802.1D and 802.1W BPDU Format

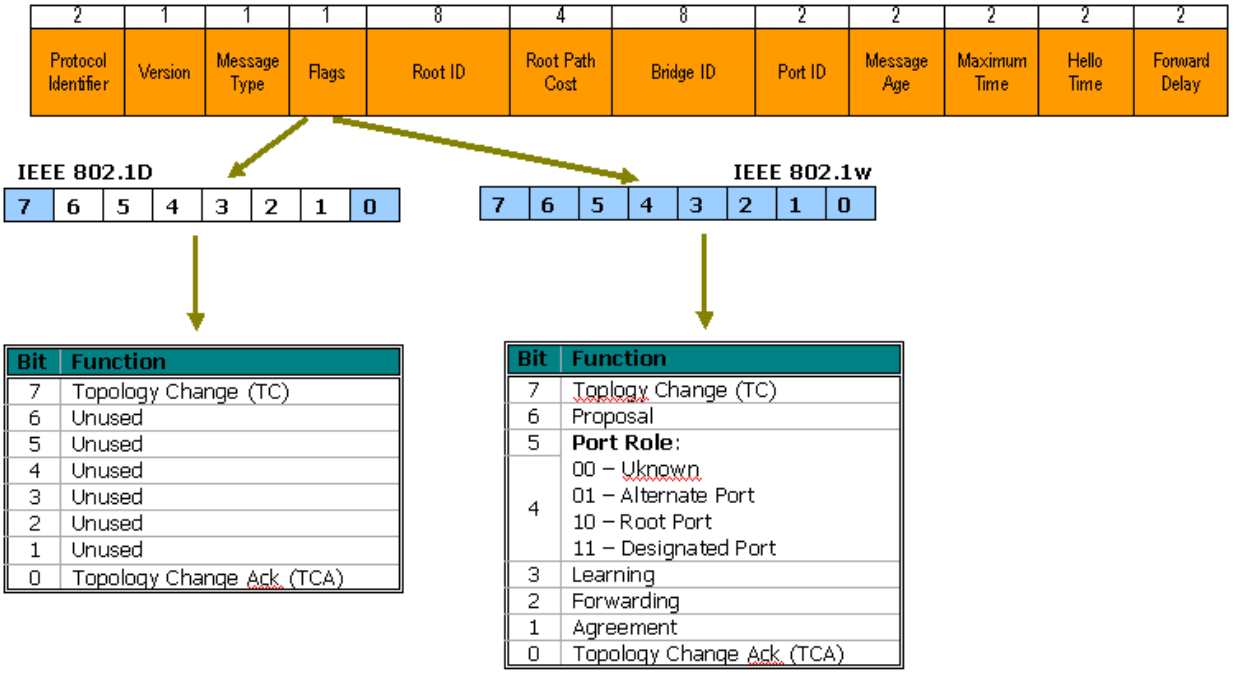

Bit	Function
7	Topology Change (TC)
6	Unused
5	Unused
4	Unused
3	Unused
2	Unused
1	Unused
0	Topology Change Ack (TCA)

Bit	Function
7	Toplogy Change (TC)
6	Proposal
5	**Port Role**:
4	00 – Uknown 01 – Alternate Port 10 – Root Port 11 – Designated Port
3	Learning
2	Forwarding
1	Agreement
0	Topology Change Ack (TCA)

802.1D STP BPDU Format

```
▶ Frame 9: 60 bytes on wire (480 bits), 60 bytes captured (480 bits)
▼ IEEE 802.3 Ethernet
  ▶ Destination: Spanning-tree-(for-bridges)_00 (01:80:c2:00:00:00)
  ▶ Source: Cisco_f2:c3:19 (00:12:da:f2:c3:19)
    Length: 38
    Padding: 0000000000000000
▶ Logical-Link Control
▼ Spanning Tree Protocol
    Protocol Identifier: Spanning Tree Protocol (0x0000)
    Protocol Version Identifier: Spanning Tree (0)
    BPDU Type: Configuration (0x00)
  ▶ BPDU flags: 0x00
  ▶ Root Identifier: 4096 / 1 / 00:12:da:f2:c3:00
    Root Path Cost: 0
  ▶ Bridge Identifier: 4096 / 1 / 00:12:da:f2:c3:00
    Port identifier: 0x8019
    Message Age: 0
    Max Age: 20
    Hello Time: 2
    Forward Delay: 15
```

