Network Working Group
Internet Draft: 03
Category: Standards Track

M. Beckman
U.S. Department of Defense
22 February 2007

IPv6 Dynamic Flow Label Switching (FLS)
draft-martinbeckman-ietf-ipv6-fls-ipv6flowswitching-03.txt

Status of this Memo

Copyright Notice

Abstract

This document seeks to establish a standard for the utilization of the Class of Service Field and the us of the Flow Label Field within the IPv6 Header and establish a methodology of switching packets through routers using the first 32-bits of the IPv6 header using Flow Label Switching on packets rather than full routing of packets. Within the first 32-bits of an IPv6 header exists the requisite information to allow for the immediate switching on an ingress packet of a router, allowing for

Label Switching of a native IPv6 packet. This allows the establishment of VPN circuits in a dynamic manner across transit networks. The establishment of Flows based upon the 20-bit Flow Label value can be done dynamically with minimal effort and configuration of the end-point routers of the flow. The flows can be managed or open, encrypted or in the clear, and will allow for greater scalability, security, and agility in the management and operation of networks.

Comments are solicited and should be addressed to martin.beckman@disa.mil

Table of Contents

1. Introduction and Abstract

To traverse the Internet or any large enterprise network, each router hop represents a decision point about the life cycle of each datagram. A major latency inducing function is the look-up of the destination of the packet in the routing table of each router along the way. This is for simplistic routing. If there are additional considerations, such as queuing or filtering, the process can become more laborious. Additionally, two or more networks requiring secure communications require the establishment of a VPN tunnel to assure security of the traffic as it traverses the backbone or in most cases,a carriers internetworking autonomous system. In all cases, the entire IPv6 header of 320 bits must be read, cached, and

processed at each router along the path etween the networks. What is proposed is a methodology of determining the destination port for a packet at is enters a router within the first 32-bits of information. This can be done using a hierarchical methodology of applying values to the Traffic Class Field (8-bits) and switching the packet based upon he value of the Flow Label Field (20-bits) based upon a flow label switching table within the router. The only requirement is that all routers along the paths available can read the Traffic Class Field and are capable of Flow Label Switching.

The Flow Switch Path is dynamically established by the two end-point routers with simple recognition of the flow by the intervening Next-Hop-Routers along the paths between the two End Point Routers. The flows are capable of being controlled either manually or through a Flow Label Server within an autonomous system. This is essential for the secure functioning of a network or conflicting Flow Labels will result. Finally, the Flow establishment and operation is encrypt-able, allowing for secure establishment and operation between the two end point routers of the flow.

Succinctly put, packets can be switched based upon Flow Label Value allowing for a myriad of possibilities in both topologies and secure network operations across carriers across the globe. The end result is a limiting of the need for VPN servers, IPv6 tunnels, and greater mobility of entire networks within an enterprise if proper planning and considerations are understood. Since the packet remains "IPv6 native" the ability to monitor and secure traffic becomes less problematic compared to label switching" within the MPLS context. Instead of converting and non-native IPv6 packet in MPLS form for read and analysis, the packet is handled as any other packet on the network. This is critical when networks use IP/IPv6 packet encryption since an MPLS packet is neither IP or IPv6 and cannot be handled by the encryption device with removing the MPLS shim and thereby wrecking the overall end-to-end secure transmission process.

2. Definitions and Conventions used in this document

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [7].

Flow Initiating Router (FIR) The FIR is the router that initiates a flow label switch path. The FIR sets the Traffic Class Field and Flow Label Values to the required value to set the flow up across the routing fabric between the two end points.

Flow Destination Router (FDR) The FDR is the router the FIR seeks to establish a flow with. The FDR resets the Traffic Class Field and Flow Label Values to the required value to send the packet to its final destination based upon the path determined by the local routing table.

Next-Hop-Router (NHR) The NHR established and maintains the Flow Switch Path using a Flow Switch Table that is maintained based upon instructions from the FIR and its own local routing table.

