



www.huawei.com

IEEE 1904.2 Discovery Protocol

This is NOT a Baseline Proposal

Authors: Hesham ElBakoury

Version: 1.0

HUAWEI TECHNOLOGIES CO., LTD.



GOAL

Present a high-level outline of 1904.2 discovery Protocol and its parameters assuming the access network definition of IEEE 802.1CF.

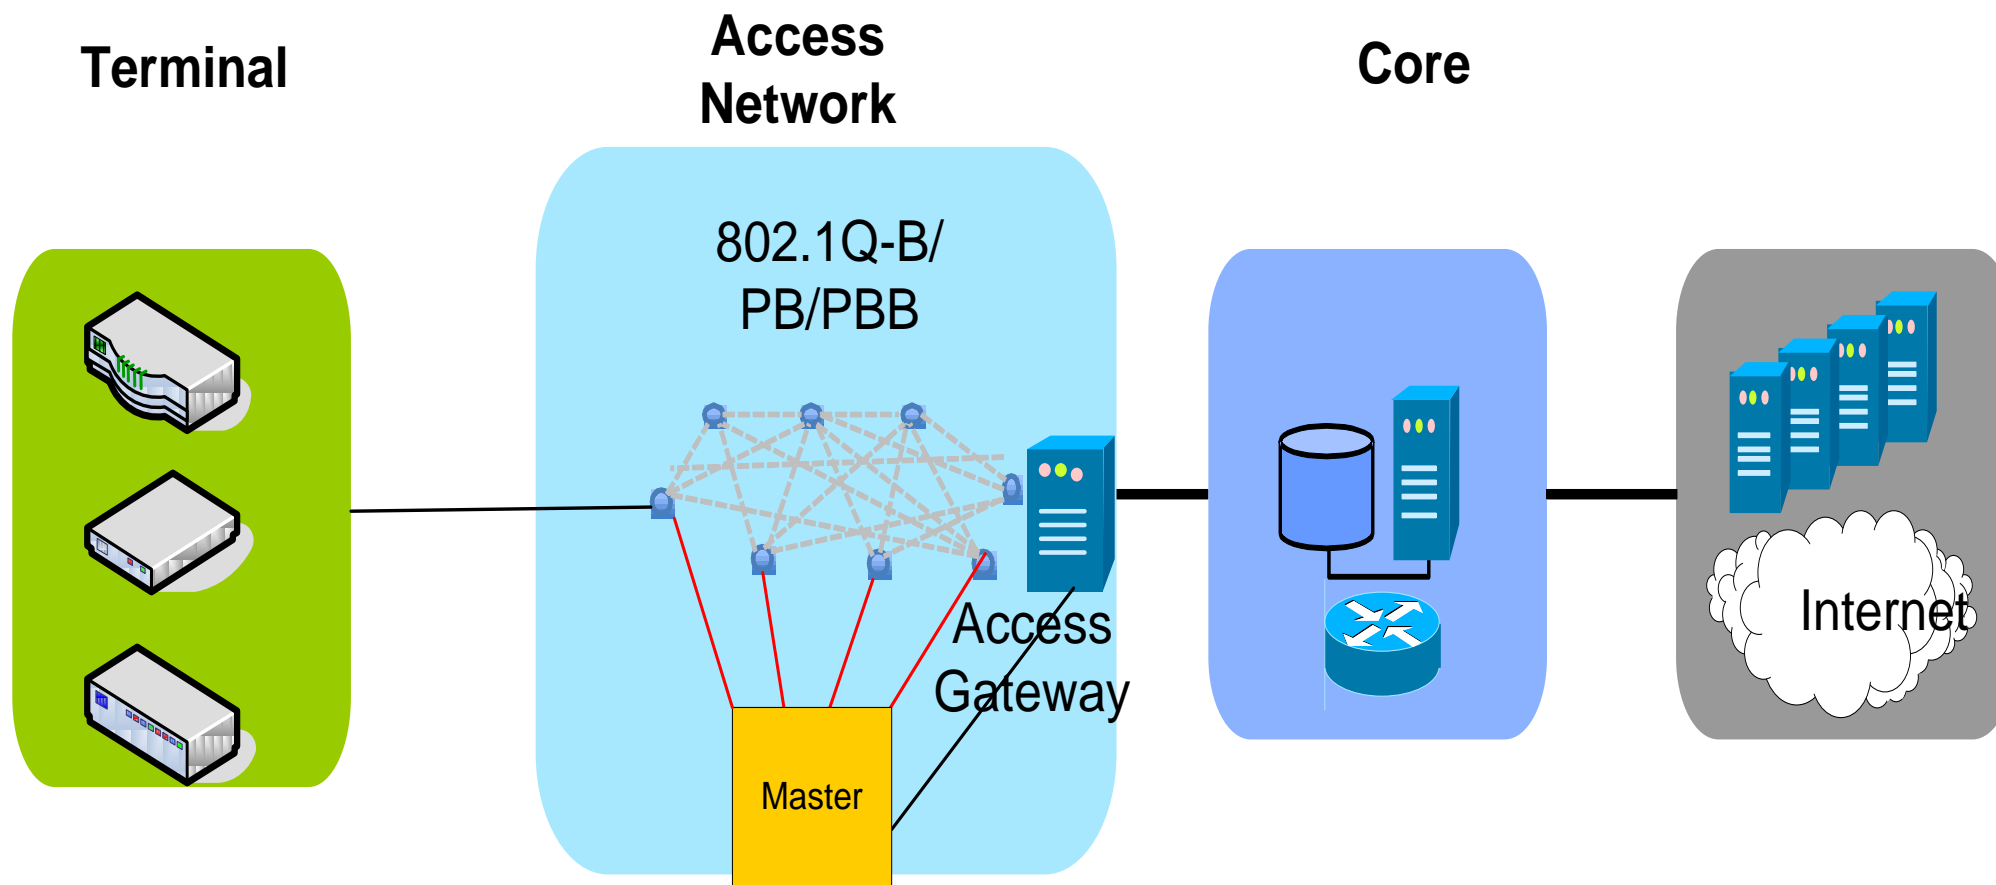
The protocol uses both Master advertisement and client solicitation messages.

