Insert new clause 44 as shown:

44. Equal Cost Multiple Paths

44.1 SPBM ECMP

SPB supports two approaches to spreading traffic load. The first approach, ECT VLANs, spreads traffic by assigning each service instance to one of the supported SPT Sets. Each SPT Set uses a different equal cost tie-breaking algorithm and thus may select different shortest paths in the SPT Region. This approach is applicable to both SPBV and SPBM. The Base VID is assigned to an ECT-ALGORITHM (tie-breaker) that defines the SPT Set for that VLAN. In case of a PBBN, each I-SID is assigned to a Base VID. This approach can spread traffic load while maintaining traffic path congruence in each VLAN.

The second approach, ECMP, spreads traffic by employing multiple equal cost paths. This approach is only applicable to SPBM. Unicast frame filtering selects dynamically from the set of equal cost next hops for the frame's destination MAC address. Multicast frame filtering entries are generated by calculations that can distribute the multicast trees over a set of equal cost trees in the SPT Region.

44.1.1 Equal Cost Multiple Paths Operation

SPBM equal cost multiple paths (ECMP) provides a mechanism to spread traffic load across a set of equal cost shortest paths in an SPT Region. Rather than using an equal cost tie-breaking function to create a single SPT for all individual addresses, ECMP uses a hash function to spread individual address forwarding over a set of equal cost forwarding ports.

If flow filtering (44.1) is not supported, the ECMP ECT Algorithm selects one port from the set of equal cost forwarding ports for each individual address (28.8.2) and creates a Dynamic Filtering Entry for the address indicating the selected port.

If flow filtering (44.1) is supported, a set of equal cost forwarding ports may be recorded in the FDB in a Dynamic Filtering Entry (8.8.3). During unicast frame filtering the MAC Relay Entity selects one forwarding port from the applicable set using a dynamic port selection process (44.3).

The individual address forwarding port selection processes, both with and without flow filtering support, are designed to

a) spread traffic (i.e., make varied choices) for different flows traversing the SPT Region, and
b) maintain frame ordering for each flow by making the same choice for all frames belonging to any single flow.

ECMP is supported by Dynamic Filtering Entries in the filtering database and, optionally, the Flow Filtering tag (F-TAG, 44.1.1) that carries Flow Hash and TTL information. The flow hash serves to both distinguish frames belonging to different flows and identify frames that belong to (or may belong to) the same flow. The flow hash is used as input to the dynamic port selection process to support the design goals noted above. The TTL is used to mitigate traffic loops, if they happen to occur (44.4).

While unicast frames are spread across equal cost paths, multicast frames cannot be treated this way. Instead multicast traffic load is spread by selecting varied SPTs for each group address or set of group addresses. The default ECMP ECT Algorithm selects a source specific SPT for each node that may source multicast traffic. That source SPT is used for all group addresses sourced from that node. Finer grained multicast load spreading is supported by providing configurable options. These options include:

c) Multiple source specific SPTs per node,
d) Shared trees, and
e) Head-end replication.

Up to 16 source specific SPTs can be supported per node by using 16 different tie-break masks in calculating the SPT. During SPT calculation the tie-break mask is used to select one from the set of equal cost parents (neighbor nodes closer to the root) for each node (28.8.2). The default multicast treatment is equivalent to using a tie-break mask of zero. An I-SID on a CBP can be assigned to a specific SPT by configuring a Tie-Break Mask for that I-SID [item h) in 12.25.8.1.2 and item i) in 28.12.10].

To reduce the amount of multicast forwarding state required, shared (spanning) trees are supported. Up to 16 shared trees can be configured using the 16 tie-break masks. For shared trees the tie-break mask is used both to select a shared tree root Bridge and to select from the set of equal cost parents for each node in the tree. The shared tree root for a given tie-break mask is determined by selecting the best (least) masked Bridge ID (i.e., the Bridge ID value masked using the tie-break mask). The SPT calculation from the shared tree root also uses the tie-break mask to select between equal cost parents (again selecting the least masked Bridge ID). Each I-SID on a CBP can be assigned to a shared tree by setting a shared tree transmit indicator bit and configuring the tie-break mask to identify the shared tree (28.12.10). All CBPs for a given I-SID that transmit on a shared tree use the same group address (Figure 27-4), and share the forwarding state for this address in the FDB.