Switched Flow Path (SFP) is the switched path taken by packet across a Routed fabric based upon the value of the Flow Label and, if used, flow set.

Flow Set (FS) a group of flows through a router identified by the FS value in the Traffic Class.

Flow Path Server (FPS) is a physical or virtual host on the network the FIR, FDR, and NHRs use to validate Flow Path setup requests.

3. The Flow Label Switching Traffic Class.

The first requirement is to establish a flow path across a routed fabric based upon a traffic class value within the defined parameters of RFC 2474. RFC 2474 currently defines three pools for traffic class use within IPv6:

Pool	Codepoint space	Assignment Policy
1	xxxxx0	Standards Action
2	xxxx11	EXP/LU
3	xxxx01	EXP/LU (*)

The codepoint space uses the first six bits of the 8 bits in the traffic class field. Flow Label Switching uses the "top half" of the Traffic Class field by setting the first bit to "1". Pool "128" would use codepoint space of 1abcxxyy, where a,b, and c have values as list below. When "c" is set to "0" and DF codepoint space is in use within a routed domain, "xx" are direct mappings of pools 1, 2, and 3 into the traffic class field. The values for "yy" are reserved.

The second requirement is to identify a packet as being flow switched versus routed. To accomplish this, the Traffic Class Field is used.

In any event the packet is either routed of flow switched . Therefore, the differentiation is set in the first bit of the Traffic Class Field, which is set to 1 for flow switched. This leaves the lower half values of the Traffic class (0-127) available of use in routing. The remaining values of the Traffic class Field of a Flow Switched packet are as follows:

version	Traffic Class	Flow Label
1 2 3 4	1 2 3 4 5 6 7 8	20 bits
0 1 1 0	1 a b c d e f g	1 - 1,048,574

Value a 0 = Open / 1 = Managed
Value b 0 = Clear / 1 = Encrypted
Value c 0 = Data Traffic / 1 = Flow Management Message
Values d through f are dependant upon the value of c .

Note: When "c" is set to "0" and RFC 2474 is in use, pools 1, 2, or 3 are manipulated per RFC 2474 and RFC 3168; therefore, the FIR and FDR map d through g directly into the Traffic Class field.

Pool 128 has a codepoint value of" 1dddxx with an assignment policy of Flow Label Switching where "d" is the defined value per this document and "x" is the value defined in RFC 2474. Pool 128 has a range of 128 to 255.

When the fourth bit (c) is set to "0" the packet is user traffic moving across the flow. The balance of 4 bits is used for priority, differentiating between inter-AS or intra AS Flows, or a combination of both when RFC 2474 Differentiated Service (DS) and RFC 3168, Explicit Congestion Notification (ECN) is not in use. This allows for 16 priorities, sixteen different set of flows, or a combination of differing flow sets with internal priority queues.

When it is a combination of both, priority is set first, and the flow set is set second. As an example, two flow sets (blue and red) are set in field g with blue or red being a value of 0 and the other a value of 1 .
Each flow set then has 3 bits for setting priority using d f . As a cautionary note; by not following RFC 2474 and RFC 3168, Explicit Congestion Notification cannot be used.

When c = 1, the packet is a Flow Management Packet between the two end point routers (the FIR and FDR) as well as for the intervening NHR s along the flow path. The follow are the Values of d though g in this circumstance and are covered in the mechanics of setting of a Flow Switch Path:

d e f g	Decimal	Purpose
0 0 0 0	0	Set up an Asymmetric Flow
0 0 0 1	1	Set up a Symmetric Flow
0 0 1 0	2	NHR Acknowledgment
0 0 1 1	3	NHR Failed
0 1 0 0	4	Restart Flow
0 1 0 1	5	Keep Alive from FIR
0 1 1 0	6	Keep Alive from FDR
0 1 1 1	7	Flow Tear Down
1 n n 0	8-14	FPS Management
1 1 1 1	15	Reserved

