



# FCoE Discovery

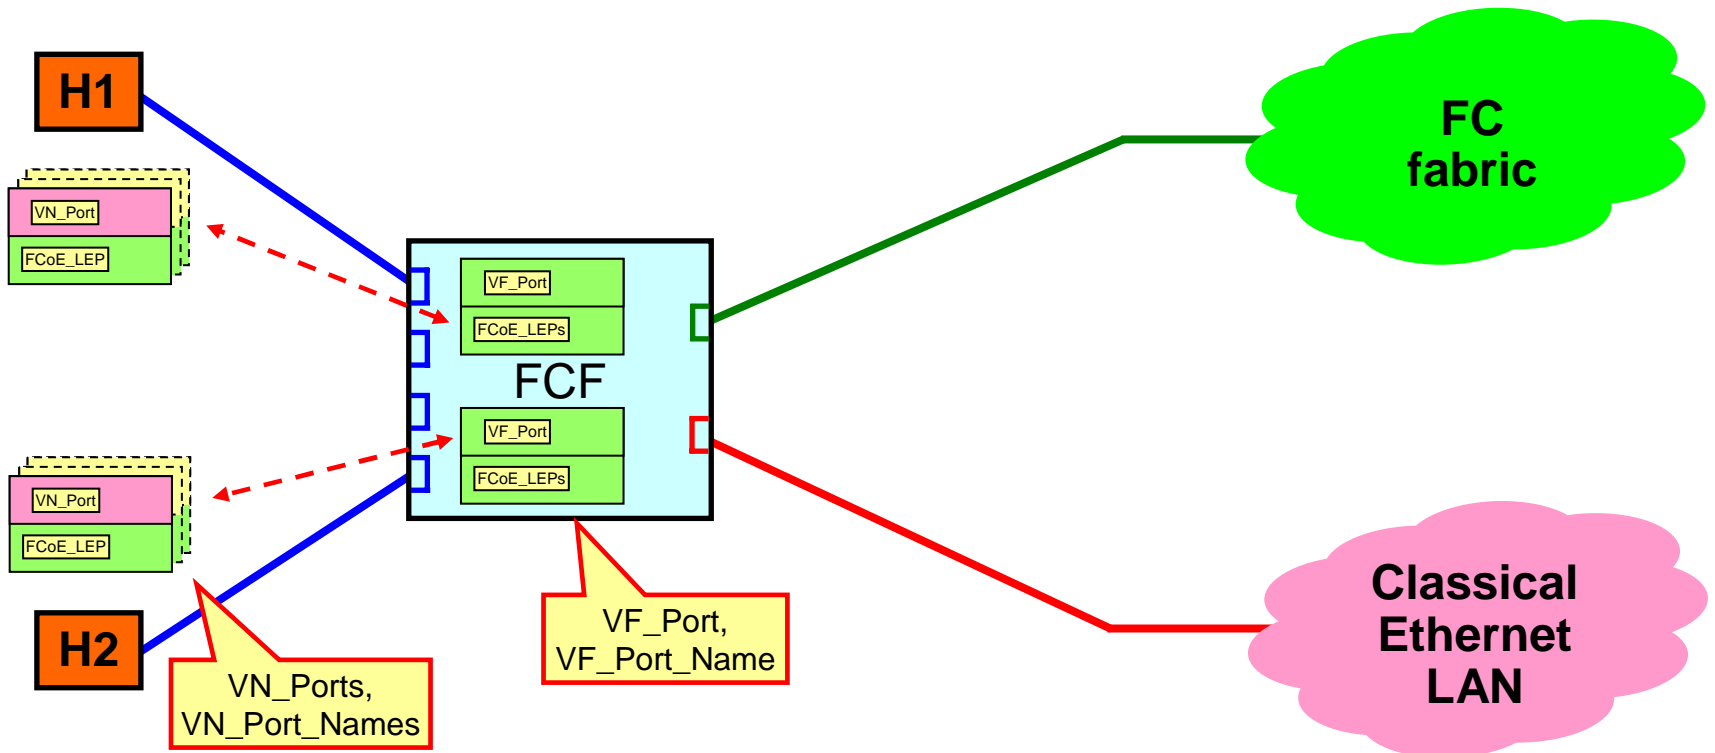
T11/07-572v0, October 2007

Claudio DeSanti

# Agenda

- **Topologies and Discovery**
- **A Case Study**
- **FCoE Discovery Protocol**

# Simple Topology



# Simple 'Discovery'

- **When an ENode is directly connected to an FCF, the FCF discovery is not needed**

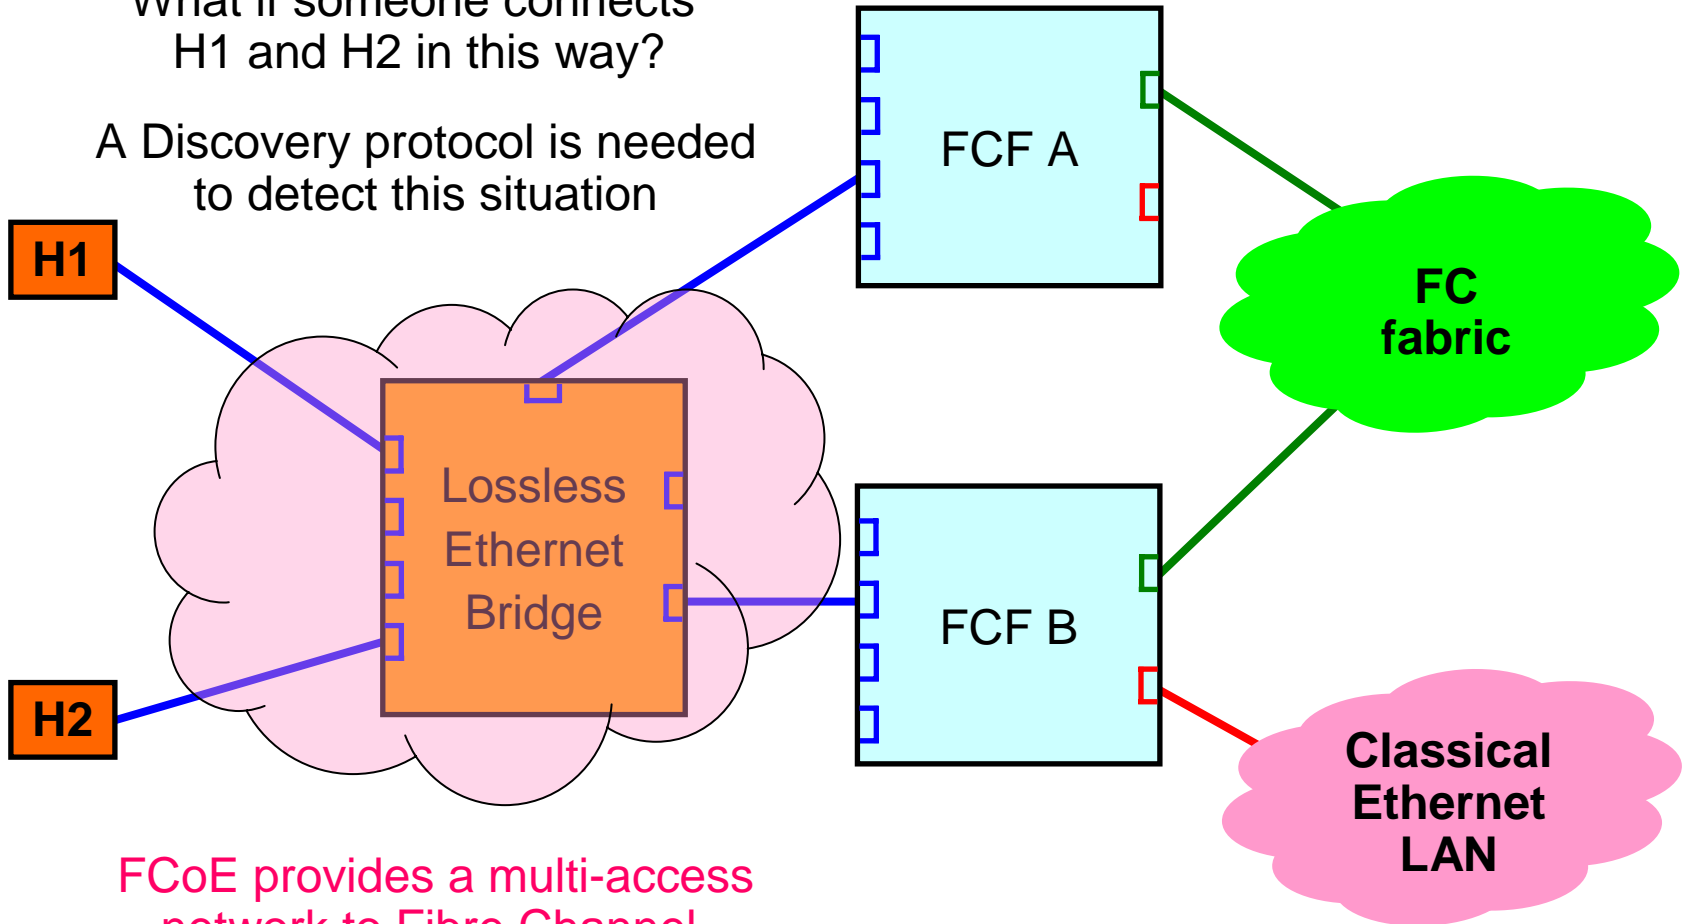
In fact it has not been discussed in the past...

