



FC-SW-7

Leaf or PSA?

Leafing through the Fabric

Letter Ballot Comment Resolution

Author

Howard L. Johnson (Broadcom)

Dave Peterson (Broadcom)

T11-2019-00015-v001 (as modified)



Table of Contents

- Provide “as modified” text
 - Based on February FC-SW-7 discussion



Proposed Comment Resolution – “as modified” text

6.2.9.3 Link State Header format – Table 33

- Add a “bit” of clarity
 - Define Bit-1 in the LSR Flags field
 - Describes the class-F behavior, which clarifies the purpose of the Leaf Switch
- LSR Flags:
 - Bit-0: Leaf Switch:
 - If this bit is set to one, then the Switch shall be a route termination point (I.e., the Switch is a Leaf Switch). As a result, the Leaf Switch shall not establish routes through this Switch.
 - If this bit is set to zero, then the Switch is not a route termination point.
 - Bits-1: Class F Routing:
 - This bit is valid only if bit 0 is set to one (I.e., the Switch is a Leaf Switch).
 - If this bit is set to one, then the Leaf Switch shall allow all Class F frames to be routed through the Switch.
 - If this bit is set to zero, then the Leaf Switch shall not allow Class F frames to be routed through the Switch.
 - Bits 2-7: Reserved

| Bit | Description |
|--------|---|
| 0 | <p>Leaf Switch:</p> <p>If this bit is set to one, then the Switch shall be a route termination point (I.e., the Switch is a Leaf Switch). As a result, the Leaf Switch shall not establish routes through this Switch.</p> <p>If this bit is set to zero, then the Switch is not a route termination point.</p> |
| 1 | <p>Class F Routing:</p> <p>This bit is valid only if bit 0 is set to one (I.e., the Switch is a Leaf Switch).</p> <p>If this bit is set to one, then the Leaf Switch shall allow all Class F frames to be routed through the Switch.</p> <p>If this bit is set to zero, then the Leaf Switch shall not allow Class F frames to be routed through the Switch.</p> |
| 2 to 7 | Reserved |

Thank You

Howard Johnson

This slide intentionally left blank

References

Notes from February 2019 Meeting

(T11-2019-00015-v000)

Letter Ballot Comment

Cisco #1

- Reference
 - Leaf Switch Text (Clause 18 – T11-2018-00094-v003)
- Comment
 - Cisco-L#1 The Leaf Switch functionality needs more clarification. With Leaf switches, we can end up with a situation where a domain exists but fabric cannot route to it. This case could have implications on several features which require communication with the Domain controller of “other” OR “all” switches.
 - (a) The fabric distributed services having a Work Category of “1-to-All” defined as: “A 1-to-All request is a request that is unable to be handled entirely by the Entry Switch, but for which the Entry Switch is unable to identify the set of remote Switches to query. The Entry Switch sends request frames directly to the Domain Controller of all Switches in the Fabric”
 - (b) Switch Trace Route (STR): “After receiving a FC Trace Route (FTR) request (see FC-GS-8), the Managing Switch shall send a STR request to the Domain Controller of the Switch to which the source port is attached”
 - (c) There seem to be a lot more...

