Internet Engineering Task Force
Internet-Draft
Expires: January 10, 2009

S. Chakravorty
J. Bush
The MITRE Corporation
J. Bound
NAv6TF

July 9, 2008

IPv6 Label Switching Architecture
draft-chakravorty-6lsa-03

Status of this Memo

Abstract

This specification provides an architectural framework, called IPv6 Label Switching Architecture or 6LSA, for an end-to-end, IP-centric packet transmission technique that uses the IPv6 packet header Flow Label to establish IPv6-based label switched paths. The label switched paths, called 6LSPs, provide application and user specified routes for efficient transport of packets and as means for quality of service (QoS) delivery, IPv4 tunneling, VPN and other mechanisms. Through look-ups of 20-bit labels instead of 128-bit IPv6 addresses, the architecture may provide potential memory and processing savings, the latter through significantly reduced address fetches for the low-

powered, handheld devices. The label has two components comprising Global Label value and Local Label value. The Global Label value from the source is delivered to the destination unmodified. However, the intermediate network nodes in 6LSA are allowed to temporarily replace the Local Label value with a value of local significance. This enables 6LSA flows to be hop-specific although session-based and as such a unique QoS delivery technique for bandwidth constrained media. 6LSA also enhances security since label generation and assignment algorithms can be modified periodically.

Finally, it must be pointed out that the 6LSA concept of temporary flow label assignment is applicable to the 6LSA domain only. The concept is not applicable to domains outside the 6LSA.

Table of Contents

1. Introduction

Several approaches have been developed over the past decade to provide QoS in large networks. These include DiffServ, RSVP, MPLS, ATM, extensions to routing protocols, and proprietary mechanisms. Some of these techniques can also be applied to IPv6 transport but not directly. IPv6 offers strong features to enable QoS delivery, most significant among them are the Flow Label and the Routing Header extensions.

The IPv6 Label Switching Architecture (6LSA) specified here enables fast switching using 20-bit labels instead of 128-bit IPv6 addresses and thereby provides memory and processing savings, the latter because address fetches for low-powered, handheld devices, which are mostly 32-bit architectures, could be up to four times as fast. The architecture allows temporary replacement of labels inserted by a source node with the proviso that the original labels are reinserted prior to the packet arriving at its destination.

This document introduces an architectural framework for the use of IPv6 packet header Flow Labels for setting up labeled paths of local significance to provide services that cannot be provided by a purely routed, connectionless approach. The 6LSA allows local label generation in a network node. It also allows a network management entity updating available label tables, across the network to reduce man-in-the-middle attacks.

This local label generation coupled with dynamic labeling helps the 6LSA to provide the means for mobile nodes to set up labeled paths, automatically and/or manually, for end-to-end QoS delivery or for any other service delivery.

2. Overview

The IPv6 label switching mechanism makes use of the 20-bit Flow Label in the IPv6 header to assign flow IDs which are used for labeled paths, also called IPv6 label switched paths (6LSPs). The specification of IPv6 label is in conformance with RFC 3697 RFC 3697 [1], IPv6 Flow Label Specification, in which the use of Flow Label has been recommended for establishing different types of flows at the source nodes and packet classification in the intermediate nodes. The proposed mechanism of IPv6 label switching broadens the scope of the use of Flow Labels beyond its simple use for packet classification to its use of packet classification and forwarding. The component of the Flow Label, called Local label value, is locally assigned in the packet header along a path. This part of the Flow Label can be different from the corresponding part of the original

Flow Label assigned by the source and as such this component is temporary and has only per hop significance unless such significance is extended to multiple hops.

The 6LSA concept of Flow Label assignment is applicable to the 6LSA domain only. Whether the concept is applicable to domains outside the 6LSA is of no concern to 6LSA unless there is one-to-one mapping in the labels and implied significance.

In a conventional router, the forwarding decision is independently made in each router as the packet travels from one router to the next. In each router, the packet header is analyzed using a network layer routing algorithm to find the best next-hop that is often the shortest distance to the next router. Since this forwarding in each router is "independent" of how the previous packet for the same destination was processed and forwarded, the routing is considered connectionless.

The process specified in this document supports two functions: first, grouping of packets of similar flow requirements into a Forwarding Equivalence Class (FEC), and second, forwarding all packets belonging to an FEC along the same path. The forwarding of packets does not necessarily have to be along the paths as determined by the routing algorithms or by manual configuration provided there is no looping of packets caused by the 6LSA forwarding.

In 6LSA, the FEC is encoded in the Flow Label field as a non-zero value, which is also called label in this document. The label is available in one of 3 ways: (1) locally generated, based on certain algorithm or policy, (2) in the incoming packet, as Flow Label from a source node, or (3) distributed, through a label distribution process. The assignment of a label to an FEC is identified as a binding. Once the binding of a label is created and provided to an FEC, the forwarding policy of the packet may be represented through the label and is maintained at a minimum for the duration of the session.