- **However there is no way to enforce the simple topology only...**

# Generic Topology (1)

What if someone connects  
H1 and H2 in this way?

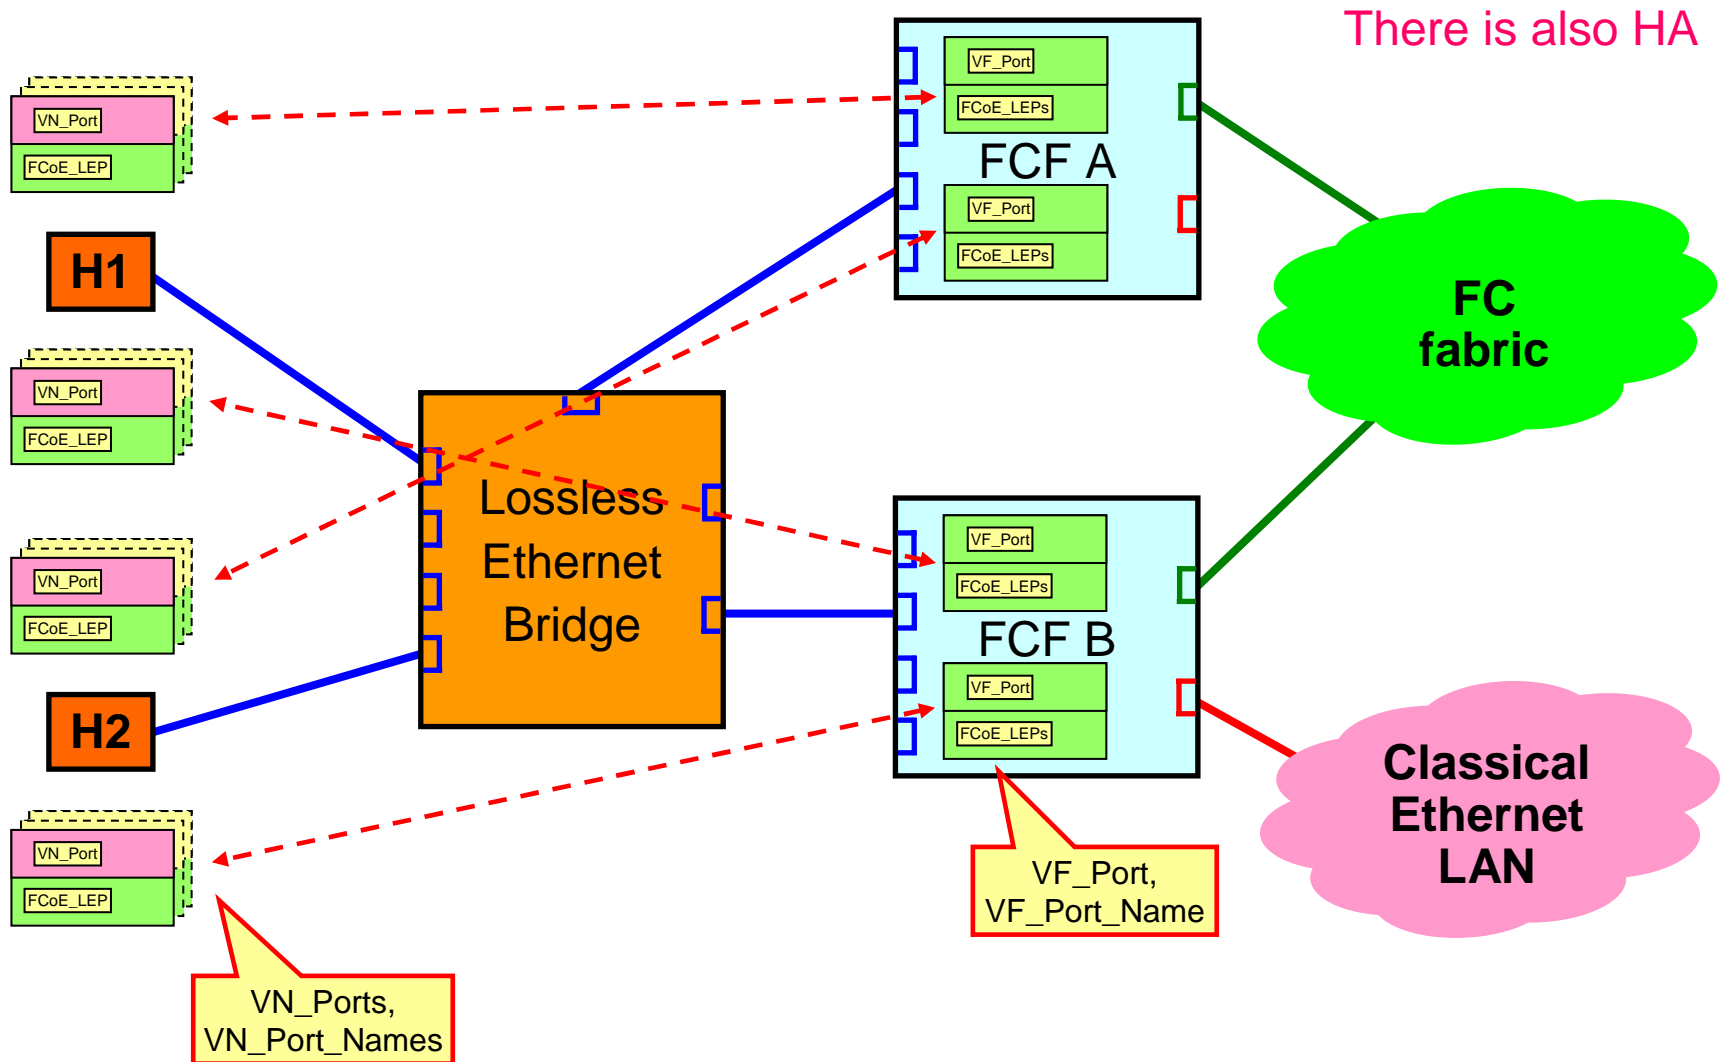
A Discovery protocol is needed  
to detect this situation



FCoE provides a multi-access  
network to Fibre Channel,  
traditionally point-to-point only

This configuration could work...