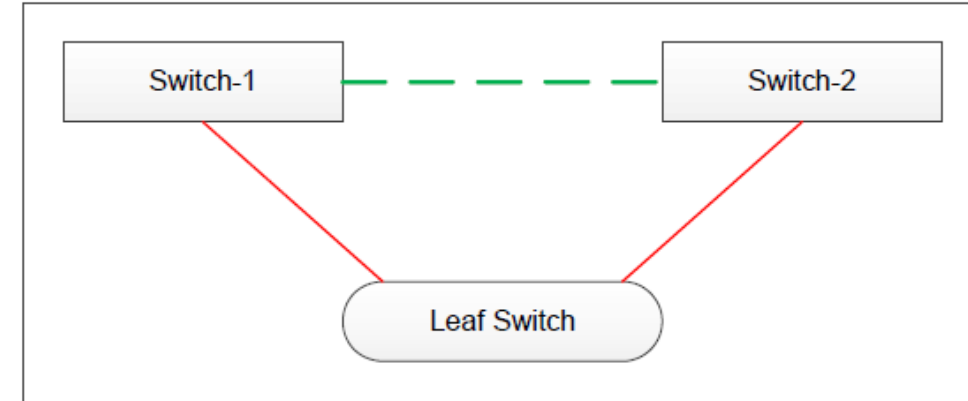


Figure 57 – Leaf Switch example

Intended Behavior

Address class-F traffic

- Distributed Services
 - Class-F vs all other classes of traffic
- Trace Route
 - Um, that's kind of the point 😊
- Others
 - Need a little more guidance here
- Resolution
 - Depends on the purpose of the Leaf Switch
 - Is it an appliance or terminal switch (i.e., PSA or Leaf)?
 - An appliance may not want to “merge” the fabrics
 - The solution is to allow for a range of restricted behavior
 - I.e., explicitly describe how Class-F traffic is affected

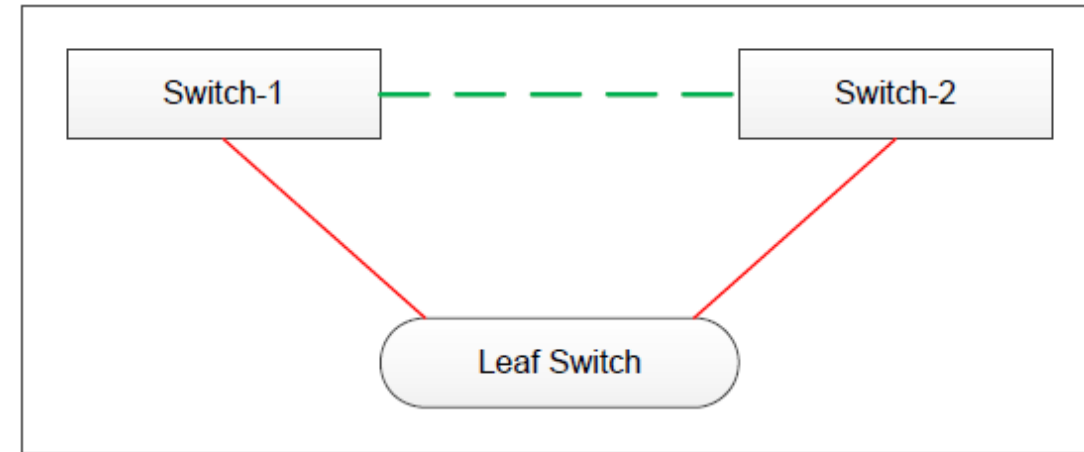


Figure 57 – Leaf Switch example

Proposed Comment Resolution

6.2.9.3 Link State Header format – Table 33

- Add a “bit” of clarity
 - Define Bit-1 in the LSR Flags field
 - Describes the class-F behavior, which clarifies the purpose of the Leaf Switch
- LSR Flags:
 - Bit-0: Leaf Switch
 - The Switch shall be a route termination point
 - Other Switches in the Fabric shall not route through this Switch
 - Bits-1: Class-F pass-through
 - Valid only when bit-0 is set to one
 - The Switch shall route Class-F frames
 - All other classes of service shall not be routed
 - Bits 2-7: Reserved

| Bit | Description |
|--------|--|
| 0 | Leaf Switch: If this bit is set to one, then the Switch shall be a route termination point and other Switches in the Fabric shall not route through this Switch. If this bit is set to zero, then the Switch shall allow other Switches in the Fabric to establish routes through this Switch. |
| 1 | Class-F Pass-through: If this bit is set to one, then the Switch shall allow all Class-F frames to pass-through the Switch unrestricted. If this bit is set to zero, then the Switch should restrict all frames through the Switch. This bit is valid only when Bit-0 is set to one. |
| 2 to 7 | Reserved |

References


Notes from October 2018 Meeting

(T11-2018-00024-v003)

Summary of Changes

3.1.58 Leaf Switch

- Define a Leaf Switch
 - a Switch operating as a termination point for routes (i.e., does not allow routes through the Switch)



Summary of Changes

6.2.9.3 Link State Header format

- Modify Link State Header format
 - See Table 33
- Redefine first “Reserved” Field
 - “LSR Flags”
- Define “LSR Flags”
 - **LSR Flags:**
 - Bit-0: Leaf Switch
 - The Switch shall be a route termination point
 - Other Switches in the Fabric shall not route through this Switch
 - Bits 1-7: Reserved.