The selection of the FEC is based on one or more flow characteristics. The selection of flow characteristics and therefore of FEC is an administrative, service or policy decision, or a combination thereof. Such a decision is meant to support certain traffic requirements, such as availability bandwidths on certain links, VPN (Virtual Private Network) configuration, values encoded or configured into the Traffic Class field of IPv6 packet headers, or requirements conveyed by the source or administrative entity or by other means. A superset of global FEC selections and corresponding label values are included in this document.

Merging of flows on a single 6LSP is possible with the consequence that two or more flows are inseparable and indistinguishable inside a 6LSA domain from the merger node to the egress or penultimate node. Merging of flows may lessen the amount of state maintained within 6LSA nodes, however, since there is no stacking of labels in 6LSA and each packet's label has to be examined individually, the processing saving is insignificant.

The 6LSA label supports the forwarding of packets belonging to the same FEC along an IPv6 Label Switched Path (6LSP). For the lead packet of each flow, traffic class, source and destination addresses, and possibly other special handling requirements conveyed by the source (e.g., by a control plane protocol) may be examined for FEC identification and label selection. This decision process is independently carried out on each 6LSA node transited by the lead packet across a 6LSA domain. Subsequent packets of the same flow (or a group of flows) typically do not go through the same process of label selection and assignment because the label binding to an FEC and egress interface has been cached. So, the subsequent packets can be rapidly forwarded.

The 6LSA domain routers are not cognizant of labels outside of their domain.

3. Terminology

This section provides a general overview of alphabetically arranged terms that are used in this document. Some of these terms are more precisely defined in the later sections of this document.

6LSA	a set of nodes that perform 6LSA routing and forwarding operations and are in one 6LSA domain
6LSP	label switched path, a virtual path, through a pair or more 6LSA nodes
6LSR	IPv6 label switching router that is capable of forwarding IPv6 packets based on FEC attributes
FEC	Forwarding Equivalent Class, collection of IP packets that receive the same forwarding treatment and are forwarded over the same 6LSP
Flow	sequence of packets identified by the Flow Label
Switching table	forwarding table that comprises packet forwarding information

4. Distinguishing Characteristics of 6LSA

The 6LSA characteristics justify its application wherever IPv6 is deployed and where QoS network performance or other service delivery is required, or where other available packet forwarding mechanisms cannot deliver packet performance end-to-end. The following special characteristics of 6LSA are listed here in no particular order:

- 6LSA technology offers a methodology for use of flow label in the IPv6 header which is as yet unused.

- No Extraneous Label Bindings - 6LSA eliminates the need for using extraneous labels (labels from other than Layer 3) and/or eliminates the need for associated label distribution across the network

- No IPSec Constraints - 6LSA does not need to use the Layer 4 ports to define a flow and therefore can be used with IPSec VPNs or other IPSec based services.

- Free of Layer 2 Overheads - Being a layer 3 packet forwarding solution, the 6LSA does not need a layer 2 packet forwarding mechanism such as ATM and as such 6LSA avoids ATM's fragmentation and re-assembly delays and associated header overhead. It also avoids the need for added signaling and state machine mechanisms to provide

ATM switched paths and ship-by-night capabilities. Additionally, being based on routing protocols (6LSRs peer in Layer 3), 6LSA tends to avoid the potential for O(N2) and O(N3) problems.

o Deployment Ease - The 6LSA can be deployed over all layer 2 protocols such as Ethernet and PPP. There is no layer 2 interface limitation in 6LSA.

o Extensive QoS Label Space - The 20-bit Flow Label in addition to the 8-bit Traffic Class field can provide a huge traffic classification space, both the fields may be used together in the 6LSA.

o Feature Visibility - All of the IPv6 packet header features are available while the packet is enroute to destination as such IPSec or similar packet encryption does not disable the delivery of QoS or other similar services that 6LSA can provide.

o IPv6 Transition Support - During the transition from IPv4 to IPv6, the latter is generally deployed as network-islands surrounding IPv4/MPLS core. The implementation of 6LSA in native IPv6 edge networks will greatly facilitate traffic engineering between the edge and the core by mapping 6LSA Global Label value in the Flow Label to be carried over to MPLS label in the core if the FEC representation by the flow label is common in both domains, for example, if the label represents a given bandwidth, say.

o Security Enhancement - Since 6LSA allows node-local generation of labels, such generation, where adopted, can be made totally random or periodically synchronized across the 6LSA domain to considerably reduce man-in-the-middle attacks.

To summarize, 6LSA provides a significant layer 3 capability for switching IP packets - with little or no added overhead for signaling.

5. Routing Versus Switching of IP Traffic

The routing of packets on the Internet has the following essential characteristics in addition to connectionless, non-sequential packet transport: 1) The paths are not dedicated virtual paths, and 2) There is generally no delivery or QoS guarantee - only a single class of traffic is available.