The protocol discussion should help come up with a baseline proposal for the next IEEE 1904.2 meeting.

Agenda

- **IEEE 802.1CF Definition of Access Network Boundary**
- **The 1904.2 discovery protocol is run before the discovery/registration procedure that is implemented by the protocol tunneled by 1904.2 UMT.**
- **IEEE 1904.2 Discovery Protocol**
 - IEEE 1904.2 Master Specification.
 - IEEE 1904.2 Client Specification.
- **Discuss protocol timers and parameters.**
- **Forwarding and Identifying IEEE 1904.2 Discovery Protocol Messages.**

IEEE 802.1CF Definition of Access Network Boundary



- More than one Master can be available for redundancy and load balancing.
- Master can be connected to the access network via one or more interfaces (?)
- A client can be connected to the access network via one or more interfaces (?)

IEEE 1904.2 Master Discovery Protocol

- **Management Master/Controller discovery uses the following messages:**
 - **Master Advertisements.**
 - **Master Solicitations.**
- **The master periodically multicasts/broadcasts a “Master Advertisement” message from the interface(s) connected to the access network.**
 - The advertisement message announces the MAC address of the interface and master capabilities.
- **Clients discover the addresses of the Management Master by listening for advertisements.**
- **When a client starts up, it may multicast/Broadcast a “Master Solicitation” message to ask for immediate advertisements instead of waiting for the next periodic ones to arrive.**

1

Master Specification

Advertisement Parameters

- **MaxAdvertisementInterval**
 - The maximum time allowed between sending multicast Advertisements from the interface, in seconds.
 - Must be no less than 4 seconds and no greater than 1800 seconds (?)
 - Default: 10 Minutes (?)
- **MinAdvertisementInterval**
 - The minimum time allowed between sending unsolicited multicast Advertisements from the interface, in seconds.
 - Must be no less than 3 seconds and no greater than MaxAdvertisementInterval.
 - Default: $0.7 * \text{MaxAdvertisementInterval}$ (?)

Advertisement Message

- **Master Advertisements contain the following information:**
 - In the destination address field of the MAC header: the interface's configured **MasterAdvertisementAddress** multicast address.
 - The source address field is the MAC address of the Master.
 - **The Lifetime field:** the interface's configured **MasterAdvertisementLifetime**. This field, specifies the maximum length of time that the advertised address is to be considered as valid Master address by clients in the absence of further advertisements.
 - This field is used to ensure that clients eventually forget about Masters that fail, become unreachable, or stop acting as Masters.
 - This field must be no less than **MaxAdvertisementInterval** and no greater than 3 hours (?). The default should be $3 * \text{MaxAdvertisementInterval}$ (30 minutes ?)
 - **Master Capabilities** (e.g. the management protocols it uses, openflow, SNMP, TR-69).
 - When to expect the next advertisement.

