802.1Qcc:
Topology Discovery

Rodney Cummings

National Instruments

Agenda

- Introduction
  - Assumptions for what are we trying to do
- Research
  - Overview of what is out there now
- Recommendations

Introduction

Introduction

- Qcc supports centralized configuration/control
  - In addition to distributed
- All TSN features have managed objects
  - e.g. Qbv, Qbu, AS, CB, …
  - Transitioning MIB to YANG
- Remote management is one clear near-term solution

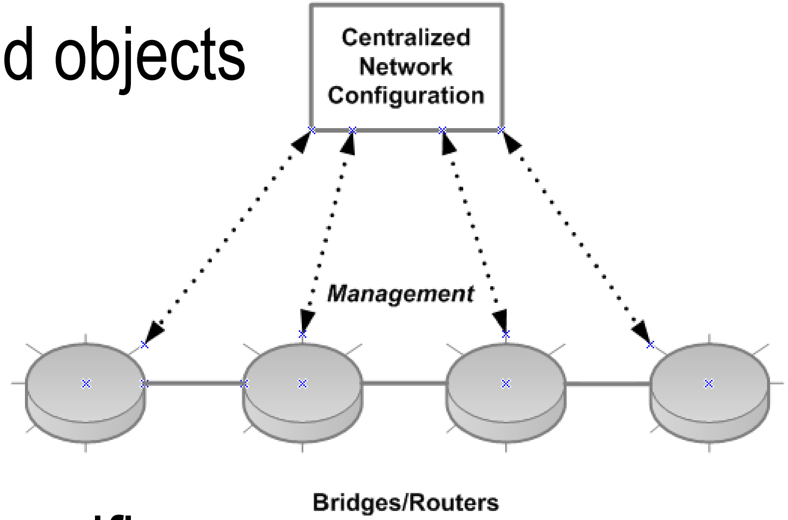

- Like any Q amendment, Qcc specifies managed objects to meet its goals
  - Gaps: static reservation, bridge delay, TE-MSTID (nail-up)
  - Last gap to resolve: Physical topology discovery

Summary of Goals from TSN Meetings

1. Discover systems: end stations, bridges <u>and</u> routers
    - Including router that does not support 802.1 protocols
2. No protocol mandates in Qcc
    - Support what is out there
3. Keep it simple and complete (always works)
    - Focus on common standards
4. Fundamental information needed is:
    - Persistent ID for each TSN-capable system and its ports
        - Persistent (non-volatile) as long as port exists, including reboots
    - Physical connectivity of each port to its neighboring port
    - Address for use with remote management protocol(s)

New Proposed Goals

5. Re-use discovered IDs in Qcc UNI
6. Support discovery from out-of-box
   - Do not require IT-style management as a precondition
   - Sensor with 2 ports, connect to industrial controller, and go

Research

Three Categories

- Management data models
  - MIB and YANG
- Protocols to control active topology
  - Spanning tree protocols, Interior Gateway Protocols (IGP), …
- Protocols for topology discovery
  - LLDP, …

Management: MIB (1 of 2)

- Management address
  - Typically IP (UDP), but MAC is possible
    - E.g. SNMP over 802 (RFC 4789)
  - Presumably not guaranteed to be persistent
- IEEE8021-BRIDGE-MIB
  - ComponentID: Multiple per bridge, each of which has
    - Bridge Address: MAC address of bridge
  - Port has
    - Port Number (1..n)
      - Not required to be consecutive (“holes” can exist)
    - Individual MAC address
    - ifIndex: For use with IETF IF-MIB

Management: MIB (2 of 2)

- IETF RFC 2863 (Interface MIB, aka IF-MIB)
  - ifTable of interfaces
  - ifIndex (index to ifTable) not required to be consecutive
    - Use SNMP GetNext to skip over holes
    - Too dynamic for topology discovery (see RFC 2922 Design Goals)
  - Each ifEntry (interface entry) has
    - ifName: Read-only name assigned by hardware
      - Multiple interfaces can use same ifName
    - ifAlias: Persistent name writable by management; empty out-of-box
    - ifPhysAddress: For 802 this is MAC address
- IETF RFC 6933 (Entity MIB v4)
  - Physical (and logical) info about router and its ports

Management: YANG

- IETF RFC 7223 (Interface Mgmt, analogous to IF-MIB)
  - List of interfaces keyed by name, not index
    - YANG is not limited to indexed table; List eliminates holes
  - Name must be unique to the server (bridge/router)
    - System-controlled interface: Bridge/router decides the name
    - User-controlled interface: Name provided to ‘create’
    - Presumably persistent in startup datastore
  - Layered: Physical and logical interfaces
  - Relationship to IF-MIB
    - “interface-ref” typedef for name is same as IF-MIB ifName, but only if the system did not support same ifName for different interfaces
    - “if-index” is a read-only variable for ifIndex (for IF-MIB support)
    - Optional “description” is similar to ifAlias