4. Flow Label Switching Setup and Management

Across a routed fabric, a switched flow is initiated by a Flow Initiation Router (FIR). To accomplish this, the router has a virtual interface established with a routable 128-bit Unicast address. The Flow Destination Router has the same setup with a different routable 128-bit Unicast address. The initiating packet from the FIR to the FDR is as follows:

version	Traffic Class	Flow Label
1 2 3 4	1 2 3 4 5 6 7 8	20 bits
0 1 1 0	1 0 0 1 0 0 0 0	0-FE

Payload Length	Next Hdr 59	Hop Limit

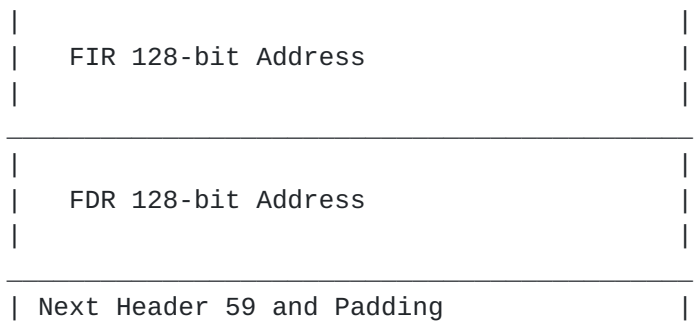
FIR 128-bit Address
FDR 128-bit Address
Next Header 59 and Padding

This establishes a simple, asymmetric Flow Path. The FIR send the packet via the destination port of the FDR based upon the route listed in the routing table.

The FIR then sets the flow label value with the end-points into a flow switch table and marks the label as the router being an end-point for the flow. The Next Hop Router (NHR) receives the packet and established an entry in the flow switch table based upon the routing table as port to the FIR and FDR associated with the flow label. Since this flow in asymmetric, the ports used by the flow path could be dissimilar is the best paths per the routing table have an asymmetric pattern. This is possible for Flows over ASN s where BGP parameters may make ingress and egress to another AS asymmetric. For Symmetric flows, bit 5 is set to one, and the NHR simply duplicates the Flow Switch Table Entry reversing the ingress/egress ports for the flow label association. Once the flow switch table is updated by the NHR, the packet is sent to the next NHR on the routed path, each updating its own Flow Switch Table. The NHR then sends an acknowledgement to the sending router with a TC field of:

1 n n 1 0 0 1 0, where n is the value of the TC filed received.

This is of importance later when flows are setup as managed, with or without encryption. The receiving this acknowledgement then marks the Flow Switch Table entries as active. This process through the NHR s continues until the packet is received by the FDR. Since the destination address is local to the router, the FDR then sets the flow label value with the end-points into a flow switch table and marks the label as the router being an end-point for the flow. The FDR then sends a keep-alive to the FIR with a TC value of 1 n n 1 0 1 1 0 via the flow path established.

The FIR will send a keep-alive with a TC value of 1 n n 1 0 1 0 1. Both the FIR and FDR will send their respective keep-alive packets over the flow path on a varying interval of 1-180 seconds. If the end point routers do not receive a keep-alive from their respective end-point, the FIR and/or FDR will send a restart message using a TC Value of:

1 n n 1 0 1 0 0.

This initiates the Flow over the NHR path. The purpose of the restart message is to force the NHRs on the path to revalidate the Flow Switch table entry for that particular flow. During the startup phase of the flow. If there is a duplicate flow label entry in an NHR along the path (Example: The Network Administrator attempts to use the same flow label values for two different sets of end points, that NHR sends back a NHR Fail message with a TC value of 1 n n 1 0 0 1 1. Any Reviving NHR then drops that entry from the flow switch table and forwards the messages back to the FIR. The FIR then logs to console and drops the flow setup. The Flow Switch Table entries for Next Hop Routers (NHRs) remains valid for 1 to 30 minutes if there are no packets matching the entry. The purpose for this control is to purge unused flow paths from the routed path automatically; however, care should be taken to ensure the FIR/FDR Keep-Alive messages transpire within the purge time set.