Switching of packets has the following basic characteristics: 1) The routing of packets is connection-oriented; the flow of packets comprises sequential packets, 2) The packets travel over dedicated,

virtual paths or virtual circuits (VCs) which are either temporary or permanent, and 3) Switching is generally associated with certain delivery guarantees and traffic classification.

The issues with switching are:

* There are delays caused by VC setup time across the network.

* There is a need for VC state maintenance.

* IP address to VC translation latency is always present however small.

* Potential is there for large resource wastage in case of a link or node failure in IP virtual network over switched network, for example IP over ATM that may cause N2 and N3 problems.

The design of 6LSA overcomes the above disadvantages of switching by making the establishment of the switched path hop-specific but flow session bound.

6. 6LSA Basic Components and Their Attributes

In this section, the 6LSA basic components and their attributes are defined.

6.1. Flow

A flow in 6LSA is identified by the label value in the Flow Label field in the IPv6 packet header. All packets belonging to the same flow must be sent with the same Flow Label per hop, at a minimum, and from the source node to the destination node, at a maximum. The label in the lead packet and that in the subsequent packets of a flow may be different or the same depending upon the algorithm selected. In 6LSA domain, non-zero label identifies a best effort delivery.

6.2. Forwarding Equivalence Class (FEC)

The FEC represents a group of IPv6 packets that all receive the same forwarding treatment. The forwarding treatment may be based on one or more attributes associated with an FEC or other processing requirements imposed on the flow externally.

Several flows from multiple sources may receive the same forwarding treatment and thus belong to the same FEC. For example, if multiple flows are to be processed in the same outgoing queue because they all deliver the same service, they are identified by the same FEC.

6.3. Label

A label is the 20-bit, fixed length Flow Label identifier in the IPv6 header. A node in the 6LSA binds the Flow Label to a Forwarding Equivalency Class (FEC) and uses it to forward the packet. The label may thus represent the FEC. In the 6LSA, the FEC indicates the forwarding treatment a group or class of packets (with a given label) receives. This label and the FEC it represents have only local (per hop) significance unless the attributes associated with a FEC are shared among multiple hops.

A label is thus associated with the characteristics of a flow which is represented in the FEC.

A flow label is assigned to a flow by the flow's source node. New flow labels must be chosen (pseudo-)randomly and uniformly. The purpose of the random allocation is to make any set of bits within the Flow Label field suitable for use as a hash key by routers, for looking up the state associated with the flow.

A label in an incoming packet is called an "incoming label", and that in an outgoing packet is called an "outgoing label".

The 6LSA nodes are permitted, but not required, to verify that the flow conditions are satisfied. If a violation is detected, it should be reported to the source by an ICMP parameter Problem message, Code 0, pointing to the high-order octet of the Flow Label field.

A source node must not reuse a flow label for a new flow within the maximum lifetime of any flow-handling state that might have been established for the prior use of that flow label.

The first 7 bits of the Flow Label are of global nature that all nodes in a 6LSA domain need to treat identically. This component of the Flow Label thus must be maintained end-to-end without any change. The security implications of maintaining Global Labels end-to-end are clear and therefore will not be addressed further in this document.

A primary selection of Global Label value is presented here below for FEC usage as guidance. Note that the most significant bit is the left most bit. In the selection provided below, only the first 3 leftmost bits identify the superset of the flow characteristics. The next four bits provide further identify granularity of the flow characteristics.

The selection for the Enterprise Specific category allows an enterprise such as a corporation, service provider, or governmental agency to have their own specific Global Value component to take

advantage of enterprise's unique feature sets and for added security.

The last 13 bits can be locally selected by each node or a group of nodes; they are typically hop-specific. This hop specificity may extend for more than one hop. Any node not associated with this hop need not adhere to the selected FEC. The Local Label Value is of local significance only unless it is extended to more than the local hop associated nodes.

The Global Label value provides space for 8 major categories each with 16 different subcategories. A substantive subcategory of reserved space is also available for future or application specific use.

The total Local Label value space available to 6LSA nodes is between 1 and 2^13 times yielding the ability to uniquely identify 8,192 6LSPs per physical (or virtual) port per hop. The label space applies to each physical interface.