Active Topology: 802.1

- Bridge identified using its MAC address
  - ISIS-SPB System ID
  - RSTP/MSTP Bridge ID
- Port identified using its Port Number
- No management address

Active Topology: IETF IGP

- OSPF-TE (RFC 3630)
  - Router Address TLV
    - "stable IP address of advertising router that is always reachable if there is any connectivity to it"
    - "typically implemented as a "loopback address" "
    - "known in other standards as "router ID" "
    - Unclear if globally unique; unclear if usable as management address
  - Link TLV uses "interface IP address" for physical port
- GMPLS (RFC 4202)
  - Unnumbered link: For interface to point-to-point link, 32-bit number assigned locally by router
    - Similar concept to Port Number and ifIndex

Active Topology: IETF IGP

- IS-IS TE (RFC 3784)
  - IP address for each interface
  - 4-octet Router ID (same as OSPF TE)
- IS-IS for GMPLS (RFC 4205)
  - Link ID uses “unnumbered link” of GMPLS (same as OSPF)

Topology Discovery: LLDP (1 of 5)

- 802.1AB: Simple protocol with a MIB
- Transmit my local info (bridge and its ports)
  - Local LLDP MIB
- Receive neighbor's local info
  - Store in Remote LLDP MIB
  - No propagation
    - (typical)
- MIB is superset of IETF RFC 2922 (PTOPO MIB)

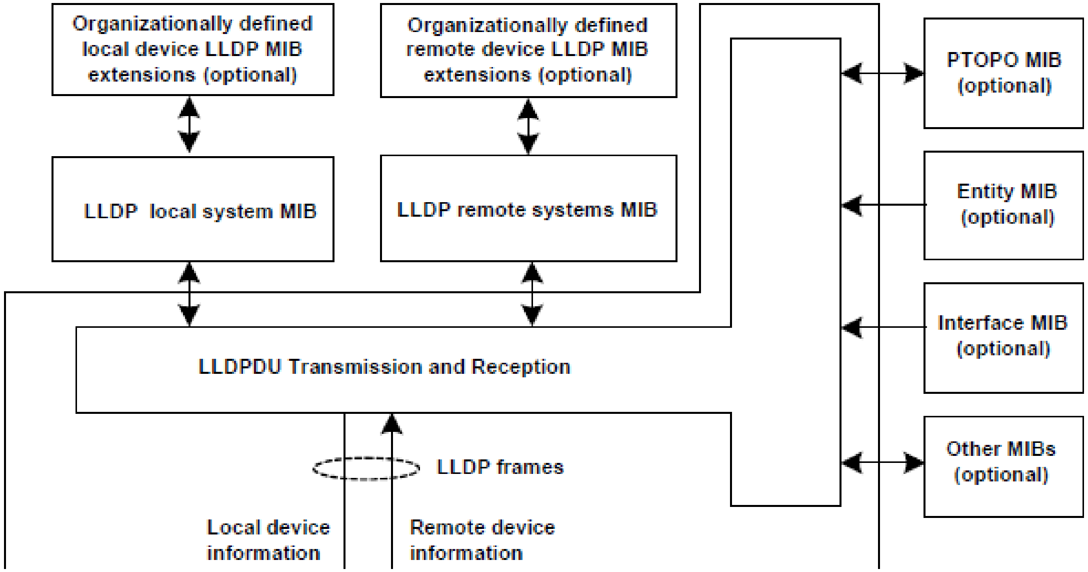

Topology Discovery: LLDP (2 of 5)

- Mandatory TLVs
  - Chassis ID (bridge/router)
  - Port ID
  - Time To Live, End of LLDPDU
- Optional TLVs
  - Management Address
  - System Capabilities

Topology Discovery: LLDP (3 of 5)

- Management Address
  - Subtype is [IANA Address Family Number]
  - "returned address should be the most appropriate for management use, typically a layer 3 address such as the IPv4 address"
    - IPv6 and MAC also possible
  - Needed for TSN
  - Optional OID and interface num (ifIndex or Port Num)
    - Not necessarily needed for YANG (TSN)
- System Capabilities
  - 16 bit map, one for (End) Station Only
    - Needed for TSN to identify potential talkers/listeners

Topology Discovery: LLDP (4 of 5)

- Port ID
  - One subtype by “preferred use”

Table 8-3—port ID subtype enumeration

ID subtype	ID basis	References
0	Reserved	—
1	Interface alias	ifAlias (IETF RFC 2863)
2	Port component	entPhysicalAlias when entPhysicalClass has a value ‘port(10)’ or ‘backplane(4)’ (IETF RFC 4133)
3	MAC address	MAC address (IEEE Std 802)
4	Network address	networkAddress [a]
5	Interface name	ifName (IETF RFC 2863)
6	Agent circuit ID	agent circuit ID (IETF RFC 3046)
7	Locally assigned	local [b]
8–255	Reserved	—