5. Managed Flow Label Switching

In the proceeding section, the flows are openly established from one FIR to and FDR with automatic processing by the intervening NHRs along the routed path. While convenient and possibly applicable within a large enterprise network, the management of possibly over 1 million flows will become problematic. Further, while Flow Label Switching is generally for routers, flows could conceivably be established between hosts on the network for a variety of purposes such as server-to-server updating and archiving, true peer-to-peer networking where latency of service is problematic; however, the openness of open flow label switching allows for greater risks to the routed infrastructure. To mitigate these risks and allow for more centralized management, the second bit of the TC filed can be set to one making the establishment of Flow Switch Paths centrally controlled.

As a methodology, Managed Flow Switching is simple. The second bit of the TC field is set to 1. Caching the packet, the receiving router then and requests a validation of the Flow Path from a flow path server (FPS) on the network. Multiple Flow Path Servers (FPS) are required for redundancy.

The recommended methodology would to imbed the server as an internal service on a set of routers within the infrastructure with a common 128-bit Anycast address for the server.

The transaction for setup should be simplistic and allow for secure means of authentication between the routers and the FPS devices on the network.

The conceptual transaction methodology is as follows:

- A Flow Path Server is established on the network with a predetermined Anycast Address available to only the routers or specified devices on the network.

- Each router in the fabric has the Anycast address loaded in the configuration to request a Flow Path Lookup. Additionally, each router should be configurable to globally deny non-managed Flow Path Switching request, yet have the option of permitting individual

- A Flow Path is loaded into the server with the Flow Label, Flow Set, Priority, FIR Unicast Address, and the FDR Unicast Address.

- The Flow Label with Flow Set, Priority, and FDR Address are setup in the FIR.

- The FIR requests validation of the Flow Path from the FPS.

- Once the FPS validates the Flow requested by the FIR and responds with an acknowledgement, the NHR sends the set packet to the next NHR on the Flow Path per the routing table.

- Caching the packet, the first NHR then and requests a validation of the Flow Path from a flow path server (FPS) on the network. When the Flow is validated, the request is forwarded to the next NHR on the path per the local route table. Each NHR responds with an acknowledgement to the requesting router as in the unmanaged flow operation.

- The process repeats through the chain of NHRs until the request is received by the FDR. Caching the packet, the FDR then and requests a validation of the Flow Path from a flow path server (FPS) on the network.
  Once acknowledged, the FDR has acknowledgement, it sends a Keep-Alive
  to the FIR as in the unmanaged flow.

Once the Flow Switching Path is established, the FPS is no longer used. The validity on the Flow Switch Path continues to be maintained via keep-alive packets between the endpoint routers and timers on the NHRs along the path.

Inter-FPS updating for multiple FPS on a routed fabric is essential when using Anycasting. Each FPS will belong to a hierarchy of servers, with one being designated as the root server in a fashion similar to DNS; however, the exchange need to take place via TCP in a point-to-point fashion. If a flow is configured into a secondary server, the root server is notified.
In the event of a root server failure, the next server will assume the role as root server. The recommended approach is to prioritize based upon lowest MAC address or unicast end-station address or the servers.
Since updates are not immediate, A Flow Path Validation request will query the closest FPS per Anycasting methodologies. If the Flow is not found, the FPS queries the root server for an update. If not found the validation fails, yet if the root FPS has the entry, is sends a validation to the secondary server. The secondary server then updates its Flow Path Database.
The root FPS will send an initial full database update to the secondary FPS and will only send adds and drop on a periodic basis after that. If a new secondary FPS is placed into the service, the root server must be manually configured with the address on the secondary server s unicast address. The root FPS will then send the full database to the secondary FPS. A secondary FPS will not request and update. This precludes a rouge FPS from hijacking the FPS database.