Table 33 – Link State Header format

| Item | Size (bytes) |
|-------------------------------|--------------|
| LSR Type | 1 |
| LSR Flags | |
| LSR Age | 2 |
| Reserved | 4 |
| Link State Identifier | 4 |
| Advertising Domain_ID | 4 |
| Link State Incarnation Number | 4 |
| Checksum | 2 |
| LSR Length | 2 |

Summary of Changes

8.5.1 LSR Flags

- Describe LSR Flags
 - The Leaf Switch bit (i.e., bit 0) in the LSR Flags
 - Operates as a Leaf Switch (see clause 18)
 - Indicates how the Switch handles frames addressed to Domain_IDs that do not belong to the Switch
 - Leaf Switches do not allow frames to be routed through the Switch to other Switches in the Fabric
 - (i.e., no valid paths through this Switch).

Table 33 – Link State Header format

| Item | Size (bytes) |
|-------------------------------|--------------|
| LSR Type | 1 |
| LSR Flags | |
| LSR Age | 2 |
| Reserved | 4 |
| Link State Identifier | 4 |
| Advertising Domain_ID | 4 |
| Link State Incarnation Number | 4 |
| Checksum | 2 |
| LSR Length | 2 |

Summary

Clause 18 Leaf Switch

- A Leaf Switch operates at the edge of a Fabric
 - Cannot be used to access other Switches in the Fabric
 - Functions as a terminating point
 - Do not include paths through a Leaf Switch to reach any other Switch in the Fabric
- Identified by the Leaf Switch bit in the Link State Record LSR Flags
 - Leaf Switch identified during the FSPF LSR database synchronization
- Paths through a Leaf Switch are excluded from the path selection results
- Example
 - “Green” ISL is not active
 - “Red” ISLs are active
 - The path through the Leaf Switch is not valid




Figure 57 – Leaf Switch example

Move to incorporate FC-SW-7

- Leaf Switch Text
 - T11-2018-00094-v003

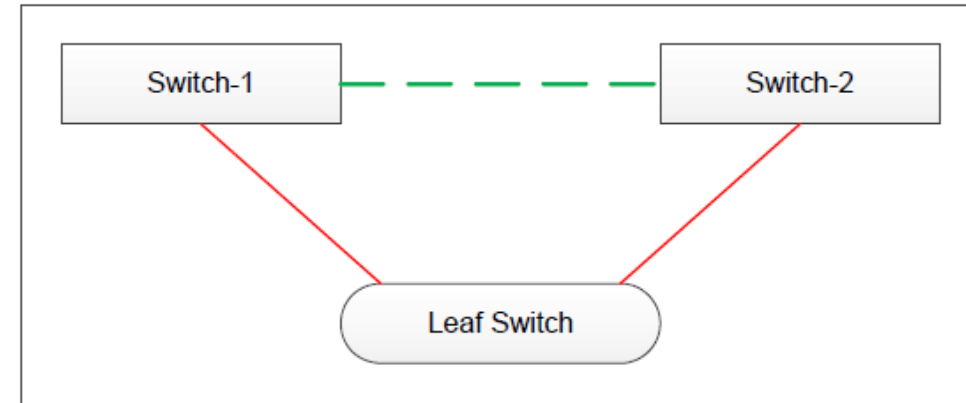


Figure 57 – Leaf Switch example

References


Notes from August 2018 Meeting

(T11-2018-00024-v002)

Where were we?

Talkin' bout Leaf Switches

- Leaf Switch Attribute
 - “Non-Pass-Through” mode
 - “Leaf Switches” are discoverable
 - “Leaf Switches” are the termination of routes
- FSPF Behaviors
 - Routes through “Leaf Switches” are not included
 - Paths to “Leaf Switches” are available
 - Paths through “Leaf Switches” are excluded
 - Unreachable domains
 - Ignore switches accessible only “through” Leaf Switches



Solution Summary

Description of “Leaf Switch”

- 6.2.9.3 Link State Header format
 - Modify Link State Header format
 - See Table 33
 - Redefine first “Reserved” Field
 - “LSR Flags”
 - Define “LSR Flags”
 - **LSR Flags:** The format of the LSR Flags field (7:0) is as follows:
 - Bit-0: Leaf Switch
 - When this bit is one, the domain shall be a terminating point (i.e. a “Leaf Switch”) and other domains in the fabric are not reachable through this domain.
 - When this bit is zero, the domain shall allow other domains in the fabric to be reachable through this domain.
 - Bits 1-7 reserved.