It is possible to avoid using group addresses altogether and handle multicast frames by head-end replication. An I-SID on a CBP using head-end replication sends a unicast frame to each destination instead of sending a group addressed frame. If an I-SID is using head-end replication it is configured to indicate no multicast transmission and no group address forwarding state will be installed for that I-SID (28.8.2). An I-SID on a CBP using head-end replication should indicate that it wishes to receive multicast frames since other CBPs supporting the same I-SID may still use multicast.

NOTE: ISIS-SPB does not check for the use of the same tree type for all endpoints of an I-SID, but use of different tree types within a single I-SID is not recommended.

44.1.2 ECMP ECT Algorithm

The content of FDB entries for individual addresses in ECMP depends on whether or not flow filtering (44.1) is supported by a Bridge.

If flow filtering is not supported or the ECT_ALGORITHM value indicates no flow filtering (Table 28-2), then for each individual address the ECMP ECT Algorithm selects one port from the set of equal cost forwarding ports using a form of the FNV non-cryptographic hash algorithm and creates a Dynamic Filtering Entry (8.8.3) indicating that port as forwarding. The forwarding port is selected by the following process:

For a set of *k* equal cost forwarding Ports, let **P** be an ordered list of the forwarding Ports sorted from greatest to least neighbor SPB System Identifier and P_i be the *i*th Port in **P**. Thus P_0 is the Port whose neighbor has the greatest SPB System Identifier and P_{k-1} is the Port whose neighbor has the least SPB System Identifier. The following algorithm is used to select a value for *i* for the individual MAC address, thus selecting Port P_i to be the forwarding port:

a) Set a 32-bit variable *bestHash32* to 0xFFFFFFFF,
b) Set *i* to 0 (zero),
c) For each *n* in 0..*k*-1 (in order) perform the following steps 1 through 5
    1) Set a 32-bit variable *hash32* to 0x811C9DC5,
    2) For each octet O_j (extended to 32 bits with leading 0's) in the Bridge's SPB System Identifier (8.13.8), starting with the least significant octet (octet 5) and progressing to the most significant

octet (octet 0),
set *hash32* to ((*hash32* xor O_j) × 0x01000193),

3) For each octet O_j (extended to 32 bits with leading 0's) in the neighbor SPB System Identifier for Port P_n, starting with the least significant octet (octet 5) and progressing to the most significant octet (octet 0),
set *hash32* to ((*hash32* xor O_j) × 0x01000193)

4) For each octet O_j (extended to 32 bits with leading 0's) in the individual MAC address, starting with the least significant octet (octet 5) and progressing to the most significant octet (octet 0),
set *hash32* to ((*hash32* xor O_j) × 0x01000193)

5) If *hash32* < *bestHash32* set *i* to *n*.

If flow filtering (44.1) is supported and the ECT_ALGORITHM value indicates flow filtering (Table 28-2), ISIS-SPB configures the Flow Filtering Control Table (44.2) for each port on which the ECMP SPBM VID is registered (27.13). For CBPs the Base VID is configured with both Flow Filtering and Flow Hash Generation enabled. For PNPs the Base VID is configured with Flow Filtering enabled and Flow Hash Generation disabled. For each individual address the ECMP ECT Algorithm calculates a Port Map specifying forwarding for a set of equal cost hops to that address and creates a corresponding Dynamic Filtering Entry (8.8.3). These Dynamic Filtering Entries are used in the dynamic selection of forwarding ports by the flow filtering process (44.3).

For each I-SID on a CBP transmitting multicast frames the ECMP ECT Algorithm may calculate a source rooted SPT. If equal cost paths are found to any node on the tree, the path whose parent node (neighbor node toward the root) has the least Bridge Identifier is selected. A group MAC Address Registration Entry (8.8.4) is created for the I-SID's source rooted tree SPBM Group MAC Address (Figure 27-3) if the node performing the calculation is located between the source node and any multicast receiver(s) for the I-SID. The group MAC Registration Entry indicates forwarding for each port that leads to at least one multicast receiver for the I-SID on the calculated SPT (except the port from the source).