802.1W RSTP BPDU Format

- Frame 19: 60 bytes on wire (480 bits), 60 bytes captured (480 bits) on interface 0
- IEEE 802.3 Ethernet
  - Destination: Spanning-tree-(for-bridges)_00 (01:80:c2:00:00:00)
  - Source: Cisco_f2:c3:19 (00:12:da:f2:c3:19)
    Length: 39
    Padding: 00000000000000
- Logical-Link Control
- Spanning Tree Protocol
    Protocol Identifier: Spanning Tree Protocol (0x0000)
    Protocol Version Identifier: Rapid Spanning Tree (2)
    BPDU Type: Rapid/Multiple Spanning Tree (0x02)
  - BPDU flags: 0x3c (Forwarding, Learning, Port Role: Designated)
  - Root Identifier: 0 / 1 / 00:12:da:f2:c3:00
    Root Path Cost: 0
  - Bridge Identifier: 0 / 1 / 00:12:da:f2:c3:00
    Port identifier: 0x8019
    Message Age: 0
    Max Age: 20
    Hello Time: 2
    Forward Delay: 15
    Version 1 Length: 0

802.1S MSTP BPDU Format [1]

Common and Internal Spanning Tree (CIST)

Every MSTP region runs special instance of spanning-tree known as IST or Internal Spanning Tree (=MSTI0). This instance mainly serves the purpose of disseminating STP topology information for MSTIs. IST has a root bridge, elected based on the lowest Bridge ID (Bridge Priority + MAC address). The situation changes with multiple MSTP regions in the network. When a switch detects BPDU messages sourced from another region (or STP/PVST+ BPDU), it marks the corresponding port as MSTP **boundary**. For the convenience, we would call all other ports as “internal”. A switch that has boundary ports is known as boundary switch.

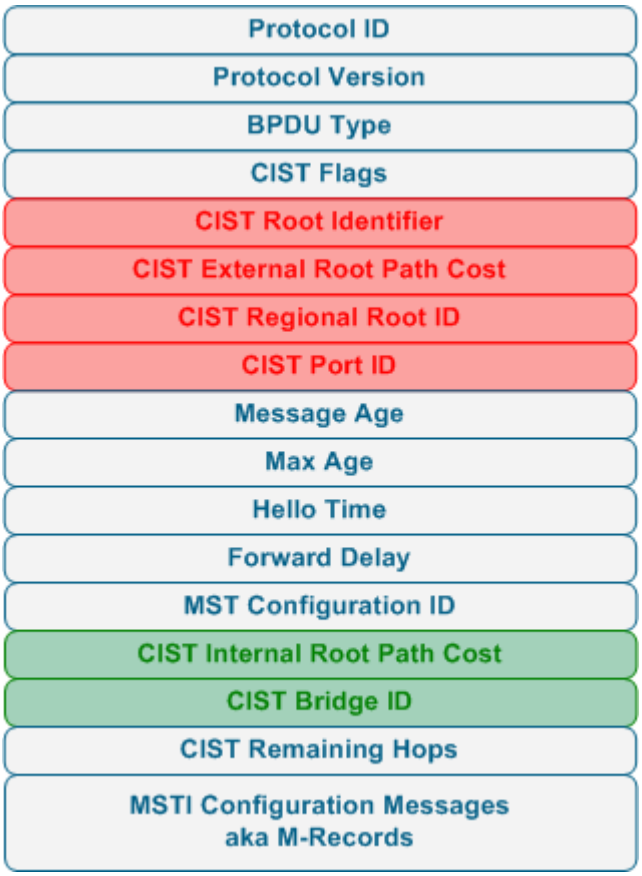

The MSTP BPDU contains two important block of information. One, highlighted in red, is related to CIST Root and CIST Regional Root election. As you will see later, CIST Root is elected among all regions and CIST Regional Root is elected in every region. The green block outlines the information about CIST Regional Root (which becomes the IST Root in presence of multiple regions). The CIST Internal Root path cost is the intra-region cost to reach the CIST Regional Root.

[1]Petr Lapukhov http://blog.ine.com/2010/02/22/understanding-mstp/

It is important to keep in mind that IST Root = CIST Regional Root in case where multiple regions interoperate. This transformation is explained further in the text. Now, to define the CIST Root and CIST Regional Root roles:

- CIST Root is the bridge that has the lowest Bridge ID among ALL regions. This could be a bridge inside a region or a boundary switch in a region.
- CIST Regional Root is a **boundary switch** elected for every region based on the shortest external path cost to reach the CIST Root. Path cost is calculated based on costs of the links connecting the regions, **excluding** the internal regional paths. CIST Regional Root becomes the root of the IST for the given region as well.