FEC (First 3 Bits)	Next 4 Bits	Purpose
No FEC (000)	0000	
Domain Specific (000)	0001 - 1111	

VPN (001)	0001	(IPSec - Tunnel Mode)
	0010	(IPSec - Transport Mode)
	0011	(Special Encryption)
	0100	(VRF)
	0101	(End Network Specific)
	0110 - 1111	(Reserved)

TE Subset/	0001	(DiffServ)
QoS Enhancement (010)	0010	(RSVP)
	0011	(RSVP-TE)
	0100	(SIP)
	0101	(H323)
	0110	(Large File)
	0111	(Circuit Emulation)
	1000	(Fixed Bandwidth)
	1001	(Video Streaming)
	1010	(Multicast)
	1011	(Anycast)
	1100	(Queue Weighting)
	1101	(Precedence Sensitive)
	1110	(Enterprise Signal)
	1111	(Reserved)

Encapsulation (011)	0001	(IPv6 in IPv6)
	0010	(IPv4 in IPv6)
	0011	(Other in IPv6)
	0100	(Enterprise Specific)
	0101 - 1111	(Reserved)

Enterprise Specific(111)	0000 - 1111	(Reserved)

As an example, between nodes A and B in a 6LSA domain, if there is a group of packets that are tunneling IPv4 packets in IPv6, then the first 7 bits of the label for all these packets will be marked 0110011. A 6LSA node may choose to use the remaining bits to signify different physical or logical ports or such other characteristics that are locally assignable.

For each of the above selections, up to 8,192 additional local selections are available for each FEC group for one or more hops as represented by the remaining 13-bit space.

A total of 16 selections are available to an enterprise for use as desired. For each of these 16 selections, an enterprise may use additional 13 bits in the remaining flow label space for node specific processing needs.

It is envisioned that an application, middleware, end-system socket software or network node software will be enabled to encode the needed FEC for a flow just as DSCP marking is encoded for DiffServ.

In this document, the nomenclature of label or Flow Label is variously used though the meaning is the same and while referring to the word label used in other technologies, such as in MPLS, the context of the technology is also mentioned to distinguish the application and meaning of the word.

6.4. Labeled Packet

In 6LSA, a labeled packet is an IPv6 packet whose header has a non-zero Flow Label value.

If an incoming packet into a 6LSA node has a zero label value and it is not a lead packet, it is assigned a non-zero label value before it is forwarded. This is because in 6LSA, a zero label is used to indicate that the packet is a best-effort packet.

The label from a flow in the 6LSA may be transferred to a non-6LSA network layer or to a non-6LSA data link layer as long as there is a field available that can carry the 6LSA label. The particular encoding technique and its significance must be agreed by both the layers - layer 3, the network layer, and layer 2, the data link layer. Specifics of the method of this encoding are outside the scope of this document.

6.5. IPv6 Label Switching Router (6LSR)

A router operating in the 6LSA is an IPv6 label switching router, called 6LSR, which performs as a minimum the two 6LSA functions of: 1) replacing (swapping) an incoming label in a packet with an outgoing label, and 2) forwarding the packet based on the appropriate forwarding treatment.

6.6. Lead Packet

A packet arriving at a 6LSA node is a lead packet if it is the first packet of the flow that this node has received. A lead packet may or may not be the first packet of the flow that the upstream router or 6LSR forwarded to this 6LSR. It is possible that the first packet in the flow is lost or misdirected, or for that matter, first few

packets in the flow are lost or misdirected.

All lead packets, whether they are the first packet from the upstream 6LSA node or not, are treated like they were the first packet of the flow.

Lead packets have no existing entry in the switching table of a 6LSR or 6LSA node.

6.7. Switching Table

A switching table in 6SLA is often called the forwarding table in the networking literature. This table comprises the following information as they become available:

-- Label value from the lead packet, if there is a non-zero label otherwise a zero value is entered
-- Incoming interface, that is, interface on which the packet arrived
-- Next hop IPv6 address
-- Outgoing interface, that is, the interface on which the packet is forwarded to the next-hop 6LSR
-- FEC value that identifies the forwarding operation that needs to be performed on a packet

The outgoing label entered in the switching table has to be unique so that this node or 6LSR is able to identify a unique LSP for the packet.

Additional information that may be available in the switching table is as follows.

* The data link layer and encapsulation to use
* How the label value, typically the Local Label value, needs to be encoded in the label field
* Timers associated with the packet
* Last packet in the flow
* Information with regard to how to discard labels or packets

The format and content of the switching table entry are implementation and configuration specific and are not specified here.

7. Attributes of Label Binding

The binding of a label to a FEC is based on known attributes of the packet or on externally applied constraints. The binding of a label

to a FEC is local to a 6LSA node unless such binding is promulgated across the network through some exterior means such as a using a label distribution protocol (LDP) commonly used for MPLS. LDP is outside the scope of this document.

8. IPv6 Label Switched Path (6LSP)

A label switched path in the 6LSA is called 6LSP and identifies a virtual circuit through which one or more flows are routed to the next-hop 6LSR.

The sequence of 6LSA nodes, that is, the sequence of nodal interfaces through which labeled packets are transported represents a 6LSP. A 6LSP is thus represented by a label between any two nodes, and by one or more sequence of labels between multiple nodes, or between two or more hosts, or any combination thereof. Conversely, a 6LSP may also represent the FEC binding of each flow in each of the 6LSR on the 6LSP. In this regard, 6LSP closely resembles a virtual circuit (VC) in ATM, and LSP in MPLS.