The ECMP ECT Algorithm supports two additional parameters that control multicast forwarding state calculation. First, a Tie-Break Mask value may be specified for an I-SID on a CBP [item i) in 28.12.10 and item h) in 12.25.8.1.2]. When a non-zero Tie-Break Mask is advertised, a 64-bit mask (formed by repeating the 4-bit Tie-Break Mask value in each nibble) is applied to (xor'ed with) each Bridge Identifier before selecting the lowest value as described in the SPT calculation in the paragraph above. Thus the default behavior described above is equivalent to using a Tie-Break Mask of zero. This parameter enables up to 16 different SPTs to be calculated from a given source node. Each I-SID on a CBP may be assigned independently to one of these 16 source trees.

The second parameter controlling multicast forwarding state calculation is the shared tree indicator. An I-SID on a CBP may be configured to use a shared tree instead of a source specific SPT [item h) in 28.12.10 and item g) in 12.25.8.1.2]. There may be up to 16 shared trees, corresponding to the 16 Tie-Break Mask values. Each shared tree is calculated by selecting a root Bridge using the least masked Bridge Identifier. If a node advertises a Bridge Priority for the given Tie-Break Mask (28.12.6.1) the advertised value is used instead of the normal Bridge Priority in forming its Bridge Identifier (before masking). In calculating the SPT from the selected root the same Tie-Break mask is used to form masked Bridge Identifiers (using the normal Bridge Priority) and the least value of masked Bridge Identifier is used to chose between equal cost parents for any node. If an I-SID on a CBP is configured to use a shared tree, the group address used in the group MAC Address Registration Entry is constructed using the Tie-Break Mask and the I-SID (27.15, Figure 27-4). In this case the ECMP ECT Algorithm creates a group MAC Address Registration Entry if the node performing the calculation is between the CBP and any multicast receivers on the SPT calculated from the selected root node. The Port Map indicates forwarding for all Ports that lead from a multicast transmitter to at least one multicast receiver for the I-SID.

If an I-SID on a CBP does not indicate transmission of multicast frames (e.g., head-end replication is to be used instead) then no group MAC Address Registration Entry is configured for multicast from that CBP for that I-SID.

Table 28-2 shows the ECT-ALGORITHM values for the ECMP ECT Algorithm.

Table 44-2—ECMP ECT-ALGORITHM values

ECT-ALGORITHM	ECMP Behavior
00-80-C2-11	ECMP without flow filtering
00-80-C2-12	ECMP with flow filtering

44.1.3 Loop prevention for ECMP

Loop prevention for unicast ECMP operates as specified in Clause 13 with the extension of allowing multiple Root Ports in case of multiple equal cost paths in order to support the per frame tie-breaking of ECMP. The ECMP connectivity to a destination bridge is comprised of the shortest paths destined to that bridge when the topology is Agreed in the meaning of the Agreement Protocol. If a bridge is not the destination bridge, then the one or more ports that ISIS-SPB has calculated as respectively providing a shortest path to the destination is a Root Port in the ECMP connectivity. Each port, other than a Root Port, is a Designated Port if ISIS-SPB has calculated that the port provides a shortest path for frames forwarded from the attached bridge to the destination in the ECMP connectivity. With these extensions to the definitions of Port Roles, Clause 13 specifications apply to ECMP as well.

Note 1 - Despite of allowing more than one Root Port for an SPT Bridge, a single Root Port is enforced on per frame basis as ECMP operation ensures that a frame is only forwarded through one of the Root Ports towards the destination.

Loop prevention for the shared trees used for ECMP multicast operates as specified by Clause 13 for shared spanning trees except for that instead of a spanning tree protocol entity it is ISIS-SPB that configures the per tree Port States and Roles based on its calculations and the operation of the Agreement Protocol; furthermore, the Learning Process (8.7) neither creates nor deletes Dynamic Filtering Entries.