Alias must be set by mgmt; Bad for goal #6 (out-of-box)

MAC/IP; Must search for it; Bad for goal #3 (simple)

YANG IF name (if unique); **Preferred for TSN?**

Topology Discovery: LLDP (5 of 5)

- Chassis ID
  - One subtype by "preferred use"

Table 8-2—chassis ID subtype enumeration

ID subtype	ID basis	Reference
0	Reserved	—
1	Chassis component	EntPhysicalAlias when entPhysClass has a value of 'chassis(3)' (IETF RFC 4133)
2	Interface alias	IfAlias (IETF RFC 2863)
3	Port component	EntPhysicalAlias when entPhysicalClass has a value 'port(10)' or 'backplane(4)' (IETF RFC 4133)
4	MAC address	MAC address (IEEE Std 802)
5	Network address	networkAddress [a]
6	Interface name	ifName (IETF RFC 2863)
7	Locally assigned	local[b]
8–255	Reserved	—

Alias must be set by mgmt; Bad for goal #6 (out-of-box)

MAC is persistent and globally unique; **Preferred for TSN?**

YANG IF name is not globally unique

IP address possibly not persistent (unless it is "router ID")

Topology Discovery: Others

- Routers discover topology using the IGP (IS-IS, OSPF)
  - Router does not run LLDP unless it also supports bridging
- Several proprietary protocols
- When all else fails, use a toolbox
  - Ping, Traceroute, DNS, monitoring, …

- YANG data model for network topologies
  - draft-clemm-i2rs-yang-network-topo
  - Each "level" has details in augment (e.g. "Service", L2, IS-IS)
  - Assumes a central topology-discovery entity

Summary: End Station

- Typically do not run management server
  - LLDP or similar for nearest bridge/router to explicitly discover
- MAC address is persistent and globally unique
  - Each interface has MAC address
- IP addresses not necessarily persistent
  - Exception: IPv6 unicast using MAC address as interface ID

Summary: Bridge

- We can assume a management server
- MAC address is persistent and globally unique
  - Available in management, active topology, and LLDP
- Port identification can use YANG IF name
  - MIB and active topology use number, but
    YANG management and LLDP can focus on name
- We can assume LLDP
  - Provides management address, chassis ID, and port ID
    - <u>Everything we need</u>
  - For TSN, we may want to specify an LLDP “profile”

Summary: Router

- “Router ID” seems to be consistent in IGP
  - Presumably persistent and unique to area
  - If we use IANA Address Family, this is covered by IPv4
- Port identification can use YANG IF name
  - Same rationale as bridge
  - IGP port identification (number or IP address) can be mapped to this name
- Management address is a challenge
  - IGP doesn’t explicitly provide management address
  - We cannot mandate LLDP protocol for routers

Recommendations

Rec #1: LLDP profile for bridging

- Organization can specify a profile of LLDP for TSN
  - LLDP protocol required for end stations and bridges
  - Management required for bridges (not end stations)
  - Require some optional features
    - Tx and Rx
    - Management Address TLV
    - System Capabilities TLV (to detect end station)
    - Chassis ID subtype = MAC address
    - Port ID subtype = ifName (unique)
- Fully interoperable

Rec #2: YANG Net Topology to CNC

- As YANG Data Model for Network Topologies takes shape, help ensure that augments for L2/L3 provide
  - Persistent ID for end station, bridge, and router
  - Management address for bridge/router
- TSN CNC can specify this data model as an input
  - Presumably defer this specification until RFC is ready
- Solves topology discovery for some router use cases
  - Central topology-discovery entity
    uses proprietary and/or toolbox techniques
  - Not necessarily interoperable

Rec #3: New RFC for PTOPO YANG

- IETF RFC 2922 specified a MIB for physical topology (PTOPO MIB)
  - MIB but not “mechanisms” (protocol or toolbox)
  - Predecessor to LLDP MIB (local and remote)
  - Informative
- New IETF project could refresh this concept for YANG
  - Router can discover its own connectivity (e.g. IGP)
  - Central topology-discovery entity can also populate
  - Ideally, this YANG is the same as 802.1AB’s YANG
- Note: This was originally proposed as part of 802.1Qcc, but it is best done in an IETF RFC and/or 802.1AB

Rec #4: IDs for 802.1Qcc UNI

- Purpose is for CNC to identify a port in its topology
- End station interface
  - MAC address as the primary type
  - Optional persistent IDs added on: IP address, number, …
- Bridge/router
  - Single persistent address as primary type
    - MAC or IP (“router ID”)
  - Management address uses IANA Family (like LLDP)
- Port of bridge/router
  - YANG interface name as primary type
  - Optional persistent IDs added on: MAC, IP, number, …

Thank You