# Generic Topology (2)



# FCoE Discovery

- **Before FLOGI, each Ethernet port of an ENode sends a solicitation ELS to the multicast address All-FCoE-Forwarders**  
Carrying its Universal burned-in MAC address
- **Each FCF replies with a unicast advertisement ELS**  
Carrying its FCF-MAC address (and other parameters)
- **ENodes builds a list of the reachable FCFs**
- **Each ENode performs a unicast FLOGI to a subset of the available FCFs**  
To keep the FLOGI process unchanged  
FDISC to the same FCF-MAC address provides additional N\_Port\_IDs
- **Periodically FCFs send a multicast advertisement ELS to the multicast address All-FCoE-ENodes**  
To allow discovery of a newly connected FCF

# Considerations

- **Advertisement messages may be used to carry additional information to ENodes**

E.g., the value of the FC-OUI, the Switch\_Name, the Fabric\_Name

- **ENodes may use the value of the FC-OUI parameter to construct the FCoE Mapped MAC address**

The FC-OUI is a network configuration parameter

- **The provided information may be used by ENodes to detect anomalous configurations or to select the FCFs where to perform FLOGI among the available ones**

E.g., receiving two different Fabric\_Names or two different FC-OUI values indicates something “strange”

- **The FCoE Discovery protocol is functionally associated with FCoE Controllers**

May be used to verify the connectivity between FCoE Controllers

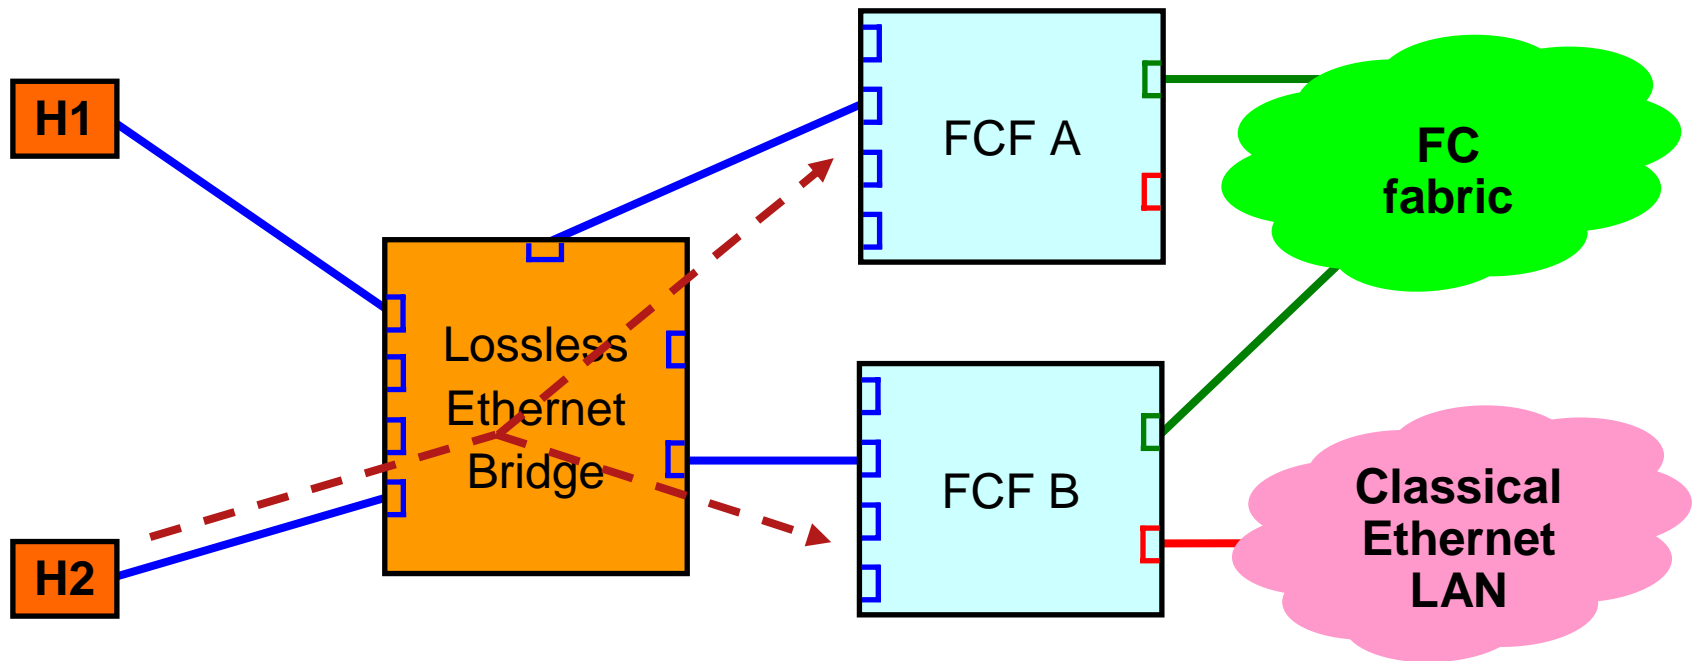
# FCF Discovery

- **The same discovery protocol may be used also by FCFs to discover each other**
- **When connected to a network, an FCF sends a solicitation ELS to the multicast address All-FCoE-Forwarders**
  - Carrying its FCF-MAC address and other parameters
- **Each FCF replies with a unicast advertisement ELS**
  - Carrying its FCF-MAC address and other parameters
- **The FCF creates an instance of a VE\_Port per each of the discovered FCFs**
  - Running the FC initialization protocols
- **Periodically FCFs send a multicast advertisement ELS to the multicast address All-FCoE-Forwarders**
  - To allow discovery of a newly connected FCF

# Agenda


- **Topologies and Discovery**
- **A Case Study**
- **FCoE Discovery Protocol**

# Solicitation from H2



All-FCFs
MAC(H2)
Solicitation
[MAC(H2)]

# Advertisements from A and B




H2's FCF list:

FCF-MAC(A)  
FCF-MAC(B)

MAC(H2)
FCF-MAC(A)
Advertisement
[FCF-MAC(A), FC-OUI value]

MAC(H2)
FCF-MAC(B)
Advertisement
[FCF-MAC(B), FC-OUI value]

# FLOGI/FDISC Requests



H2's FCF list:

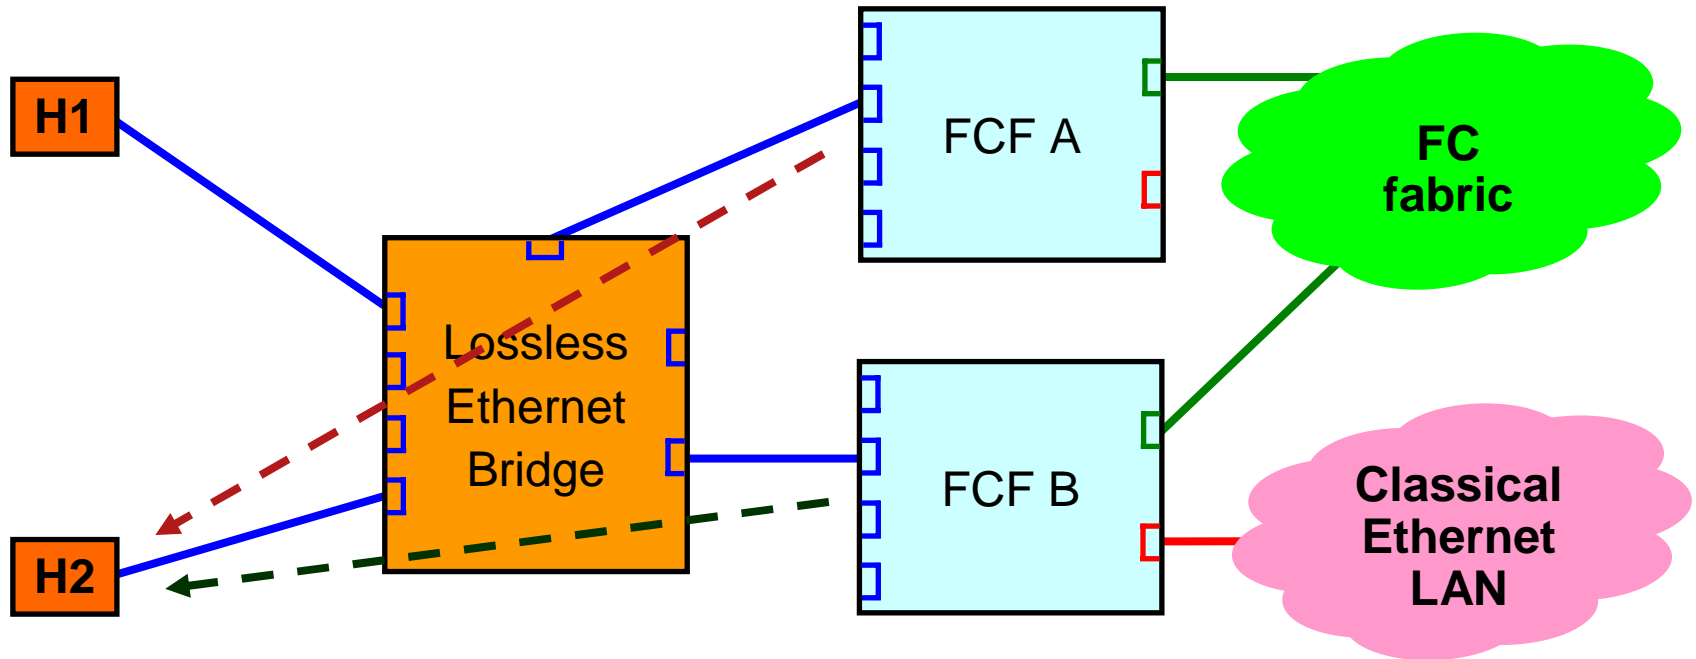
FCF-MAC(A)  
FCF-MAC(B)

FCF-MAC(A)
MAC(H2)
Encap. FC frame
D_ID = FF.FF.FE
S_ID = 00.00.00

FCF-MAC(B)
MAC(H2)
Encap. FC frame
D_ID = FF.FF.FE
S_ID = 00.00.00

The SA of all FLOGI/FDISC Requests is the Universal MAC address of the station

# FLOGI/FDISC Accepts



H2's FCF list:

FCF-MAC(A)  
FCF-MAC(B)

MAC(H2)
FCF-MAC(A)
Encap. FC frame
D_ID = XX.XX.XX
S_ID = FF.FF.FE

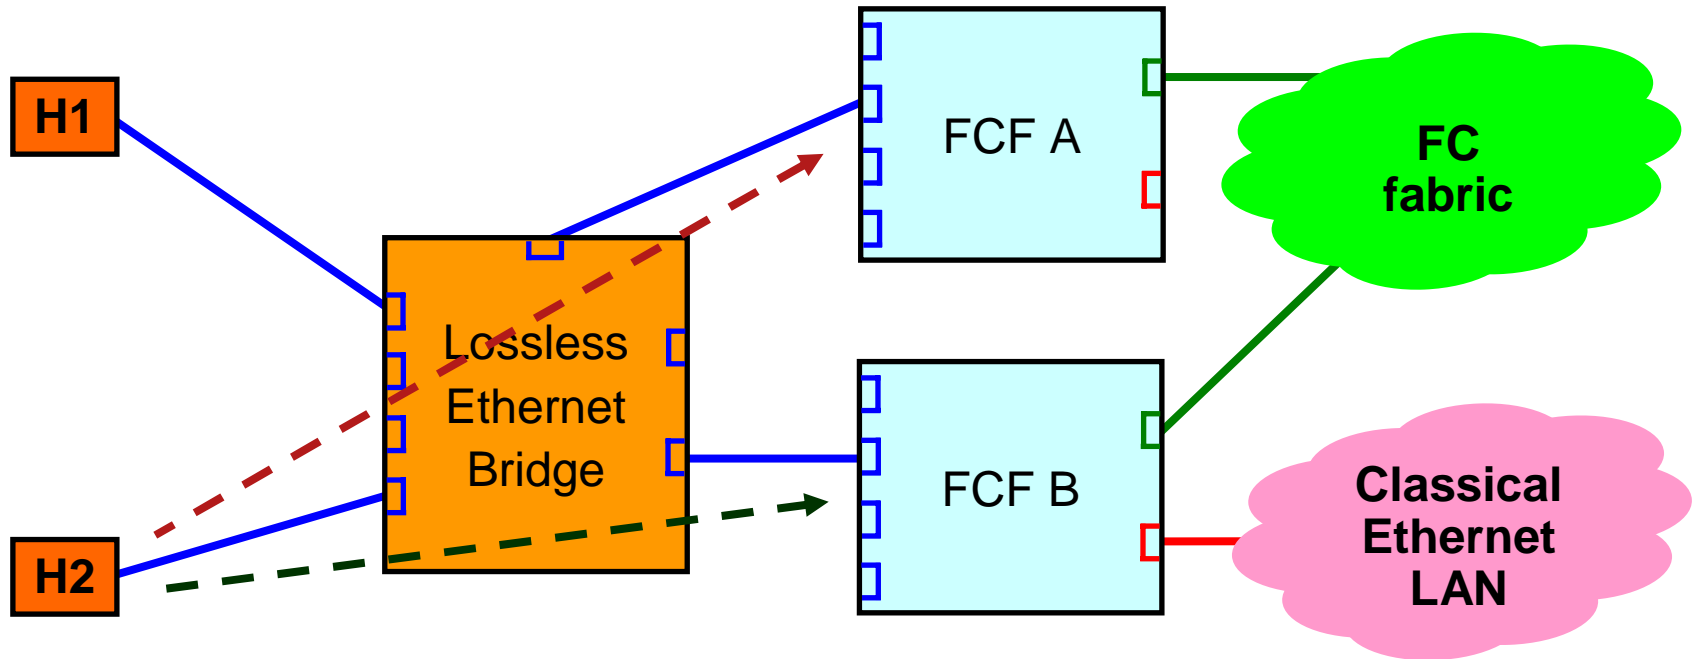
MAC(H2)
FCF-MAC(B)
Encap. FC frame
D_ID = YY.YY.YY
S_ID = FF.FF.FE

H2's FCoE\_LEPs:

FC-OUI || XX.XX.XX, FCF-MAC(A)  
FC-OUI || YY.YY.YY, FCF-MAC(B)

FC-OUI value  
provided by the fabric

# Subsequent FCoE Frames (1)



H2's FCF list:

FCF-MAC(A)  
FCF-MAC(B)

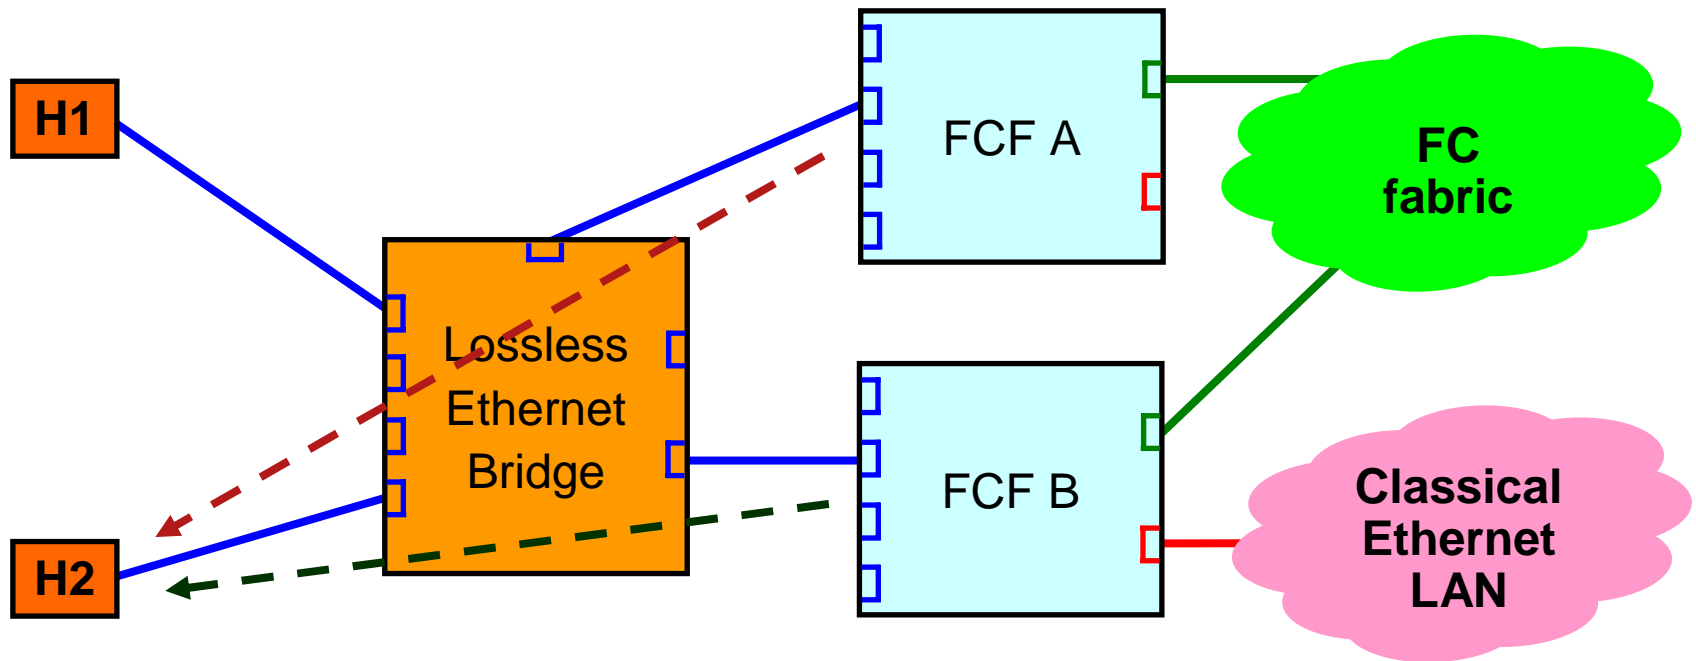
<b>FCF-MAC(A)</b>
<b>FC-OUI  XX.XX.XX</b>
<b>Encap. FC frame</b>
<b>D_ID = ZZ.ZZ.ZZ</b>
<b>S_ID = XX.XX.XX</b>

<b>FCF-MAC(B)</b>
<b>FC-OUI  YY.YY.YY</b>
<b>Encap. FC frame</b>
<b>D_ID = VV.VV.VV</b>
<b>S_ID = YY.YY.YY</b>

H2's FCoE\_LEPs:

FC-OUI || XX.XX.XX, FCF-MAC(A)  
FC-OUI || YY.YY.YY, FCF-MAC(B)

# Subsequent FCoE Frames (2)



H2's FCF list:

FCF-MAC(A)  
FCF-MAC(B)

FC-OUI  XX.XX.XX
FCF-MAC(A)
Encap. FC frame
D_ID = XX.XX.XX
S_ID = ZZ.ZZ.ZZ

FC-OUI  YY.YY.YY
FCF-MAC(B)
Encap. FC frame
D_ID = YY.YY.YY
S_ID = VV.VV.VV

H2's FCoE\_LEPs:

FC-OUI || XX.XX.XX, FCF-MAC(A)  
FC-OUI || YY.YY.YY, FCF-MAC(B)

# Management Considerations

- **The Universal burn-in MAC address of a station may be used to relate together all VN\_Ports residing behind that station**

**A function similar to the Permanent Port\_Name for physical N\_Ports**

**That address is used as SA in all FLOGI/FDISC Requests**

- **A new NAA Name type, for MAC addresses, may allow to register the Universal burn-in MAC address of a station as the Permanent Port\_Name of each VN\_Port**