Note 2 - The per tree Port States and Roles do not belong to an MSTI having a distinct MSTID because an ECMP B-VID is assigned to the SPBM MSTI. The Port States and Roles are associated with a single SPT used as a shared spanning tree supporting bidirectional forwarding (not as a unidirectional tree as for SPB).

44.1.4 Connectivity Fault Management for ECMP

An ECMP path MA is used to continuously test multiple paths between two endpoints. This type of MA has only two MEPs, at two SPT Region boundary ports. The MA is identified by a TE-SID comprising an ESP in each direction between the two endpoints. An ECMP path MEP cycles through a set of flow hash values intended to direct the CCM frames along (and thus test) different paths to the other MEP in the MA. Each flow hash is used in multiple consecutive CCM frames so that a fault on a particular path will be detected and reported back via an RDI signal. ECMP path MEPs are not used to generate LTM or LBM frames, since these can be generated from an SPBM MEP, and therefore there are no MHFs associated with ECMP path MAs.

44.1.4.1 ECMP path MEP placement in a Bridge Port

Since ECMP path MEPs must only see traffic from one other endpoint, these MEPs need to be further differentiated by the TESI Multiplex Entity (6.19). This locates an ECMP path MEP so that it receives

CCMs only from the other MEP in its MA. Figure 44-1 also shows an example of ECMP path MEPs on a CBP.

