wireless LAN's, microwave personal communications services, and digital cellular services will become generally available over the next few years and could make significant demands upon our routing architecture if pressed to provide mobility to systems.

In the TUNE model, mobility is provided at the transport layer. The TUNE header allows direction of transport messages without affecting the endpoint identifiers. Given two systems, EID XXX with network address NX1, and EID YYY with network address NY1, it is possible for XXX to move to a new network address (NX2) while communicating over CLNP as follows:

XXX

YYY

Disconnects from NX1/
connects to NX2

Sends dynamic update to
EID map (identical to
portability case)

Cached mapping for XXX incorrect/
messages timeout.

Cached ES info expires at
last IS router based on
holding timer.
(holding timer presumably
low for mobile host)

Cached mapping for XXX incorrect/
messages report Host Unreachable.

Failure causes reacquisition of EID

->

address mapping.

Mapping returns new information/
Messages to XXX/NX2 succeed.

For TCP connections, the reaction time to network address changes can be minimized by using a "refresh EID" flag in the TUNE header. When a system changes its network address, this flag should be set in the next datagram to each outstanding TCP connection.

[4.0](#) [Appendix A](#): EID Allocation and Format

TUNE Endpoint Identifiers are 64 bits in length, and may be represented in the familiar "dotted-decimal" notation used for IP addresses. The EIDs are allocated by the Internet Assigned Numbers Authority and its delegated authorities in accordance with the recommendations of the IETF.

The IANA should allocate EIDs in contiguous blocks of EIDs to subordinate allocation authorities as necessary to insure an adequate supply of EIDs for TUNE-using organizations. It is recommended that the IANA make use of multiple channels for distribution of EIDs including network service providers, professional organizations, standards organizations, and others as necessary.

In order to insure an plentiful supply of EID authorities, the first two octets of the EID are reserved for an allocation authority

Curran

Expires April, 1993

[Page 10]

prefix. This allows for 64K allocation authorities. It is foreseeable that the IANA might supply multiple authority prefixes to some authorities so that further delegation of authority may occur. An example of this would be if the IANA were to allocate sufficient prefixes to ISO so that each country could have a unique authority prefix.

Allocation authorities may allocate EIDs in any of the following sized blocks: 2^{**8} , 2^{**16} , 2^{**24} , 2^{**32} , 2^{**40} , and 2^{**48} . Authorities are required to maintain records of blocks assigned and provide support infrastructure (such as top-level directory servers) to allow successful EID resolution process for each block assigned. The requirements for directory services are intended to discourage small allocations by authorities, and it is expected that typical EID block allocations will consist of 2^{**16} or 2^{**24} EIDs per organization.

In order to provide interoperability between IPv4 and TUNE systems, it is necessary to "map" existing IP addresses into TUNE EIDs. It is recommended that the following EIDs be reserved for this purpose:

128.0.x.y.z.0.0.n

where

128.0	Authority Prefix
x.y.z	IP address (high 3 octets)
0.0.n	IP address (low octet)

Although this method of allocation does not provide a straightforward translation for existing IP addresses, it will allow existing network (and subnet) allocations to evolve into large EID namespaces prior to the establishment of multiple EID allocation authorities.

4.0 Appendix B: DNS Support for EID Resolution

TUNE requires a distributed mapping function from a given EID to the appropriate network service and address used to reach it. This function will be initially provided through existing DNS services in order to allow timely deployment in the Internet.

It is necessary to represent the EID information in the DNS namespace in a manner which supports delegation from the IANA to each assignment authority, and from each assignment authority to each organization. DNS has two restrictions upon how delegation may be performed: delegation may only occur on token boundaries (where tokens are delimited by periods in the name string), and information

is organized with the rightmost tokens being most significant.

Considering these requirements, the following approach is recommended for EID to network information mapping via DNS:

The IANA will establish DNS service for a new domain: IN-EID.ARPA.