9. 6LSA Architectural Functionalities

This section describes 6LSA functionalities including label acquiring, label binding to FEC, and label swapping.

a) The 6LSA comprises two basic functions: first, grouping of packets of similar flow requirements into a FEC, and second, speedily forwarding all packets belonging to an FEC along the same path - including flow merging when multiple flows have the same FEC characteristics.

b) A 6LSA domain edge 6LSR replaces the incoming Local Label in a packet with an outgoing Local Label establishing a unique 6LSP for one or more packets in the flow associating it with the same incoming Flow Label and source and destination address combination. Each 6LSA node ensures there is no duplication of 6LSPs from itself to the downstream node for the same FEC.

c) An intermediate 6LSR receives a flow on a unique 6LSP. It replaces the incoming label with a FEC bound outgoing label based on the needed local treatment of the packet. The last 13 bits in the label that this 6LSR encodes represent the local 6LSP that this 6LSR sets up. This last locally encoded 13-bit value may add unique characteristics to the treatment of the packet or its path in addition to the global FEC treatment requirements. The total label value is entered in the local switching table. The packet is then forwarded to the next downstream 6LSR. The first 7 bits representing global FEC part may not be changed by a 6LSA node.

d) Within a 6LSA domain, if a router is a 6LSR, it swaps the label, if it is not a 6LSR; it lets the flow pass without any label changes.

e) Each 6LSP is maintained at least for the duration of the session of transport of all packets in a flow. In 6LSA, labels may be maintained for a pre-determined time after a session has ceased to exist, that is, for a fixed amount of time determined a-priori after packets belonging to a flow have ceased to arrive.

There are two salient points that need to be emphasized. First, the same label is not used for any two 6LSPs from an upstream 6LSR to a downstream 6LSR for the same pair of physical ports. Second, the last 13 bits of a label may have little significance for the downstream 6LSR unless it is aware of the significance of these bits a-priori. It has relevance for the upstream 6LSR which uses it to bind an FEC to the label and then uses the 6LSP to forward all the packets in a flow. The 6LSP thus becomes a representation of the FEC and forwarding pointer for the upstream 6LSR. The only time a 6LSP and therefore the associated full label has any relevance value to the downstream 6LSR is when the label bindings are distributed across a given 6LSA domain.

A 6LSA label by default creates a tunnel for all packets in a flow since the significance of the label bound to an FEC requires a special handling by every node for all packets in the flow which is not unlike a tunnel. The processing of packets based only on the label, once a specific label has been assigned to a flow for each hop, is similar to processing of a tunneled flow where the processing is generally determined by the tunnel header bits.

The incoming flows into any of the nodes can arrive on one or more 6LSPs. If the outgoing flows are merged onto the same 6LSP, the downstream node receives the merged flow with the same label. Packets belonging to a merged flow although are from different flows, they each still need to be examined by their label. However, once the flows are merged, distinction between individual flows may not be necessary for packet forwarding.

10. Label Assignment

In the 6LSA, the decision to assign or bind a label to a particular FEC is made by a 6LSA node, generally by the ingress node if it is not the origination point for that flow. The 6LSA host or node then informs the next-hop, downstream node of this binding via this labeled IPv6 packet or via some other method depending on the nature of label generation and distribution methodologies.

Three models have been identified for acquiring label space. The first model specifies how the labels can be generated locally; the second model refers to how labels can be acquired from label distribution, and the third, how labels are acquired from incoming packet headers. Only one of these three models of label assignment is allowed in a network of 6LSA-based nodes.

Once a label binding is available, the 6LSA requires that the label binding be retained for the duration of the session.

It is quite possible that multiple flows may require the same label and label binding to a single FEC. In these cases, all such flows may be multiplexed or merged together as one outgoing flow and forwarded on one 6LSP using the same label. At the de-multiplexing 6LSA node downstream, the flows must be discernible through the unique source and destination addresses or through other means.

The 6LSA does not prevent any 6LSA node from storing any label generated at a time different from when it is being used nor does it prevent a node from using any label that was used earlier or retrieved from a flow that used an algorithm or a model other than those identified here.

10.1. Locally Generated Label Model

In this model, 6LSA allows a node to generate its own labels to be used in the IPv6 header. The specification for the algorithm(s) used to generate the 20-bit labels is beyond the scope of this document.

An example algorithm for generating flow labels is a pseudorandom

number generator outputting values within the parameters algorithm in which the bit values are within the parameters specified in Section 6.1, Label. It is envisioned that a set of labels will be generated for every service class, elastic and inelastic, such as file transfer, voice, video, etc.

The Locally Generated Label model does not preclude manual generation of a label or range of labels through a configurator or similar other means. The locally generated label may have a value related to one or more attributes that is applicable to the next-hop node or nodes across the network.