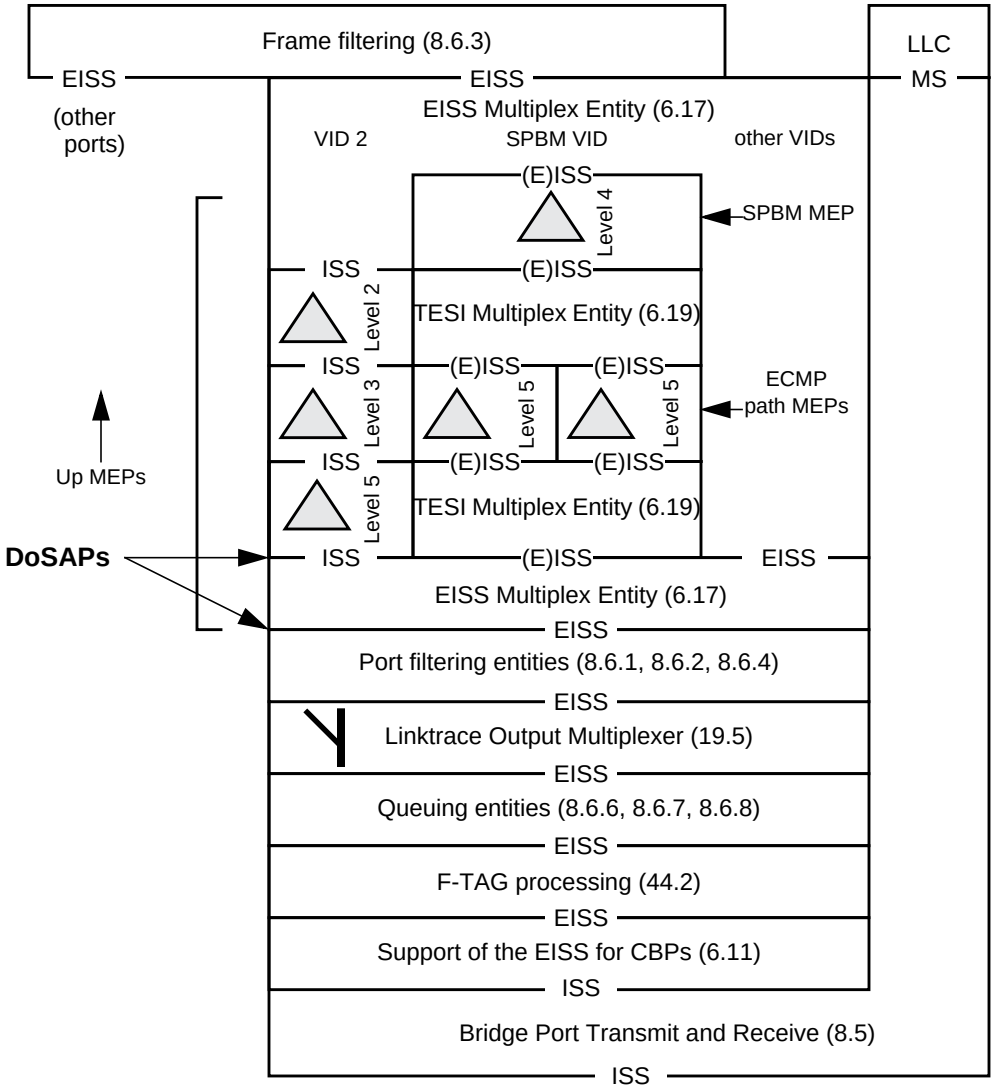

Figure 44-1—SPBM MEP and ECMP path MEP placement in a CBP

44.1.4.2 Continuity Check protocol in an ECMP path MA

The Continuity Check protocol is described in 20.1 and the corresponding state machines in Clause 20. The modifications required to realize ECMP path MAs include:

a) The MEP Continuity Check Initiator state machine adds three variables for controlling ECMP path testing (20.10.4, 20.10.5, 20.10.6) and shares two of these with the Remote MEP state machine (20-1);
b) The destination_address parameter is set to the individual MAC address of the other MEP in the MA [item a) in 20.11.1];
c) The flow_hash parameter is set to one of the values to be tested and the time_to_live parameter is set [item e) in 20.11.1];
d) The MEP Continuity Check initiator state machine is modified to support a cycle of path tests for an ECMP path MEP (20.12);

e) The Remote MEP state machine is modified to track a cycle of path test RDI values and report rMEPlastRDI as set if any RDI has been seen in the last path test cycle for an ECMP path MEP (20.19.2, 20.20).

NOTE—In an ECMP path MEP testing a set of paths, the RDI signal reflects back the state of far end CCM reception. Using the current CCM sending cycle position (pathN), the CCM sending rate, and the path round trip delay one can calculate roughly the position in the sending cycle corresponding to a received RDI. This approach makes minimal changes to the CCM state machines at the CCM sending MEP and no changes at the CCM receiving MEP. It is subject to some ambiguity based on differences in path delay, the sending rate, and cycle length. In the case of ambiguous results further testing would be required to identify a faulty path.

SPBM MAs may also implement the following enhancements to the Continuity Check protocol:

f) The procedure MEPprocessEqualCCM() on ECMP path MEPs does not include the check of the MAID on received CCMs [item b) in 20.17.1].

All other Continuity Check processes are the same as those for a VID-based MA.

44.2 Support for Flow Filtering

Flow filtering support enables Bridges to distinguish frames belonging to different client flows and to use this information in the forwarding process. Information related to client flows may be used at the boundary of an SPT Domain to generate a flow hash value. The flow hash, carried in an F-TAG, serves to distinguish frames belonging to different flows and can be used in the forwarding process to distribute frames over equal cost paths. This provides for finer granularity load spreading while maintaining frame order for each client flow.

Flow filtering behavior is controlled by ISIS-SPB. To allow Bridges that support flow filtering to be used alongside Bridges that do not support flow filtering, ISIS-SPB requires that all Bridges at the SPT Domain boundary that advertise I-SIDs for an ECMP Base VID using flow filtering support F-TAG processing. Bridges that do not advertise I-SIDs for that Base VID are not required to support flow filtering. Therefore network operators are able to upgrade selected bridges to support flow filtering and use flow filtering for selected VIDs without requiring the simultaneous upgrade of all bridges in the SPT Domain.

NOTE—Support for F-TAG processing is indicated in I-SID advertisements and so is only available via ISIS-SPB once I-SIDs are assigned to the Base-VID. A "dummy" ISID configuration can be used to discover via ISIS-SPB whether or not F-TAG processing is supported on a Bridge.