Table 33 – Link State Header format

| Item | Size (bytes) |
|-------------------------------|--------------|
| LSR Type | 1 |
| LSR Flags | |
| LSR Age | 2 |
| Reserved | 4 |
| Link State Identifier | 4 |
| Advertising Domain_ID | 4 |
| Link State Incarnation Number | 4 |
| Checksum | 2 |
| LSR Length | 2 |

Edge-Core

A topology example


- Core traffic restricted to “green” ISL
 - Core-to-Core route is performance optimized
- Avoid Core-1 to Leaf-B to Core-2
 - In the event of an outage, the preferred response is no route instead of slow route



Replication Core

A topology example


- Core traffic restricted to “green” ISL
 - Core-to-Core route is production traffic
 - Core-to-Edge route is replication traffic
- Avoid Core-1 to Leaf-B to Core-2
 - Replication route is geographically dispersed
 - Production traffic must avoid the replication route



Appliance (“PSA”)

A topology example

- Core traffic restricted to “green” ISL
 - Core-to-Core route is valid IO path
 - Core-to-Edge route is analytics path
- Avoid Core-1 to Leaf-B to Core-2
 - Core-to-Edge-to-Core route is invalid
 - Valid IO restricted to Core-to-Core route



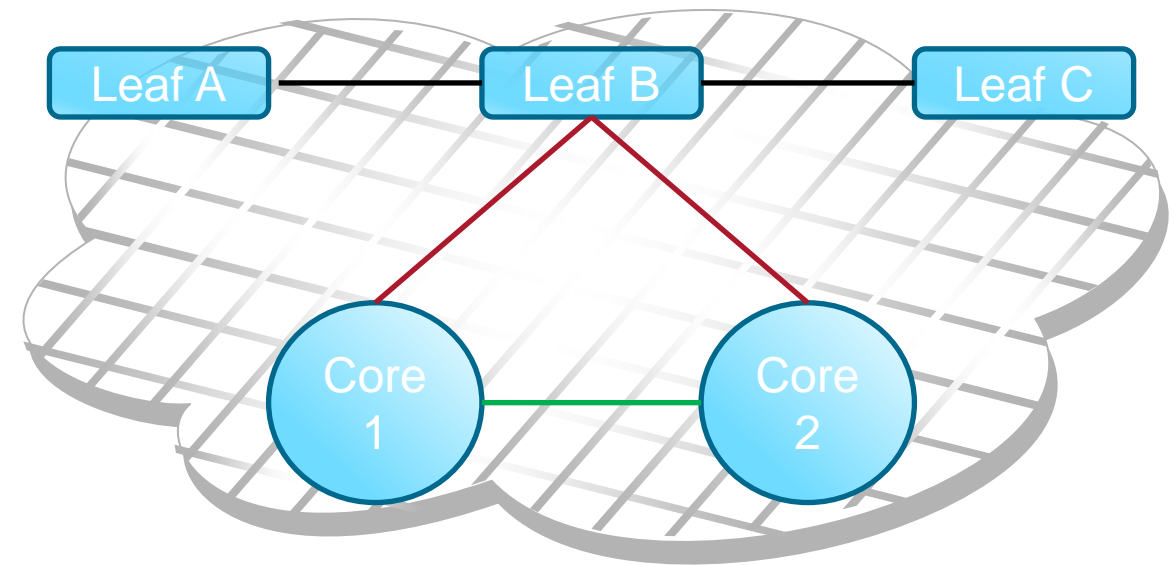
A chain, chain, chain

A topology example

- Leaf A
 - Reaches Leaf B
- Leaf B
 - Reaches Leaf A
 - Reaches Leaf C
- Leaf C
 - Reaches Leaf B