CIST Root Bridges Election Process

- When a switch boots up, it declares itself as **CIST Root** and **CIST Regional Root** and announces this fact in outgoing BPDUs. The switch will adjust its decision upon reception of better information and continue advertising the best known CIST Root and CIST Regional Root on all **internal ports**. On the boundary ports, the switch advertises **only** the CIST Root Bridge ID and CIST External Root Path Cost thus hiding the details of the region's internal topology.
- **CIST External Root Path Cost** is the cost to reach the CIST Root across the links connecting the boundary ports – i.e. the inter-region links. When a BPDU is received on an internal port, this cost is **not changed**. When a BPDU is received on a boundary port, this cost is adjusted based on the receiving boundary port cost. In result, the CIST External Root Path Cost is propagated unmodified inside any region.
- Only a boundary switch could be elected as the CIST Regional Root, and this is the switch with the **lowest cost** to reach the CIST Root. If a boundary switch hears better CIST External Root Path cost received on its internal link, it will relinquish its role of CIST Regional Root and start announcing the new metric out of its boundary ports.
- Every boundary switch needs to properly block its boundary ports. If the switch is a CIST Regional Root, it elects one of the boundary ports as the "CIST Root port" and blocks all other boundary ports. If a boundary switch is not the CIST Regional Root, it will mark the boundary ports as CIST Designated or Alternate. The boundary port on a non regional-root bridge becomes designated only if it has superior information for the CIST Root: better External Root Path cost or if the costs are equal better CIST Regional Root Bridge ID. This follows the normal rules of STP process.
- As a result of CIST construction, every region will have **one switch** having single port unblocked in the direction of the CIST Root. This switch is the CIST Regional Root. All boundary switches will advertise the region's CIST Regional Root Bridge ID out of their non-blocking boundary ports. From the outside perspective, the whole region will look like a single virtual bridge with the Bridge ID = CIST Regional Root ID and single root port elected on the CIST Regional Root switch.
- The region that contains the CIST Root will have all boundary ports unblocked and marked as CIST designated ports. Effectively the region would look like a virtual root bridge with the Bridge ID equal to CIST Root and all ports being designated. Notice that the region with CIST Root has CIST Regional Root equal to CIST Root as they share the same lowest bridge priority value across all regions.

802.1S MSTP BPDU Format

```
⊞ Frame 59: 135 bytes on wire (1080 bits), 135 bytes captured (1080 bits)
⊟ IEEE 802.3 Ethernet
  ⊞ Destination: Spanning-tree-(for-bridges)_00 (01:80:c2:00:00:00)
  ⊞ Source: Cisco_f2:c3:19 (00:12:da:f2:c3:19)
    Length: 121
⊞ Logical-Link Control
⊟ Spanning Tree Protocol
    Protocol Identifier: Spanning Tree Protocol (0x0000)
    Protocol Version Identifier: Multiple Spanning Tree (3)
    BPDU Type: Rapid/Multiple Spanning Tree (0x02)
  ⊞ BPDU flags: 0x7c (Agreement, Forwarding, Learning, Port Role: Designated)
  ⊞ Root Identifier: 32768 / 0 / 00:12:da:f2:c3:00
    Root Path Cost: 0
  ⊞ Bridge Identifier: 32768 / 0 / 00:12:da:f2:c3:00
    Port identifier: 0x8019
    Message Age: 0
    Max Age: 20
    Hello Time: 2
    Forward Delay: 15
    Version 1 Length: 0
    Version 3 Length: 80
  ⊟ MST Extension
      MST Config ID format selector: 0
      MST Config name: bpdu_tests
      MST Config revision: 1
      MST Config digest: 870555c957f1b44530b7d56fd4716adf
      CIST Internal Root Path Cost: 0
    ⊟ CIST Bridge Identifier: 32768 / 0 / 00:12:da:f2:c3:00
        CIST Bridge Priority: 32768
        CIST Bridge Identifier System ID Extension: 0
        CIST Bridge Identifier System ID: 00:12:da:f2:c3:00
      CIST Remaining hops: 20
    ⊟ MSTID 1, Regional Root Identifier 32768 / 00:12:da:f2:c3:00
      ⊞ MSTI flags: 0x7c (Agreement, Forwarding, Learning, Port Role: Designated)
        MSTID 1, priority 32768 Root Identifier 00:12:da:f2:c3:00
        Internal root path cost: 0
        Bridge Identifier Priority: 8
        Port identifier priority: 8
        Remaining hops: 20
```

Bridge ID Format