This clause specifies

a) A flow filtering tag (F-TAG) that carries a flow hash value (44.1.1);
b) F-TAG processing on Bridge ports (44.2);
c) A forwarding process using the flow hash to select one forwarding port from a set of equal cost options (44.3);
d) TTL loop mitigation (44.4); and
e) Operation of bridged LANs with selective support for flow filtering (44.5).

44.2.1 Flow filtering tag

A flow filtering tag is specified to carry additional information in support of flow filtering functions. The flow filtering tag comprises a Tag Protocol Identifier (EtherType) and Tag Control Information (TCI)..

The F-TAG TCI field (Figure 9-4) is 4 octets in length and encodes priority, drop_eligible, TTL, and flow hash information for a service primitive.

Table 44-3—F-TAG EtherType

Tag Type	Name	Value
Flow Filtering Tag	IEEE 802.1Q Flow Filtering Tag EtherType	<<TBD>>

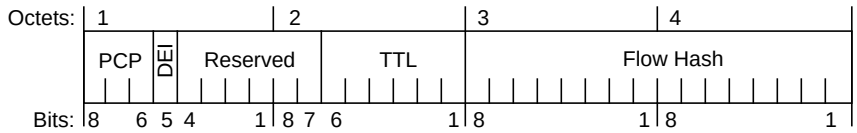

Figure 44-2—Flow FilteringTCI format

The F-TAG TCI contains the following fields:

a) *Priority Code Point (PCP)*—This 3-bit field encodes the priority and drop_eligible parameters of the service request primitive associated with this frame using the same encoding as specified for VLAN tags in 6.9.3.

b) *Drop Eligible Indicator (DEI)*—This 1-bit field carries the drop_eligible parameter of the service request primitive associated with this frame.

c) *Reserved*—This 6-bit field may be used for future format variations. The Reserved field contains a value of zero when the tag is encoded, and is ignored when the tag is decoded.

d) *Time To Live (TTL)*—This 6-bit field holds a counter value that is decremented each time the F-TAG is read (44.2.1) and provides a mechanism to limit the number of network hops a frame may take.

e) *Flow Hash*—This 16-bit field carries a flow hash value used in the dynamic port selection process for ECMP (8.8.3).

44.3 F-TAG processing

This standard specifies flow filtering behavior for SPBM ECMP (27, 27.17, 27.17.2). To support this behavior a flow hash may be conveyed in a mac_service_data_unit using an F-TAG (44.1.1). This clause specifies a Support for Flow Filtering shim that produces and processes F-TAG information. The shim has two EISS SAPs, a lower EISS toward a LAN and an upper EISS toward the MAC Relay Entity.

At the boundary of a flow filtering SPT Domain, for example at Customer Backbone Ports, the flow hash is generated for EM_UNITDATA.indication primitives. At intermediate ports, for example Provider Network Ports, the F-TAG is generated using the flow hash in EM_UNITDATA.request primitives, or processed to retrieve the flow hash value from EM_UNITDATA.indication primitives. The F-TAG also includes a TTL field, used for loop mitigation, and PCP and DEI fields that are processed by this shim.

Flow filtering support is governed by a Flow Filtering Control Table for each port containing an entry for each VID with the following fields:

a) Flow Filtering (enabled/disabled). This field indicates whether or not flow filtering behavior is enabled on the port for the VID. The default value is disabled.

b) Flow Hash Generation (enabled/disabled). This field indicates whether or not flow hash generation is enabled on the port for the VID. The default value is disabled.

c) TTL Value (1..63). This field is the initial TTL value for frames entering the flow filtering SPT Domain [item d) in 12.16.5.5.2]. The default value is 8.

The Flow Filtering Control Table is provisioned by ISIS-SPB (28.8.2). For flow filtering, a CBP will have both the Flow Filtering field and the Flow Hash Generation field set to enabled and the TTL Value field set to provide the initial time_to_live value for indication primitives. For flow filtering, a PNP will have the Flow Filtering field set to enabled and the Flow Hash Generation field set to disabled. The TTL Value field is not used on a PNP.