The 6LSA specification allows labels that are locally generated but may have non-local significance. Such attributes may be regularly or randomly refreshed by a network management or other systems. Methodologies for such refreshment are outside the scope of this specification. The Global Label part (the first 7 bits) has by default the same significance for all nodes in 6LSA. However, the Local Label value (the last 13 bits) may or may not have only local significance.

This model is not a mandatory part of 6LSA, i.e., a node is not required to implement this model in order to be considered 6LSA-compliant. However, when a 6LSA node claims to implement the Locally Generated Label Model, the implementation must conform to the specification given in this document.

The use of this model is encouraged because it is simple, efficient and avoids control plane traffic for label distribution as in the Distributed Label Model.

The Locally Generated Label Model enhances security since the labels have local significance only and can be randomly or periodically refreshed all through the 6LSA domain prior to their use.

10.2. Distributed Label Model

The 6LSA allows distribution of the Local Label (value) space generated in one or more nodes or externally in a server. The architecture also allows more than one label distribution protocol (LDP) and sharing of necessary information amongst the label distribution peers.

Mechanisms or protocols that allow a Local Label distribution and do not violate any of the 6LSA specification with regard to use of such labels and the operation of 6LSA nodes are allowed. The specifics of label distribution protocols are outside the scope of this document.

The specifics of the process by which a node binds a label to a FEC are implementation specific and thus outside the scope of this document.

The Distributed Label Model enhances security if the labels have local significance only and can be randomly or periodically refreshed all through the 6LSA domain prior to their use. For reasons of efficiency, this label model may only distribute the Local Value part of the label.

10.3. Reuse Label Model

The 6LSA allows a node to reuse existing label available in the node or obtained from labels in the incoming packets and where the flows can be associated with unique label bindings.

11. Scope and Uniqueness of Labels

A label in a packet on a 6LSP must be unique to a flow, in any given direction, between interfaces on peering nodes that are one hop apart.

A flow always originates at the upstream source node in the 6LSA domain, continues through multiple 6LSRs and terminates in the destination node which also must exist in the 6LSA. Such a flow must carry a non-zero label in its lead packet.

In the unusual event where two flows have lead packets with the same label, the follow-on labels used by the ingress routers are kept different for the two flows. In all 6LSR routers, the discriminator for these two lead packets is the physical port, source and destination address pair or such other means as allowed by the 6LSA implementer.

To summarize, the discriminator for the incoming packets from an upstream 6LSR to this 6LSR is the 6LSP and the discriminator for the outgoing packets from this 6LSR to a downstream 6LSR is the FEC. Generally, for a flow, an incoming label represents the upstream 6LSP and an outgoing label represents the downstream FEC. In many cases, the same label may be usable or used for both.

12. Label Retention Mode

If a 6LSR B receives a label binding from a 6LSR A for a particular FEC via LDP, even though B is more than one hop apart from A, then such binding may be retained or discarded by B. If the binding is

retained, then this binding may be used later when A becomes a next-hop 6LSR to B. If the binding is discarded, then B will have to acquire a new binding for traffic from A through one of the three models described in Section 10.

13. Label Stacking

The 6LSA does not allow IPv6 label stacking. There is only one label in an IPv6 packet and this label must be in the Flow Label field in the IPv6 packet header. The 6LSA allows multiple label spaces per platform; however, the use of the label must conform to the specifications stated in this document. It should be noted that this does not preclude other non-IPv6 label stacking such as layer 2 label stacking.

14. Label Swapping

Label swapping or label switching is a process in which the incoming label in a packet is replaced with an outgoing label. In this document, this process is associated with multiple other activities. These activities include matching the switching table entries with certain incoming packet header fields, binding a label to the FEC, updating the switching table and finally, forwarding the packet. However, when the swapping involves only incoming label replacement with an outgoing label, it is called label switching, which typically is a fast process and may be carried out in the interface card itself.

15. Packet Processing Algorithms

The 6LSA packet processing algorithm refers to handling of packets that arrive from hosts in a 6LSA domain. The 6LSA packet processing provides for QoS priority, VPN handling and other forwarding treatments.

At a source node when an IPv6 packet is created, the Flow Label field is encoded with a Global Value of the label depending on the nature of the application. The Global value represents the generic nature of the service or flow characteristics. It may be incognizant of and unrelated to the hop-specific flow constraints that may exist. This Global Value encoding is then followed by the hop specific Local Value label encoding. This value may be locally generated, provided by a routing or such other protocol, manually configured or a combination thereof.