For each authority prefix assignment, the IANA will delegate a subdomain of IN-EID.ARPA for management by the assignment authority. The subdomain may be determined by representing the authority prefix in dotted-decimal notation and inverting the order of tokens (as is done for the IN-ADDR domain currently)

Example:

Authority ABC is assigned prefix 0.47. A matching domain of 47.0.IN-EID.ARPA would be operated by or on behalf of the authority.

For each organization prefix assigned, the authority will delegate a subdomain of their domain for management by the organization. The subdomain may be determined by representing the organization prefix in dotted-decimal notation and inverting the order of tokens (as above).

The IANA will delegate subdomains for any organizations using EIDs based upon their IP network address allocation.

Example:

Organization XYZ is assigned prefix 0.47.255.238.0. A matching domain of 0.238.255.47.0.IN-EID.ARPA would be operated by or on behalf of organization XYZ.

Organizations are responsible for providing a pointer (PTR) record for each EID in use. The location for the PTR record may be determined by representing the EID in dotted-decimal notation and inverting the order of tokens (as above).

The PTR record will contain the name of the network service used to reach the EID, a slash ("/") character, and a network address in a format specific to the network service.

The valid names and address formats will be determined by the IANA. The use of PTR resource records requires that all address formats end with a period (".").

A network address may include a network layer protocol selector when available (e.g. NSEL in NSAP addresses) as TUNE only requires a

single dispatch point and performs transport protocol demultiplexing internally based upon the TUNE header.

TUNE hosts should only process address records for supported network protocols and should ignore all others received. The method for selecting which network service and address to use in the presence of multiple valid records is a local configuration issue, and not addressed here. It is conceivable that address selection might take local policy into consideration to provide limited source-specified routing.

Examples:

A system with EID of 128.0.192.52.71.0.0.4 reachable via CLNP might require a PTR record of the following form:

4.0.0.71.52.192.0.128.IN-EID.ARPA.

PTR CLNS/47000580FFEE00000000001100AA00040084AF10.

A system with EID of 0.47.255.238.0.0.1.178 reachable via PIP might require a PTR record of the following form:

178.1.0.0.238.255.47.0.IN-EID.ARPA

PTR PIP/7.66000.77.18.2.

5.0 Security Considerations

Security considerations are not discussed in this memo, but should be thoroughly explored before any future Internet architecture is selected.

6.0 References

- [1] "Protocol for Providing the Connectionless-Mode Network Service", ISO 8473, 1988.
- [2] Callon, R., "TCP and UDP with Bigger Addresses (TUBA), A Simple Proposal for Internet Addressing and Routing", [RFC1347](#), 1992 June.
- [3] Fuller, V.; Li, T.; Yu, J.; Varadhan, K, "Supernetting: an Address Assignment and Aggregation Strategy", [RFC1338](#), 1992 June.

- [4] Collela, R.; Gardner, E.P.; Callon, R.W., "Guidelines for OSI NSAP allocation in the internet", [RFC1237](#), 1991 July.
- [5] Plummer, D.C., "Ethernet Address Resolution Protocol: Or converting network protocol addresses to 48.bit Ethernet address for transmission on Ethernet hardware", [RFC826](#), 1982 November.
- [6] Oran, David, Ed., "OSI IS-IS Intra-domain Routing Protocol", [RFC1142](#), February 1990.
- [7] Mockapetris, P., "Domain names - Concepts and Facilities", [RFC1034](#), November 1987.
- [8] Jacobson, V., "Compressing TCP/IP headers for low-speed serial links", [RFC1144](#), 1990 February.
- [9] Croft, W.J.; Gilmore, J., "Bootstrap Protocol", [RFC951](#), 1985 September.
- [10] Reynolds, J.K., "BOOTP vendor information extensions", [RFC1084](#), 1988 December.

[7.0](#) Author's Address:

John Curran
Bolt Beranek and Newman, Inc.
10 Moulton Street
Cambridge, MA 02138
(617) 873-4398

jcurran@nic.near.net