NOTE 1—Based on the description above, for flow filtering the condition "Flow Hash Generation field indicates enabled" is equivalent to indicating a CBP and the condition "Flow Hash Generation field indicates disabled" is equivalent to indicating a PNP.

44.3.1 Data indications

On receipt of an EM_UNITDATA.indication primitive from the lower EISS on a PNP, the received frame shall be discarded if

a) For the vlan_identifier in the Flow Filtering Control Table the Flow Filtering field indicates enabled and the initial octets of the mac_service_data_unit do not contain a valid F-TAG; or

b) The value in the TTL field in the F-TAG is 0.

NOTE 2—A frame received with a TTL value of 0 will be discarded, i.e., not relayed, even if it is addressed to another port on the receiving Bridge.

Otherwise, an EM_UNITDATA.indication primitive is invoked at the upper EISS with parameter values determined from the received primitive based on the Flow Filtering Control Table for the received vlan_identifier as follows:

If the Flow Filtering field indicates disabled the invoked primitive is identical to the received primitive.

If the Flow Filtering field indicates enabled the **destination_address**, **source_address**, **vlan_identifier**, and **connection_identifier** in the invoked primitive are identical to those parameters in the received primitive.

If the Flow Filtering field indicates enabled and the Flow Hash Generation field indicates disabled, the **mac_service_data_unit** in the invoked primitive is created by removing the F-TAG from the received **mac_service_data_unit**, otherwise the **mac_service_data_unit** in the invoked primitive is identical to the received primitive.

The value of the **flow_hash** and **time_to_live** parameters are determined as follows:

a) If the Flow Hash Generation field indicates enabled the value of flow_hash is generated from the received primitive (using a process selected by the implementation). Time_to_live is set to the TTL Value from the Flow Filtering Control Table for the received vlan_identifier.

NOTE 2—Procedures for calculating the flow hash are not covered by this standard. The procedure is typically a protocol-dependent hash of the client protocol header being carried within the mac_service_data_unit.

NOTE 3—There is often a compromise between the speed of a hash calculation and the quality of value spreading. This standard makes no assurances regarding the quality of the distribution of flow_hash values.

b) If the Flow Hash Generation field indicates disabled the value of flow_hash is taken from the Flow Hash field in the received F-TAG and time_to_live is set to the TTL value in the received F-TAG decremented by 1.

The value of the **drop_eligible** and **priority** parameters are determined as follows:

c) If the **mac_service_data_unit** was tagged with an F-TAG, the value of the drop_eligible parameter and the received priority value are decoded as described in 6.9.3, using the PCP and DEI values from the F-TAG. Otherwise;
d) The received priority value and the drop_eligible parameter value are the values in the EM_UNITDATA.indication.
e) The value of the priority parameter is then regenerated from the received priority, as specified in 6.9.4.

44.3.2 Data requests

On invocation of an EM_UNITDATA.request primitive by a user of the EISS, an EM_UNITDATA.request primitive is invoked, with parameter values determined based on the Flow Filtering Control Table for the received vlan_identifier.

If the Flow Filtering field indicates disabled or the Flow Hash Generation field indicates enabled the invoked primitive is identical to the received primitive.

If the Flow Filtering field indicates enabled and the Flow Hash Generation field indicates disabled all parameters in the invoked primitive match those in the received primitive except the mac_service_data_unit. The mac_service_data_unit in the invoked primitive is formed by inserting an F-TAG as the initial octets to the received mac_service_data_unit. The values of the **flow_hash**, **time_to_live**, **priority**, and **drop_eligible** parameters are used to determine the contents of the F-TAG, in accordance with 44.1.1.

44.4 Forwarding process extension for flow filtering

Frame filtering (*8.6.3*) reduces the set of potential transmission ports in accordance with applicable Filtering Database entries. When flow filtering is enabled the ISIS-SPB control plane creates Dynamic Filtering Entries with multiple forwarding Ports in the Port Map when there are multiple equal cost hops to a given individual address. Frame filtering must select only one of these forwarding ports on which to transmit each frame with a matching destination_address and VID.