The FPS database will identify the following:
- Current Root FPS by Unicast Address

- All Secondary FPS by Unicast Address
- All Flow Path Entries including FIR by Unicast Address, FDR by Unicast Address, Flow Label Value, Flow Set Value (If used), Flow Priority (If Used), Encryption TC bit setting, Flow Symmetry Value, Time Last Keep-Alive received from FIR, keep alive interval.

The root FPS sends a Keep-Alive Query to the FIR and FDR for each flow. The FIR and FDR each respond to their respective Anycast FPS. If an FPS has not received an Acknowledgement from the End-Points within three attempts, the FPS updates is local database and sends a Flow Failure message to the root FPS. The root server takes three actions: Updates the local database by suspending the Flow Path Information, Sends an FPS Database Update to each secondary FPS, Sends a Flow Halt Message to the End-points, The FIR in turn issues a Flow Tear Down Packet to the NHRs to clear the entry from the FIR, FDR, and NHR local Flow Switch Table. The ollowing is a summary of the second half of the TC field binary settings sed with the 11n1 set first half of the TC.

Table Summary of the second half of the TC field binary settings used with the 11n1 set first half of the TC.

d e f g	Decimal	Purpose
1 0 0 0	8	End Point Keep-Alive Query to FIR/FDR
1 0 0 1	9	End Point Keep-Alive Acknowledgement from FIR/FDR
1 0 1 0	10	Flow Halt, Issue Flow Teardown Message
1 0 1 1	11	FPS Full Database Update (from root FPS to secondary FPS)
1 1 0 0	12	FPS Full Database Ack (from secondary FPS to root FPS)
1 1 0 1	13	FPS Database Update (from root FPS to secondary FPS)
1 1 1 0	14	FPS Database Ack (from secondary FPS to root FPS)
1 1 1 1	15	Flow Failure from secondary FPS to root FPS

6. Encrypted Flow Label Switching

The envisioned use of Flow Label Switching is to allow communities of interest connected to a common infrastructure to connect internally to each other without the overhead associated with tunneling or VPN arrangements; however, the Flows need to be secure from monitoring in some cases, as the packets traverse a common backbone or carrier level Autonomous System. This section deals with purpose and use of the third (3rd) bit of the TC Field for encrypting the Flows between Endpoints via either locally agreeable encryption between the endpoint routers (or hosts of the Flows are between Servers, or via a PPKI infrastructure setup.

To encrypt a Flow Path, the FIR sets the third bit of the TC field to a value of one (1). There are two possible methodologies: In the Clear Setup and Management with Encrypted Traffic or Complete Encryption. There are also two levels of Encryption: First 32-bit in the clear and the Entire IPv6 Header in the Clear. In all cases, this is not to be confused with IPv6 security and authentication headers! That is a separate function performed by the end station hosts traversing the network and is functionally performed after the actual IPv6 header is read. In this context, only the first 32-bits of the header are being read to determine a switching decision.

6.a. Encryption Methodologies

In the Clear Setup and Management with Traffic Encryption, while less scure, has logically less overhead for the intervening NHRs along the Flow Switch Path.

In this case, all Flow Setup and management is (fourth bit of the TC filed is set to one) done as previously described, except that the third bit of the TC field is set to one. Once the Flow Switch path is established between the two endpoints, the FIR and FDR exchange keys or perform another authentication and encryption algorithm. The FIR and FDR then encrypt all Traffic traversing the Flow Switch Path at either a high level or a low level. Simplistically, the transmitter encrypts and the receiver decrypts.

Complete Encryption is far more extensive in that all participating routers and, in the case of Managed Flow Label Switching, the Flow Path Servers (FPS)Encrypt all traffic, after the first 32-bits of the header.

In this case the unicast addresses of the end-points of the Flow Switch Path are hidden from view by traffic monitoring. Problematic to this is having all of the routers (as well as possibly hosts) and participating FPS devices encrypting and decrypting all Flow Label Management Packets.

This will increase processor overhead as well as add to the complexity of what is meant to be a simplistic, yet dynamic switching protocol; however, the actual traffic traversing the flow switch path only encrypted and decrypted by the end-point routers of the Flow Switch Path.