An example can best expound the labeling at a source node. Thus if a source node application intends to send a SIP packet to the far end, it determines the related Global Label value of the label which is 0100100 from the table provided in Section 6.3. Thereafter if the local hop specific treatment requires that it be sent on a specific physical port for a prioritized flow (since this signaling), it will search the configuration information provided to it for a corresponding applicable Local Label value. Let's assume, this flow has to go through the fourth port and has the highest priority handling requirement. Let's also assume that this is represented by 0000010000001. The total label value is then 0100100 00000010000001 and this then is encoded in the Flow Label field. Note that the node carries out a few other activities such as inserting in the switching table the source and destination of the packet, its Global and Local label values, and such other needed items. This switching table can also be used for holding configuration information such as FECs and related Local Label values.

The Local label value may well have been determined from the routing table as much as the Global Value from the Traffic Class field in the packet header. How this determination is made in any specific 6LSA domain or node is implementation specific.

In this example, the two characteristics of the flow represent the FEC of the flow and the selected 20-bit label is then said to be bound to this FEC.

Once the label has been bound to the FEC, the source node continues to send packets encoding them with the same label for this flow. The path of the flow then becomes the 6LSP for the flow.

At the next hop, an intermediate 6LSA node receives the packet on its third port connected to the source node. It first checks to see if either the Global or Local Label value exists in its Switching Table. If there is a match available for the Global Value label, it goes to Step 2, otherwise it follows Step 1.

Step 1: On examining the Global Label value, it determines the nature of the flow. From information provided through routing protocol, manual configuration or otherwise, it knows that such signaling flows need to be sent out via the second port of the node and that all such packets have to be queued ahead of all other packets. It identifies a proper Local Label value, which say is: 0001010000001. It replaces the incoming Local Label value with this value, queues the packet ahead of all other packets in the outgoing buffer and forwards the packet through the second port. This then presents the processing of packet in the intermediate node. As was stated earlier for the source node, additional processing takes place in the intermediate

node regarding updating of the switching table with the needed information from the outgoing and incoming packets.

Step 2: If there is a match for the Global Label value, then the Local Label value field is also searched in the Switching table and if they both match, the incoming Local Label value is switched and the packet is sent out via the second port. This continues for all packets that have both values match. However, if there is no match for the Local Value field, then the received packet is treated as the lead packet of another flow coming from the same source and the processing is identical to Step 1 since a new 6LSP has to be established for this flow.

In uncommon cases, a source node may create two flows for the same destinations with two FECs. If it does, the intermediate node has to assign two different labels and provide two label bindings for the two flows.

The processing that takes place in the second and subsequent intermediate nodes is the same as that in the first.

The process at the destination node is similar to the first intermediate node with the variation that the packets are forward to the corresponding application in this node instead of forwarding them out of this node.

Duplication of flow labeling is not possible since the local label is always unique for each flow and LSP although the same label may be used for different flows in unrelated hops out of the same node or other nodes in the same 6LSA domain.

16. Fast Switching

When a 6LSR can simply swap an incoming label with an outgoing label without going through insertion of new entry in the switching table for that packet, then this swapping is termed fast switching in this specification. This occurs when flow characteristics are well-established or deterministic enough that no additional processing is needed. One instance of this can occur in a 6LSA node where there is one incoming port and one outgoing and there is only one FEC such as is found commonly in the edge networks or other small or campus network nodes.

17. FEC Mapping

Each FEC may map to a set of flows, node and route characteristics

which may be represented in the switching table. The switching table may consist of more than one entry that a particular FEC can be mapped to and forwarded via a labeled packet.

18. Invalid Incoming Labels

An incoming or acquired label is invalid if it has a value that does not allow the 6LSA node to bind the label to a FEC. Such a label may be discarded after the lead packet is forwarded. Invalid labels may not include a zero-labeled packet.

19. Flow Aggregation or Merging

If it can be determined that two or more flows are destined for the same network or non-6LSA domain, then flow merging are implemented. The advantages are: less state is maintained in each 6LSR and there is no need to differentiate between individual flows after the merge point.

The 6LSA allows aggregation of labels when FECs represent address prefixes. Since IPv6 address prefixes are aggregatable, aggregation of FECs corresponding to aggregatable prefixes is allowed in the 6LSA. The extent of aggregation is a function of the address aggregation, granularity of service desired or both. Such aggregation may further be decided by the IPv6 packet header Traffic Class parameters.

20. Label Encodings

The 6LSA allows encoding of the label value in layer 2 protocols such as in ATM packet's VPI/VCI fields. Since only one label is used and that each such label is uniquely identifiable in the 6LSA, encoding the label in the ATM VPI/VCI field is feasible. Considerations with respect to how flows are identified, the FEC-based forwarding treatment, and flow merging issues, need careful planning in the layer 2 label encoding.

How a 6LSA label value is encoded in the ATM VPI/VCI field is outside the scope of this document.