**The F\_Port\_Name can be retrieved directly from the Name Server**

**Alternatively, the MAC address could be an additional Name Server attribute for a VN\_Port**

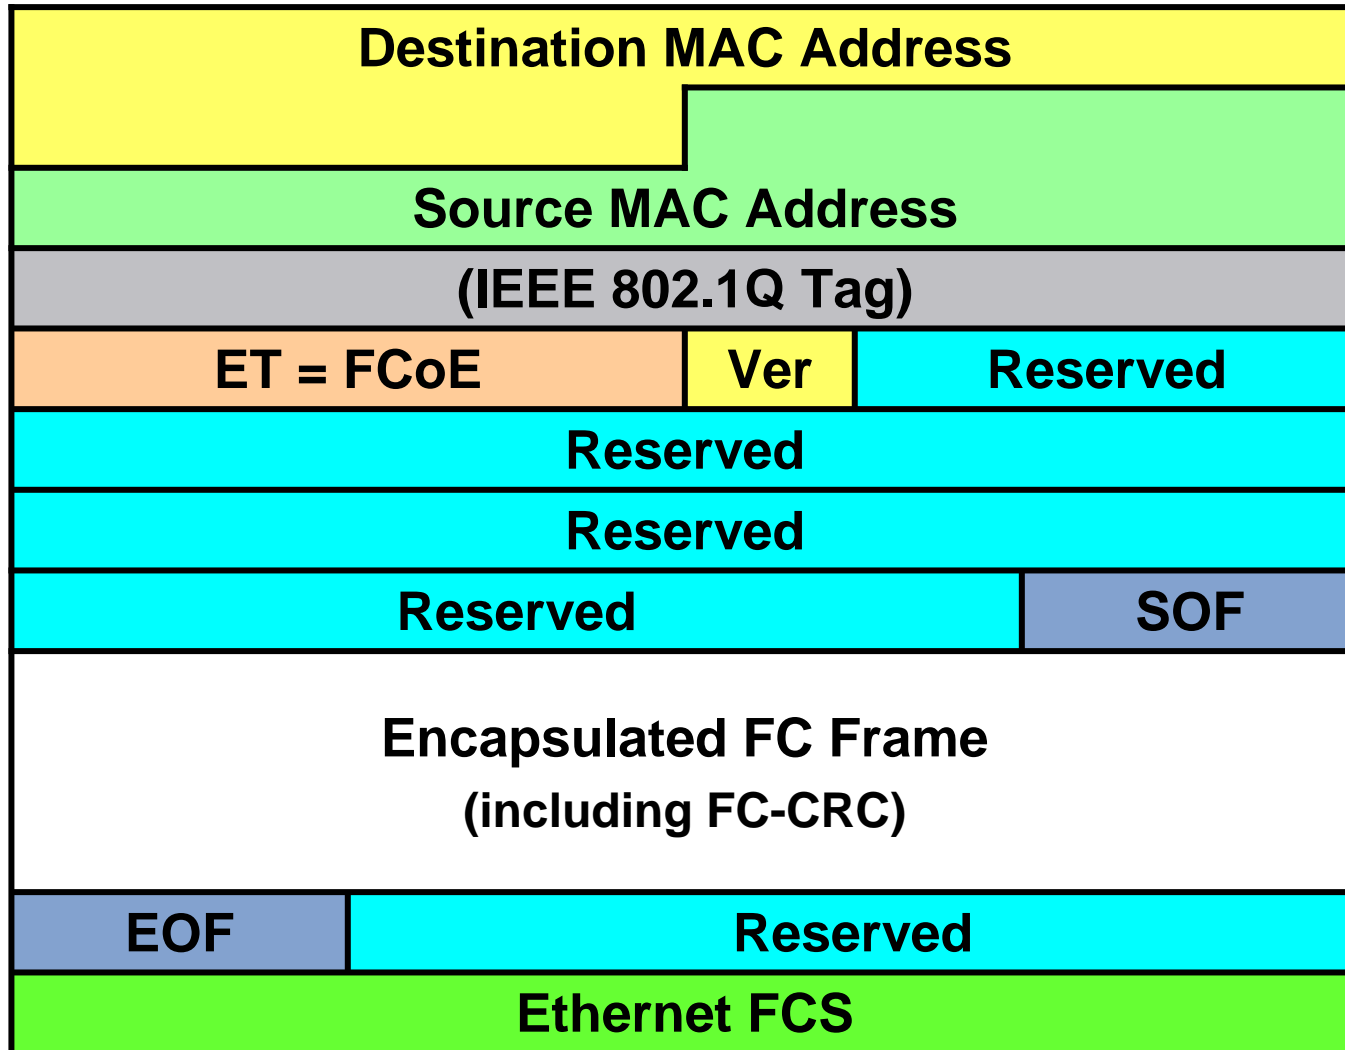
# Agenda

- **Topologies and Discovery**
- **A Case Study**
- **FCoE Discovery Protocol**

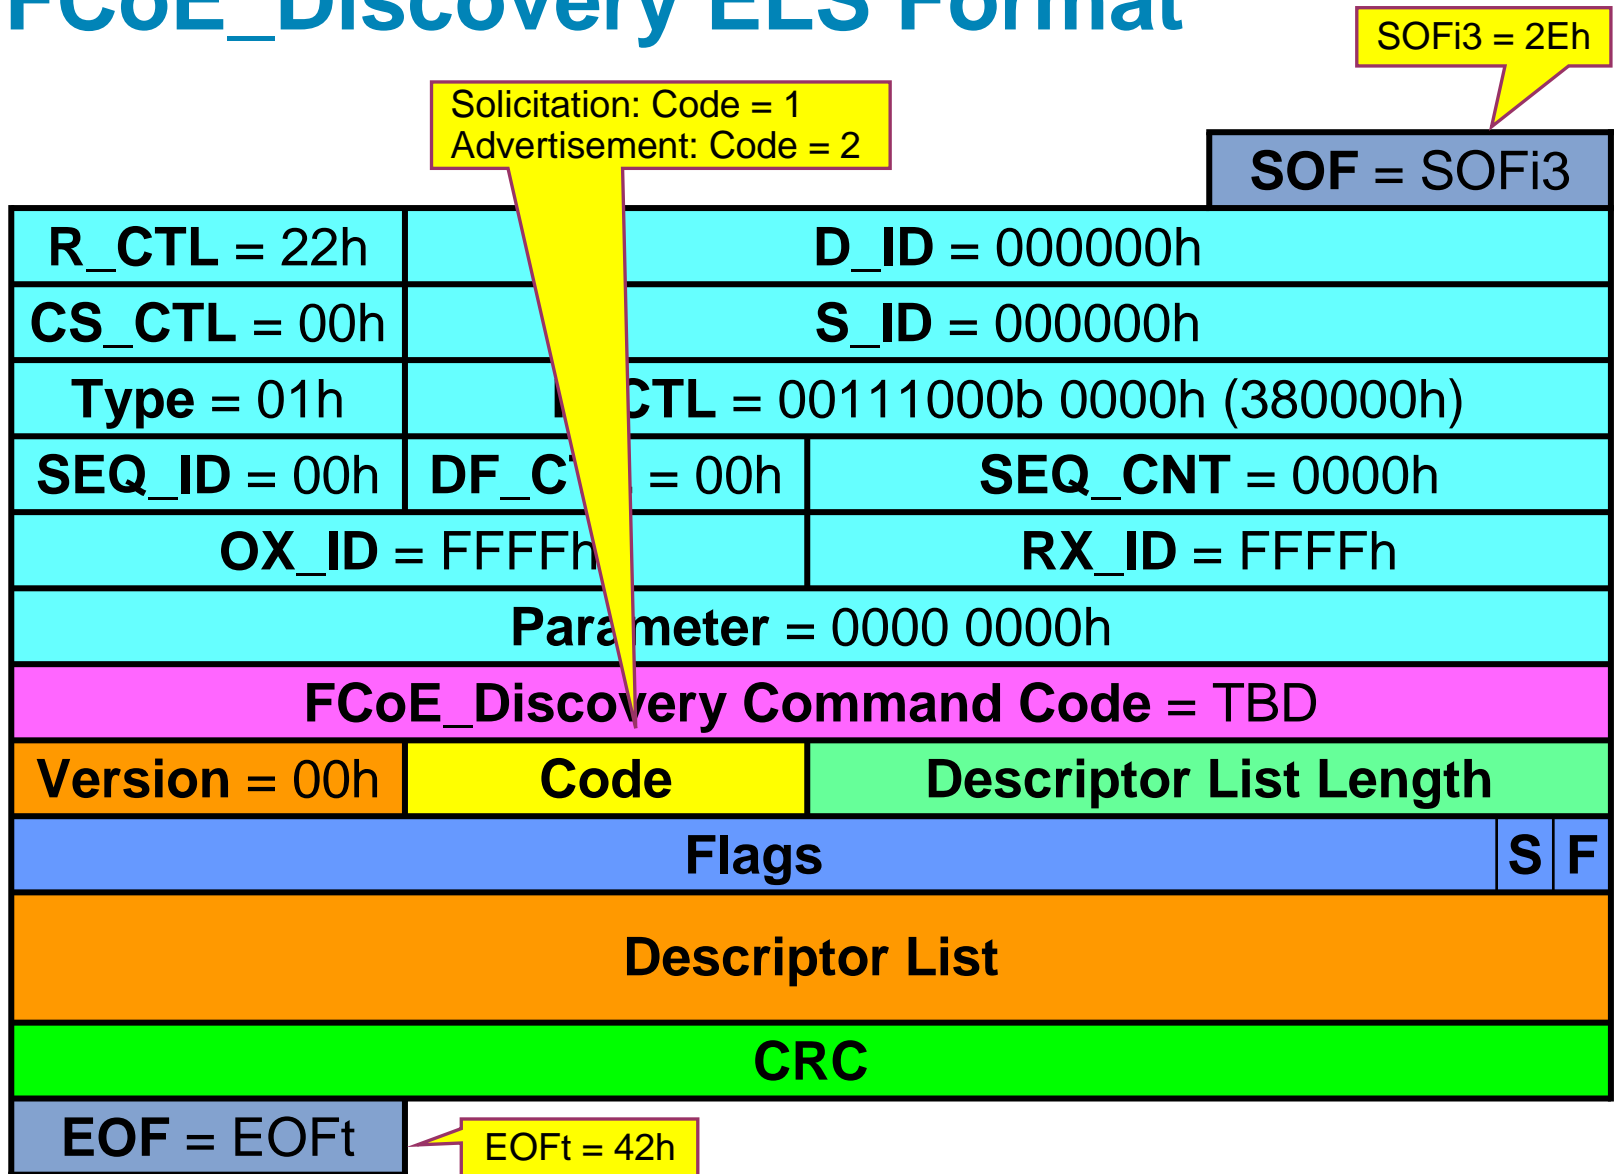
# FCoE Discovery Protocol

- **It is an Ethernet discovery protocol**
  - Using a single Ethernet frame per message
- **Each message of the FCoE Discovery protocol is defined as:**
  - an ELS
  - mapped to a single frame Class 3 Sequence
  - in a single Sequence unidirectional Exchange
  - encapsulated in a single FCoE Frame
- **D\_ID and S\_ID should be set to zero**
  - The messages are exchanged before any FC-ID is assigned
- **The protocol should be extensible**
- **A native FC\_Port should never see these messages**