- Core 1
 - Reaches Core 2 and Leaf B
- Core 2
 - Reaches Core 1 and Leaf B

- Note
 - All switches “know” all domains exist
 - Class-F flows freely



An Essay

T11-2018-00094-v001

- Clause 18 Leaf Switch*
 - A Leaf Switch in the fabric has attachment points to the of the fabric, but no traffic or domains include paths through the Leaf Switch to reach other domains. The Leaf Switch informs the other domains in the fabric that it is a leaf with the Leaf Switch bit in the Link State Record of LSR Flags field. Link State Records are created and distributed by all domains in the fabric via the FSPF Link State database synchronization and flooding protocols. The records describe the available links connecting the switch to the fabric.
 - Any switch setting the Leaf Switch bit is discoverable and the paths to it are indicated as available; however, any paths through a Leaf Switch are excluded from the path selection results. Similarly, any domain that is discoverable only by a path through a Leaf Switch is ignored and considered unreachable.
 - In Figure 1, while the “green” ISL is not active, the “red” ISL paths between Switch-1 and Switch-2 that passes through the Leaf Switch is not considered a valid path and is excluded. Therefore, Switch-1 and Switch-2 each consider the other as unreachable; but both switches treat the Leaf Switch as reachable. The Leaf Switch considers both Switch-1 and Switch-2 as reachable. Once the “green” ISL is active, Switch-1 and Switch-2 treat each other as reachable, distribute their device information, and send RSCNs to the attached devices.




Figure 1 – Leaf Switch

* Summer reading test at the October meeting

References

Notes from April 2018 Meeting

(T11-2018-00024-v001)

“Leaf Switch” Summary

Description of Non-Pass-Through mode

- 6.2.9.3 Link State Header format
 - Modify Link State Header format
 - See Table 33
 - Redefine first “Reserved” Field
 - “LSR Flags”
 - Define “LSR Flags”
 - **LSR Flags:** The format of the LSR Flags field (7:0) is as follows:
 - Bit-0: Non-Pass-Through Mode – When this bit is one, the Switch shall work in Non-Pass-Through mode and other domains in the fabric are not reachable through this Switch. When this bit is zero, the Switch shall work in Traditional mode and other domains in the fabric are reachable through this Switch.
 - Bits 1-7 reserved.
- Reference
 - T11-2018-00094-v000
 - Leafing through the Fabric (Non-pass through switch) Text

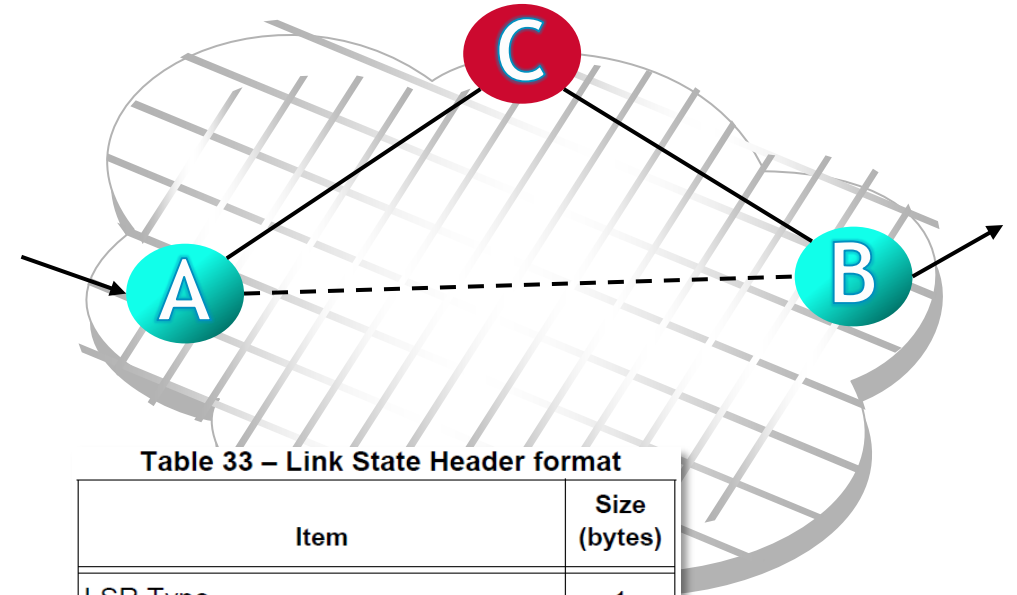


Table 33 – Link State Header format

| Item | Size (bytes) |
|-------------------------------|--------------|
| LSR Type | 1 |
| LSR Flags | |
| LSR Age | 2 |
| Reserved | 4 |
| Link State Identifier | 4 |
| Advertising Domain_ID | 4 |
| Link State Incarnation Number | 4 |
| Checksum | 2 |
| LSR Length | 2 |

References


Notes from February 2018 Meeting

(T11-2018-00024-v000)