21. Anycast in 6LSA

IPv6 defines the anycast address like a regular unicast address with a prefix specifying the subnet and an identifier that is set to all

zeroes. Anycast packets delivered to this address are delivered to one router in that subnet. There are reserved subnet anycast addresses such as for mobile IPv6 Home-Agents anycast. The 6LSA allows the use of anycast addressing. Whenever a 6LSR is a node in any anycast subnet, such a subnet may be a 6LSA, a subset of 6LSA or some other part of 6LSA.

When an anycast packet arrives in anycast subnet 6LSR where the subnet is a part or whole of 6SLA, the 6LSR binds the packet to the appropriate FEC which has anycast routing as part of the forwarding treatment attributes of the FEC. The packet is thus forwarded to a next-hop 6LSR through an interface determined by the FEC attributes related to anycast forwarding.

22. Multicast in 6LSA

IPv6 defines the multicast address by the high-order octet FF or 11111111 in binary notation and 4 bits for the scope of the multicast and an identifier bit that indicates whether the multicast address belongs to a well-known IANA multicast address group or is a temporary address.

The 6LSA allows the use of multicast addressing. A multicast tree may be a 6LSA, or a subset of 6LSA. For multicast transmission, the 6LSR binds the packet to a FEC which may represent multicast routing. The packet is thus forwarded to a next-hop 6LSR through an interface determined by the FEC attributes related to multicast delivery. See Section 6.3 above.

23. Security Considerations

The 6SLA allows Security Association (SA). If the security association partners are outside the 6SLA, then there is no effect on the 6SLA by the SA whether the mode of operation is in the transport mode or in the tunnel mode.

In the transport mode of SA, only the packet payload is subject to encryption or authentication, so the IPv6 packet header features are not affected and the 6LSA being a transport mechanism that sets up 6LSPs and provides specific FEC-driven forwarding treatment, there is no impact on the 6LSA or impact on SA operation by the 6SLA.

In the tunnel mode of SA, the SA requires an outer wrapper IPv6 packet. The sending gateway wraps the whole IPv6 packet including the content. The receiving gateway performs the checksum on the outer wrapper packet and then unwraps the packet and verifies the

checksum of the inner packet through end-to-end SA. If the outer wrapper packet conveys the Flow Label value of the inner packet, then 6SLA provides the 6LSP transport based on the inner label value, otherwise the transport indicates the outer label value. Here also, there is no impact on the 6LSA based transport of the secure packets or vice versa.

The Authentication Header (AH) is used in IPv6 for authentication of individual packets to prevent common Internet-based attacks such as IP address spoofing and session hijacking. The computation of cryptographically secure checksum over the payload as well as some fields of the IPv6 and extension headers has to take place between the SA partners. This computation does not include the Flow Label field in the packet header. This maintains label transparency in the 6SLA. Authentication can be either in the transport mode or in the tunnel mode.

The 6SLA security considerations that apply to Encrypted Security Payload (ESP) header comprise encryption modes that are categorized as transport mode or tunnel mode. In the transport mode, no encryption of the Flow Label field is performed, so the value is carried through the 6SLA. In the tunnel mode, the issues are the same as stated here above.

24. Disclaimer

Any affiliation the first two authors have with The MITRE Corporation is provided for identification purposes only, and is not intended to convey or imply MITRE's concurrence with, or support for, the positions, opinions or viewpoints expressed by the authors.

25. Informative Referneces

[1] Rajahalme, J., Conta, A., Carpenter, B., and S. Deering, "IPv6 Flow Label Specification", RFC 3697, March 2004.

[2] Rosen, E., Viswanathan, A., and R. Callon, "Multiprotocol Label Switching Architecture", RFC 3031, January 2001.

[3] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 2460, December 1998.

[4] Hinden, R. and S. Deering, "IPv6 Multicast Address Assignments", RFC 2375, July 1998.

[5] Awduche, D., Malcolm, J., Agogbua, J., O'Dell, M., and J.

McManus, "Requirements for Traffic Engineering Over MPLS", RFC 2702, September 1999.

[6] Awduche, D., Chiu, A., Elwalid, A., Widjaja, I., and X. Xiao, "Overview and Principles of Internet Traffic Engineering", RFC 3272, May 2002.

[7] Agarwal, P. and B. Akyol, "Time To Live (TTL) Processing in Multi-Protocol Label Switching (MPLS) Networks", RFC 3443, January 2003.

[8] Hinden, R. and S. Deering, "Internet Protocol Version 6 (IPv6) Addressing Architecture", RFC 3513, April 2003.

Authors' Addresses

Sham Chakravorty
The MITRE Corporation
7515 Colshire Dr.
McLean, VA 22102
USA

Email: schakra@mitre.org

Jeff Bush
The MITRE Corporation
7515 Colshire Dr.
McLean, VA 22102
USA

Email: jbush@mitre.org

Jim Bound
NAv6TF
PO Box 570
Hollis, NH 03049
USA

Email: Jim.Bound@ipv6forum.com

Full Copyright Statement

Intellectual Property