6.b. Encryption Levels

In this context, the level of encryption corresponds depth within the packet that the encryption takes place and the type of encryption (IE: Strong or Weak). The Encryption Algorithm determines the strength, the level determines how much of the header and packet is encrypted. The level is determined as part of the exchange between the end-point routers on the Flow Switch Path.

The difference between High level and Low level is that High level encryption scrambles all information after the Hop-Limit Field in the IPv6 packet, making the destination and source addresses as well as the type and content of the datagram unreadable as it passes through the NHR fabric. Low level Encryption scrambles all data after the source and destination address. This allows the destination and source addresses as well as the next header field to be monitored as the packet traverses the NHRs on the Flow Switch path.

7. Flow Sets and Queuing

Once a Flow Switching path is established, the end-points of the flow will have a TC value of: 1 m n 0 a b c d, where m = managed/open, n = encrypted/clear, and the fourth bit is set to 1. The remaining four bits (0-F) can be parsed for two uses: Flow Set Identification or Flow Priority. This feature is to allow equal flow values to be shared on a set of NHRs by differentiating them through a Flow Setvalue similar to concept of an ATM Virtual Path Identifier differentiating equal value for Virtual Circuit Identifiers (VCIs). Alternatively, the 16-bits can be

used
to prioritize which flow has priority on the routers switching based upon
Flow Value. Conceivably, a Next Hop Router in a large Transit Network with
multiple flows may receive Flow Switched packets on several ports over a
brief interval of time. This allows the switching function of the router
to queue the traffic based upon the value set in the 16 bits as the
priority level. In this case, each flow has 16 priority levels of traffic,
allowing a differentiation of latency sensitive traffic versus generic
best effort traffic. Finally, the combination of the two methodologies.
Flow Sets can be determine in the first one to three bits leaving the
remainder for Priority queuing of traffic. Alternatively, the first 1 to 3
bits can determine priority allowing for equal priority flow sets to be
established.

8. Contextual Uses and Security Considerations of Flow Label Switching

There exist two functional advantages for Flow Label Switching versus continuing with MPLS.

First, it affords an alternative to MPLS by establishing VPN circuits between to remote routers. This alternative, unlike MPLS, is dynamic and sets up flows across a routed fabric without having to reconfigure the intervening routers. Second, it allows for faster determination of a packets destiny as it ingresses into a router without resorting to mutating the IPv6 packet by adding a shim. Rather than read the entire 320-bit packet header and executing a closest match route lookup, only the first 32 bits are read and the packet is switched to an egress port, sending the packet on its way with 90% less effort in what to read to determine what to do.

Both these facets allow for some interesting capabilities for aggregation of geographically separate locations behind a single DMZ structure. Since each end-point sends and receives packets based upon Flow Label Value, forming an adjacency is formed between the two virtual Flow Label Interfaces, allowing the flow to act similar to a tunnel across a Wide Area Network. Router A sees Router B directly through their respective Flow Interfaces, allowing either A or B to act as the overall gateway for the other network.

This can extremely effective for large organizations such as the Government or Corporations who have internal organizations that each operate on differing security policies. In this context, each internal organization can be wrapped into a single security domain with a simplifying restructuring of the DMZ. This mitigates the need for VPN servers in numerous cases, and due to the dynamic setup nature of both Clear and Managed Flow Switching Paths, the mobility of entire networks can be readily achieved.

Unlike MPLS, Flow Label Switching operates within the IPv6 protocol s defined header specification. More succinctly put, the IPv6 packet may have the values of the Traffic Class and Flow Label fields manipulated, but it stills remains a native IPv6 packet, unlike MPLS which as a 32-bit shim. This is critical for government use when the data flow must traverse the newer generation of High Assurance IP Encryptor (HAIPE) devices used within US Department of Defense and elsewhere in the US Government. As stated in the name of the device; it is an IP Encryptor and not an MPLS Encryptor! MPLS poses difficult problems for this family of encryption devices currently being deployed as a replacement for link layer encryption devices.