Selection of one Port from a set indicated by the Port Map shall be performed using a form of the FNV non-cryptographic hash algorithm as follows:

For a Port Map specifying *k* potential forwarding Ports, let **P** be an ordered list of potential forwarding Ports sorted from greatest to least neighbor SPB System Identifier and P_i be the *i*th Port in **P**. Thus P_0 is the Port whose neighbor has the greatest SPB System Identifier and P_{k-1} is the Port whose neighbor has the least SPB System Identifier. The following dynamic port selection process is used to select a value for *i*, thus selecting Port P_i to be the transmission port:

a) Set a 32-bit variable *hash32* to 0x811C9DC5,
b) For each octet O_j (extended to 32 bits with leading 0's) in the Bridge's SPB System Identifier (8.13.8), starting with the least significant octet (octet 5) and progressing to the most significant octet (octet 0),
set *hash32* to ((*hash32* xor O_j) × 0x01000193),
c) For each octet O_j (extended to 32 bits with leading 0's) in the flow_hash parameter, starting with the least significant octet (octet 4 of Figure 44-1) and progressing to the most significant octet (octet 3 of Figure 44-1),
set *hash32* to ((*hash32* xor O_j) × 0x01000193)
d) Set a 16-bit variable *hash16* to ((*hash32* and 0x0000FFFF) xor (*hash32* / 0x00010000))
e) Set *i* to (*hash16* mod *k*), selecting Port P_i.

44.5 TTL Loop mitigation

The F-TAG contains a TTL field whose value is decremented each time a frame is forwarded by an SPT Bridge supporting flow filtering [item b) in 44.2]. A frame is filtered (discarded) if it is received with a TTL value of 0. This mechanism mitigates the impact of looping frames by limiting the number of hops these frames may travel after returning to any point in the network. This method of loop mitigation protects network resources by ensuring frames are forwarded at most a fixed number of times.

NOTE—The TTL carried in the F-TAG is decremented for all frames; however, the Flow Hash is only applicable to unicast frames. The Flow Hash has no effect on the filtering of group addressed frames.

44.6 Operation with selective support for flow filtering

When adding new capabilities to a network it is desirable to make changes incrementally rather than perform a wholesale upgrade (e.g., incur a "flag day" event). Flow filtering is optional functionality that can be used with selected Base-VIDs while not adversely affecting other services in the network. However, flow filtering changes frame encoding (i.e., adds new header information in an F-TAG) and therefore requires that a Base-VID using ECMP with flow filtering have all its boundary points capable of processsing the F-TAG. To enable incremental incorporation of flow filtering, ISIS-SPB requires F-TAG support for any Bridges with a CBP associated with a Base VID using ECMP with flow filtering while allowing other Bridges in the SPT region to operate without flow filtering support.

When the ECMP ECT-ALGORITHM value indicates support for flow filtering (Table 28-2) an I-SID cannot be configured to use a Base VID using ECMP with flow filtering unless the Bridge supports flow filtering (12.16.5.2). Requiring all Bridges advertising service instances using ECMP with flow filtering to support F-TAG processing provides for proper handling of F-tagged frames at the boundary of the SPT Domain.

Within the SPT Domain it is possible to use Bridges that support the ECMP ECT Algorithm but do not support flow filtering (i.e., do not recognize F-TAGs). In this case, the ECMP ECT Algorithm (28.8.2) selects one port from the set of equal cost forwarding ports for each individual address and creates a Dynamic Filtering Entry for the address indicating the selected port.

NOTE—Bridges that do not support flow filtering should not change PCP and DEI values for frames carrying F-TAGs as this may not provide the desired behavior. If the PCP and DEI values are changed in the VLAN tag only, the PCP and DEI values carried in the F-TAG will be reasserted when the F-TAG is next processed.

These ECMP behaviors enable incremental upgrade of edge bridges to support flow filtering and selective use of this capability with I-SIDs registered only at these upgraded edge bridges. ECMP operation with selective support for flow filtering enables incremental introduction of flow filtering capabilities into an existing network and earlier introduction of flow filtering support for selected services.