“Leaf Switch” Summary

Description of Non-Pass-Through mode

- Non-Pass-Through mode
 - “Leaf Switches” are discoverable
 - “Leaf Switches” are the termination of routes
- Traditional mode
 - Switches do not include routes through “Leaf Switches” in FSPF
 - Paths to “Leaf Switches” are available
 - Paths through “Leaf Switches” are excluded
 - Domains discovered only through “Leaf Switches” are ignored (i.e. unreachable)



Proposal

Designate a “Leaf Switch”

- 6.2.9.3 Link State Header format
 - Modify Link State Header format
 - See Table 33
 - Redefine first “Reserved” Field
 - “LSR Flags”
 - Define “LSR Flags”
 - **LSR Flags:** The format of the LSR Flags field (7:0) is as follows:
 - Bit-0: Non-Pass-Through Mode – When this bit is one, the Switch shall work in Non-Pass-Through mode and other domains in the fabric are not reachable through this Switch. When this bit is zero, the Switch shall work in Traditional mode and other domains in the fabric are reachable through this Switch.
 - Bits 1-7 reserved.


Table 33 – Link State Header format

| Item | Size (bytes) |
|-------------------------------|--------------|
| LSR Type | 1 |
| LSR Flags | 1 |
| LSR Age | 2 |
| Reserved | 4 |
| Link State Identifier | 4 |
| Advertising Domain_ID | 4 |
| Link State Incarnation Number | 4 |
| Checksum | 2 |
| LSR Length | 2 |

Leaf Switch

April Meeting

- Provide text updates
 - “Instructions to the Editor”



References


Notes from December 2017 Meeting

(T11-2017-00169-v001)

Leaf Switches

Characteristics

- Uses
 - Low end product
 - Limited resources
 - Provides access point to fabric
 - Minimal complexity
 - An offload engine with a fabric presence
 - Fabric based applications
 - i.e. Encryption/Compression Services, Replication Assistant, Analytics, etc
- Fabric visibility, however ...
 - Route to it, but not through it



Leaf Switches

Designation

- Link State Record descriptor
 - Leaf switch designation in the Link State Record (LSR)
 - Every switch authors a LSR record that describes the available links connected to that switch. The LSRs are distributed to all switches in the fabric via the FSPF LSDB synchronization and flooding protocols.


Table 32 – Link State Record - Link Descriptor format

| Item | Size (bytes) |
|---|--------------|
| Link State Record Header (LSR Type 01h) | 24 |
| Reserved | 2 |
| Number of Links | 2 |
| Link Descriptor #1 | 16 |
| ... | 16 |
| ... | 16 |
| Link Descriptor #n | 16 |

Leaf Switches

Behavior


- FSPF calculations
 - Leaf switches are discoverable
 - Paths to Leaf switches are available
 - Paths through Leaf switches are excluded
- Implication
 - Domains discoverable only through Leaf switches are ignored (i.e. unreachable)



Leaf Switch

Summary

- Used for low-end or PSA applications
- Designation in LSR
- Exclude in FSPF route-thru decisions




References

Notes from June 2017 Meeting

Leaf Switches

Characteristics

- Low end product
 - Limited resources
 - Provides access point to fabric
 - Minimal complexity
- Fabric visibility, but ...
 - Route to it, but not through it



Leaf Switches

Designation

- Link State Record descriptor
 - or
- Link State Record header

- Food for thought
 - August proposal

Table 33 – Link State Header format

| Item | Size (bytes) |
|-------------------------------|--------------|
| LSR Type | 1 |
| Reserved | 1 |
| LSR Age | 2 |
| Reserved | 4 |
| Link State Identifier | 4 |
| Advertising Domain_ID | 4 |
| Link State Incarnation Number | 4 |
| Checksum | 2 |
| LSR Length | 2 |

Table 32 – Link State Record - Link Descriptor format

| Item | Size (bytes) |
|---|--------------|
| Link State Record Header (LSR Type 01h) | 24 |
| Reserved | 2 |
| Number of Links | 2 |
| Link Descriptor #1 | 16 |
| ... | 16 |
| ... | 16 |
| Link Descriptor #n | 16 |