Returning to the concept of non-encapsulated tunneling, FLS paths are established using the routing tables of the routers along the path. This allows for a far more rapid fielding of flows across a routed infrastructure when compared to the implementation of MPLS. Since the flow is established between two virtual interfaces (similar to tunnel interfaces) the virtual interfaces establish link local address connectivity at layer 3 (via the FEC0::/16) routing between these two virtual interfaces is as easily achieved as it is using standard tunneling. As a practical matter, the implementation of this protocol within a routing OS should be as a subset of tunneling protocols, where the tunnel interface number may be equal to the flow label value. The ramifications of this enhancement go directly to the simplification of network operations for service providers and the reduction of costs for connectivity between geographically diverse locations within an enterprise. The following are two uses for this capacity:

Military/Government: Within the Defense Community, the major services (Army, Nay, Air Force, and Marine Corps) as well as the Joint Unified Commands and the various Defense Agencies are widely dispersed throughout the globe. Each of these various entities maintain unclassified interconnectivity via the DoD ISP NIPRNet . Since each one of these entities maintains their own security policies, each entity insists that their external traffic all originate from behind a consolidated DMZ structure. FLS simplifies this critical issue by providing secured flows between the sites to a specific DMZ. Additionally, each flow may be encrypted to where only the first 64 bits of the header are in the clear. This permits the destination and source addresses within the flow as well as the data to be hidden while the packet is switched through a common routed infrastructure to somewhere else within the enterprise s security domain. Finally, the military moves and deploys routinely. FLS permits for additional flows to be established on the fly for those deploying units permitting simplified and continuous connectivity to all domains required. This has considerable tactical, operational, and strategic value!

Commercial/Corporate: As an example, a large manufacturing corporation has numerous production facilities throughout the globe and providing secure and monitored access becomes both costly and problematic. Each site need only achieve IPv6 network access with FLS provided as a service. Each site then can be folded logically and virtually behind a single DMZ and have secure capability between sites. The sites no longer need numerous circuits enmeshing them with each other for a substantial reduction in recurring operational costs.

9. References

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

[RFC 3697] J. Rajahalme, Nokia; A. Conta, Transwitch; B. Carpenter, IBM S. Deering, Cisco; IPv6 Flow Label Specification , RFC 3697, March 2004.

[RFC 3595] B. Wijnen, Lucent Technologies; Textual Conventions for IPv6 Flow Label , RFC 3595, September 2003.

[RFC 3168] K. Ramakrishnan, TeraOptic Networks; S. Floyd, ACIRI; D. Black, EMC The Addition of Explicit Congestion Notification (ECN) to IP , RFC 3168 September 2001.

[RFC 2774 K. Nichols, Cisco Systems; S. Blake, Torrent Networking Technologies; F. Baker, Cisco Systems; D. Black, EMC Corporation, Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers , December 1998.

[RFC 3168] K. Ramakrishnan, TeraOptic Networks; S. Floyd, ACIRI; D. Black, EMC; The Addition of Explicit Congestion Notification (ECN) to IP , September 2001.

10. Acknowledgments

My thanks to Brian Carpenter (brc@zurich.ibm.com) for redirecting my efforts to ensure that inclusion of DS Field definition per RFC 2474 was properly addressed
and patiently reviewing the details.

11. Intellectual Property Statement

distribute working documents as Internet-Drafts.
Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

Individual Property Rights

By submitting this Internet-Draft, each author represents that any applicable Patent or other IPR claims of which he or she is aware have been or will be disclosed, and any of which he or she becomes aware will be disclosed, in accordance with Section 6 of BCP 79.

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 and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

12. Author's Address

Martin Beckman
Defense Information Systems Agency
5275 Leesburg Pike, 7 Skyline Place
Falls Church, VA 22041
United States of America

Phone: 703-861-6865 // 703-882-0225
EMail: martin.beckman@disa.mil