# FCoE Frame Format



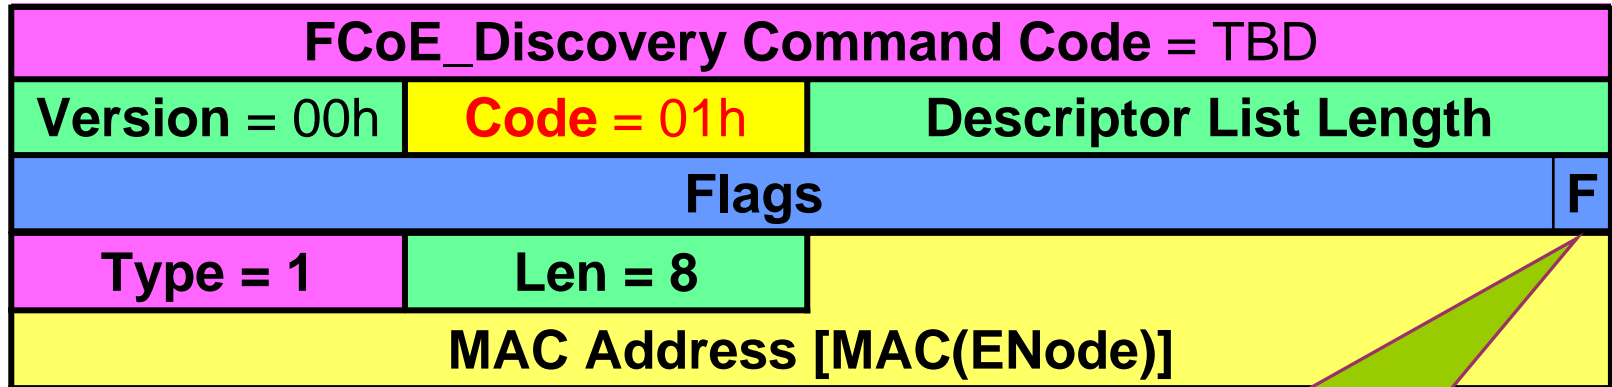
# FCoE\_Discovery ELS Format



# FCoE\_Discovery Descriptors

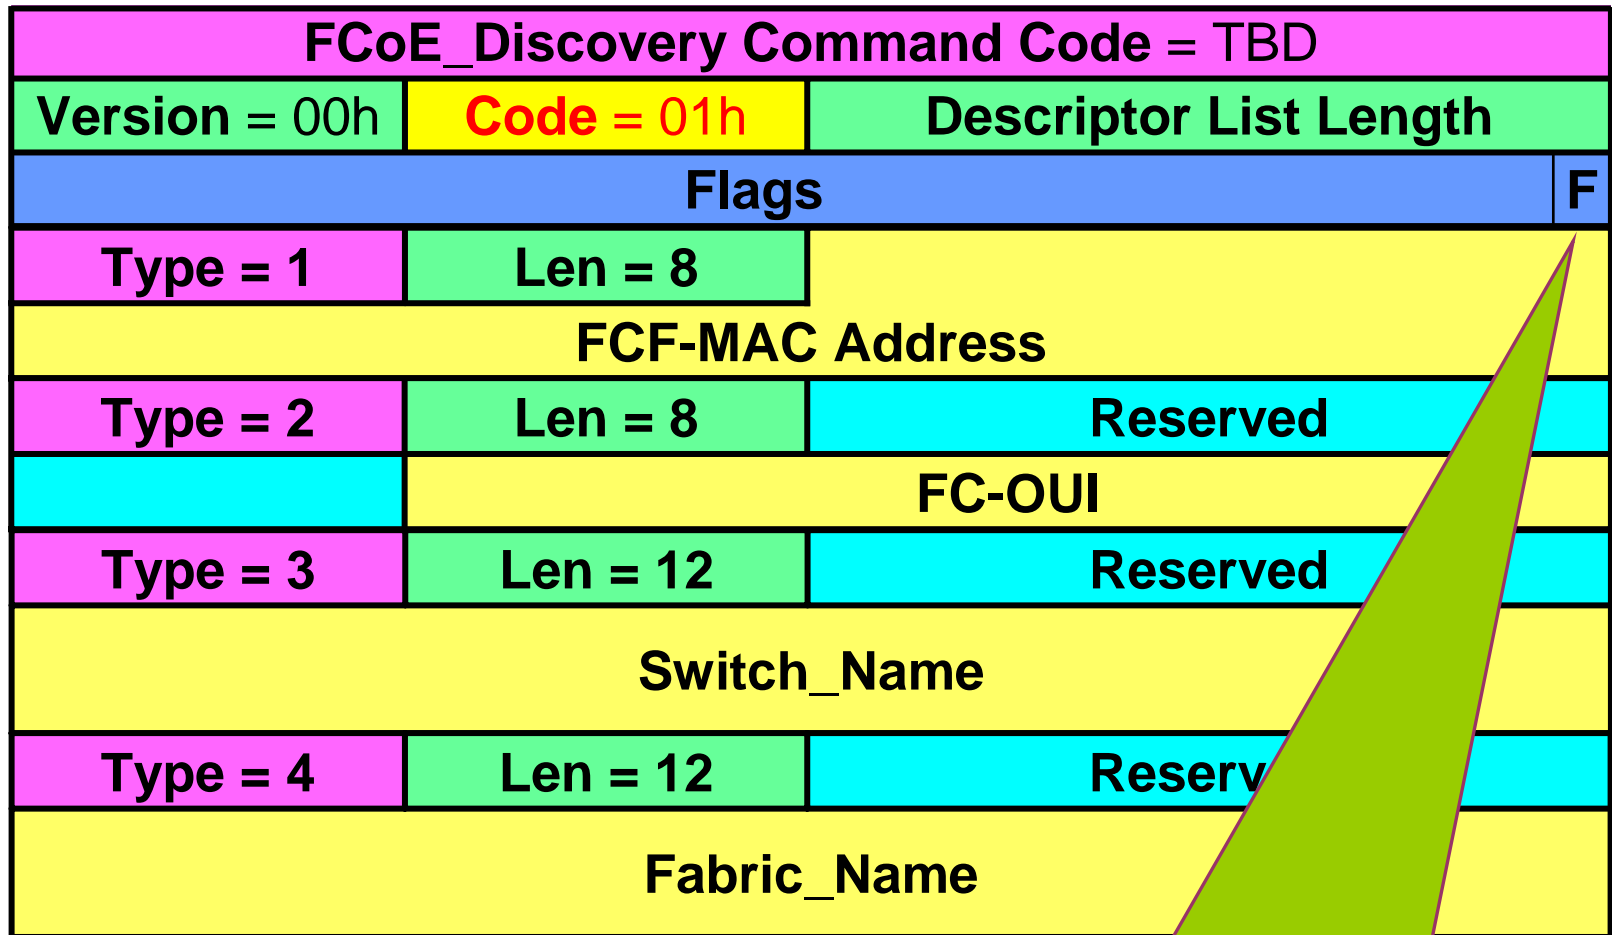


# Solicitation from an ENode



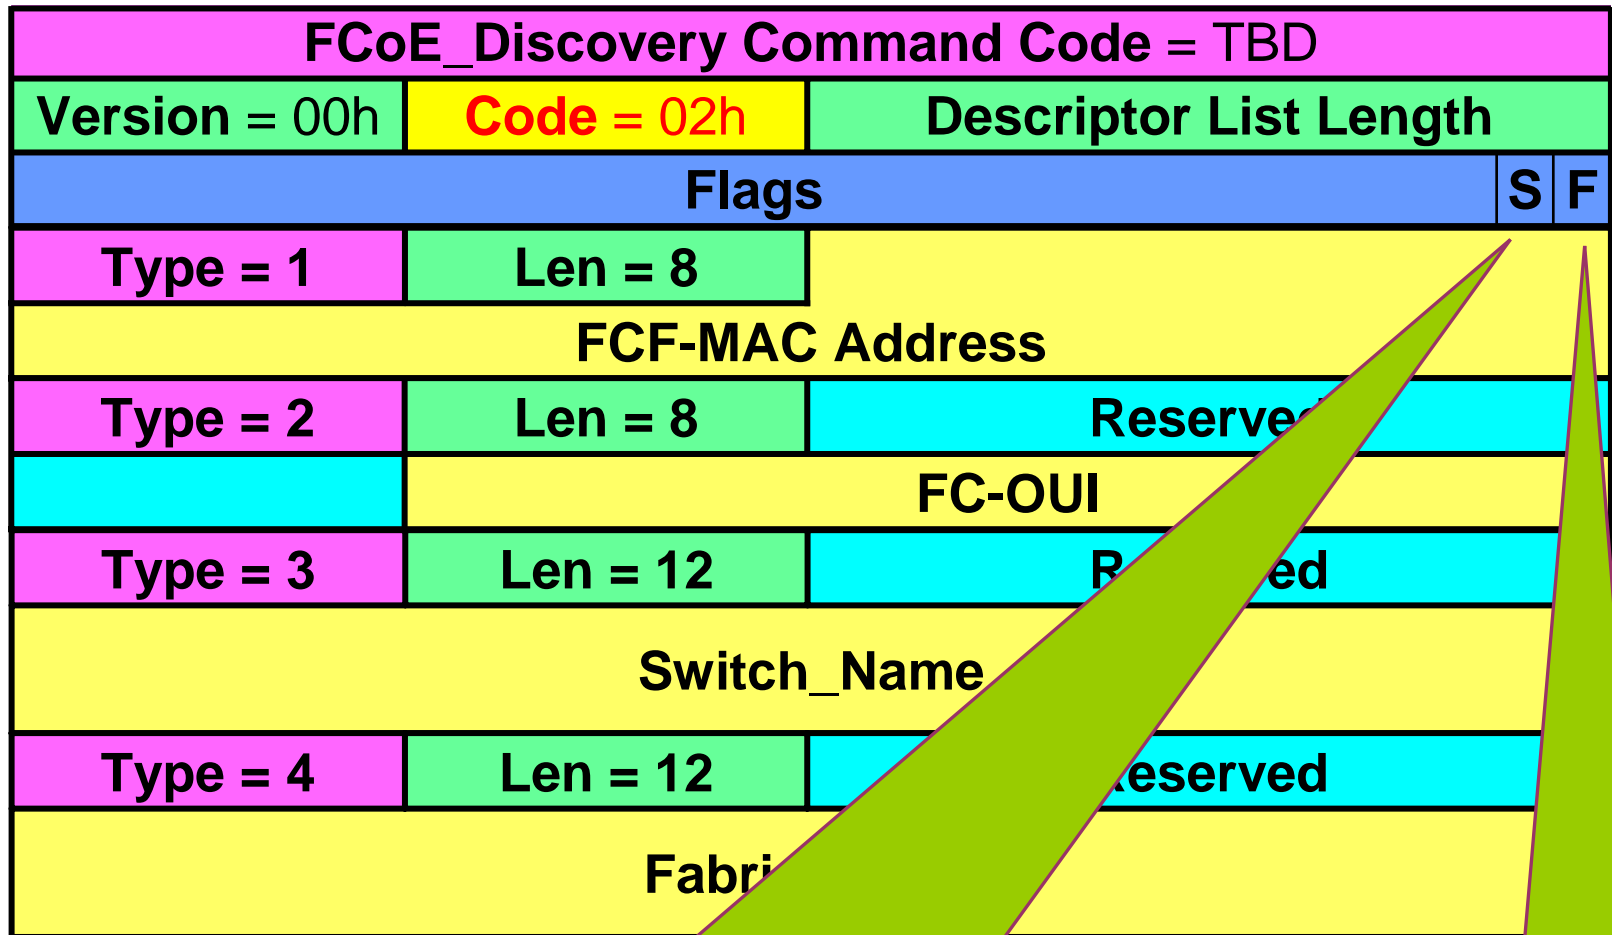
F = 0b to indicate that the Solicitation is sent by an ENode

# Solicitation from an FCF



F = 1b to indicate that the Solicitation is sent by an FCF

# Advertisements



S = 1b if the Advertisement is sent in response to a Solicitation  
 S = 0b if the Advertisement is not sent in response to a Solicitation

F = 1b (Advertisements are sent only by FCFs)

# Loss of Connectivity Detection

- **If no FCoE traffic is received from a remote FCoE Controller for more than a defined time limit (TBD), an FCoE Controller may send a unicast Solicitation message to the remote FCoE Controller**
- **If the remote FCoE Controller is alive it replies with an unicast Advertisement message with the Solicited flag set to one**
  - To explicitly indicate it is sent in response of a Solicitation**
  - Then connectivity is confirmed**
- **After a certain number (TBD) of Solicitations without a reply, the FCoE Controller may infer the remote FCoE Controller is no more reachable**
- **This is a more efficient solution than using LKA**

# Summary

- **FCoE needs a Discovery protocol**

  - To keep the FLOGI processing unchanged

  - To deal with the multi-access nature of Ethernet

    - E.g., to determine how many FCFs are connected to a certain Ethernet segment

    - Information needed to both ENodes and FCFs

- **The FCoE Discovery protocol allows early detection of potential network misconfiguration**

  - Allowing both ENodes and FCFs to react properly

- **The FCoE Discovery protocol provides configuration information for ENodes**

  - E.g., the FC-OUI value

Thank You