Master Behavior (1)

- **Master starts the transmission of periodic advertisements on an interface in the following cases:**
 - Master startup,
 - As a result of recovery from an interface failure, or through actions of system management such as:
 - Enabling the interface after it has been administratively disabled.
 - Changing the system from being a client to being a Master
- **The random interval of the first few advertisements `MAX_INITIAL_ADVERTISEMENTS` (3 ?) is limited to `MAX_INITIAL_ADVERT_INTERVAL` (16 seconds ?)**
 - Using this smaller interval for the initial advertisements increases the likelihood of a Master being discovered quickly when it first becomes available, specially in the presence of possible packet loss or temporary link partitioning.
- **Links that suffer high packet loss rates or frequent partitioning are accommodated by increasing the rate of the periodic advertisements (rather than client solicitations).**

Master Behavior (2)

- In addition to the periodic, unsolicited advertisements, a Master sends advertisements in response to valid solicitations received on any of its advertising interfaces.
- A Master may choose to unicast the response directly to the soliciting client's address or multicast it to the interface's configured **MasterAdvertisementAddress**;
 - In the latter case, the interface's interval timer is reset to a new random value, as with unsolicited advertisements.
 - A unicast response may be delayed, and a multicast response must be delayed, for a small random interval not greater than **MAX_RESPONSE_DELAY (2 seconds ?)**, in order to prevent synchronization with other responding Masters, and to allow multiple, closely-spaced solicitations to be answered with a single multicast advertisement.

Master Behavior (3)

- **A Master interface may also cease to be an advertising interface, through actions of system management such as:**
 - Administratively disabling the interface
 - Shutting down the system,
 - Changing the system from Master to Client.
- **In such cases, the Master may transmit a final multicast advertisement on the interface, identical to its previous advertisement but with a Lifetime field of zero.**

Master Advertisements on Shared Link

- **The advertisements are not strictly periodic:**
 - The interval between subsequent transmissions is randomized to reduce the probability of synchronization with the advertisements from other Masters on the same link.
 - This is done by maintaining a separate transmission interval timer for each advertising interface.
 - Each time a multicast advertisement is sent from an interface, that interface's timer is reset to a uniformly-distributed random value between the interface's configured **MinAdvertisementInterval** and **MaxAdvertisementInterval**;
 - The unique MAC address of the Master can be used as the seed used to initialize their pseudo-random number generators.
 - The actual randomly-chosen value should be in units of the highest available timer resolution
- **Expiration of the timer causes the next advertisement to be sent from the interface, and a new random value to be chosen.**

2

Client Specification

Client Solicitation Message

- **Client Capabilities:**
 - Protocols supported by Client (e.g. SNMP, TR-69, Openflow).

Client Behavior

- When a client attached to the access network starts up, it may multicast a Master Solicitation to ask for immediate advertisements, rather than waiting for the next periodic ones to arrive;
- A client is permitted (but not required) to transmit up to **MAX_SOLICITATIONS (3)** Master Solicitation messages from any of its multicast interfaces after any of the following events:
 - The interface is initialized at system startup time.
 - The system changes from being a Master to being a client by system management.
 - The interface is reinitialized after a temporary interface failure or after being temporarily disabled by system management.

Client Solicitation on a Shared Link

- If a client does choose to send a solicitation, it should delay that transmission for a random amount of time between 0 and **MAX_SOLICITATION_DELAY** (1 second) .
 - This serves to alleviate congestion when many clients start up on a link at the same time, such is the case after recovery from a power failure
 - It is recommended that clients include some unique value, such as one of their MAC addresses, in the seed used to initialize their pseudo-random number generators.
 - The actual randomly-chosen value should in units of the highest available timer resolution.

3

Multicast MAC Address

Discovery Multicast Addresses

- **Discovery frames are identified by a unique Ethertype the value.**
- **Discovery frames processing and filtering is based on the Ethertype unicast and multicast DA.**
- **Reserved Multicast address in Table 8-1 (IEEE 802.1Q-Rev/D2.0) can be used for Master Advertisement.**
 - Advertisements are processed by clients and forwarded by bridges.
- **Reserved Multicast address in Table 8-1 (IEEE 802.1Q-Rev/D2.0) can be used for Client Solicitation.**
 - Solicitations are processed by Masters and forwarded by bridges.

Thank You