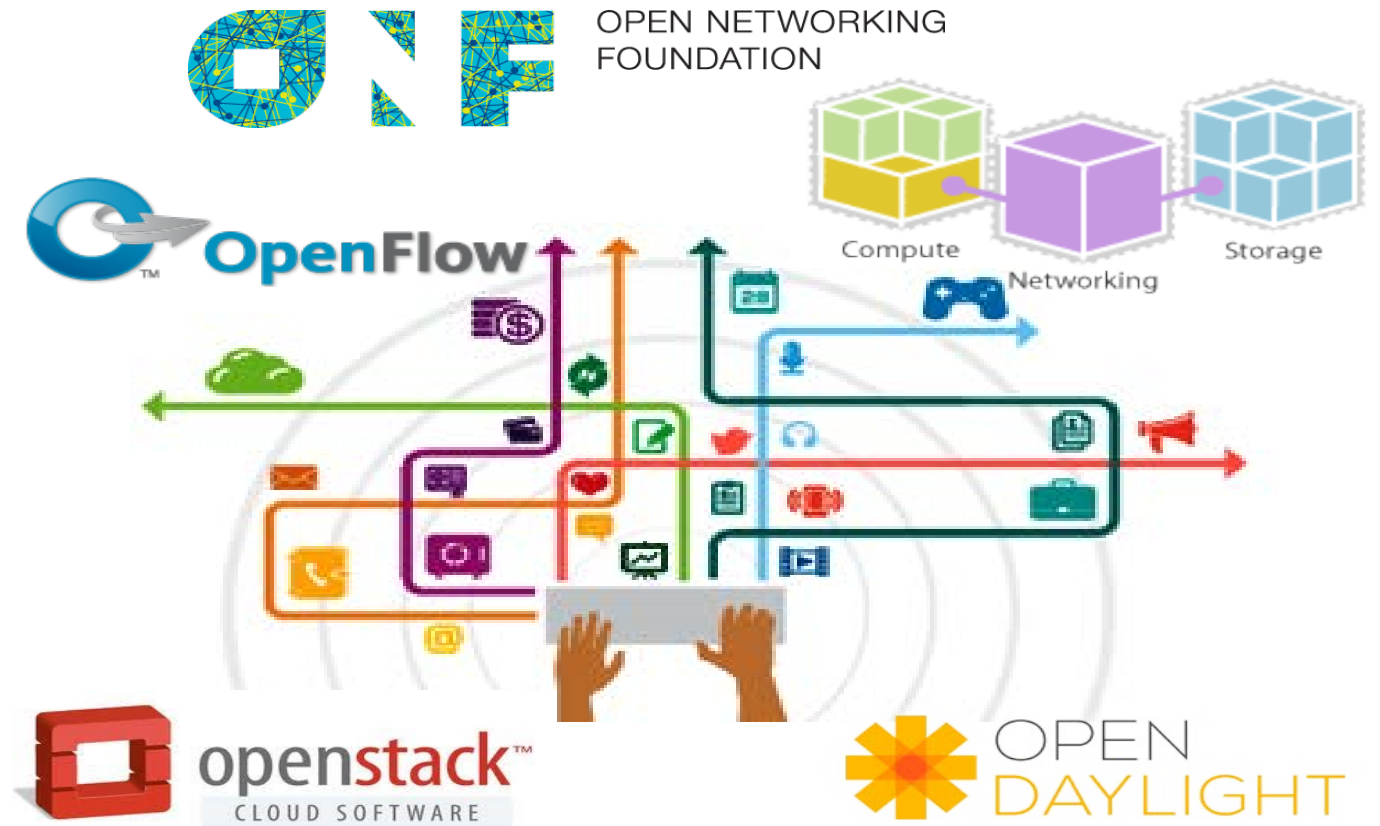
References

Notes from December 2016 Meeting (16-496v0)

Cloud Ready Fibre Channel

A Quest from the T11 Chair – Hi Steve!


- What services should FC provide?
- How do those services appear?



Expanding on Previous Conversations

Leaf Switches, RPI, Virtual Domains, Application/Discovery Services, etc


- The switch at the end of the road
 - Just another PSA or something more?
- The route (about, through, around)
 - Request Path Information/Data
- I want my vDomain
 - Nx_Port specific
- Fabric Services expansion
 - Application Services
 - Discovery Services



FC-SW-7

PSAs and other aliens

- An offload engine with a fabric presence
 - Fabric based applications
 - Encryption/Compression Services
 - Replication Assistant
 - Analytics
- Alien origins
 - Domain or Node?
 - To merge or not to merge?
 - What's my view?



End of Frame

Thank you