19.3 Connectivity Fault Management ~~Protocol~~ Protocol overview

Connectivity Fault Management can be sub-divided into the following categories:

- Fault ~~Detection~~detection
- Fault ~~Verification~~verification
- Fault ~~Isolation~~isolation
- Fault ~~Notification~~notification
- Fault ~~Recovery~~recovery

Fault ~~Detection deals with mechanism(s) that~~ detection mechanisms can detect both hard ~~failures~~faults, such as bridge failures or link ~~and node~~ failures, and soft ~~failures~~faults, such as software failure, memory corruption, mis-configuration, etc. ~~Typically a lightweight protocol is desirable to detect the~~ Following fault ~~and thus~~ detection, it ~~would~~ may be ~~prudent~~ desirable to verify the fault via ~~Fault Verification~~ fault verification mechanism before taking additional steps ~~in isolating~~ to isolate the fault. After verifying that a fault has occurred along the data path, it ~~is important to~~ may be ~~able~~ desirable to isolate the fault to a given ~~node~~ bridge or link (e.g., diagnose the fault~~). Therefore, a Fault Isolation mechanism is needed in Fault Management~~) using fault isolation mechanism. Fault ~~Notification~~ notification mechanism can be used in conjunction with ~~Fault Detection~~ fault detection mechanism to notify the ~~upstream and downstream nodes of a fault~~client layers about faults detected in server layer in layered networks. ~~Finally,~~ Fault ~~Recovery deals~~ recovery mrchanisms deal with recovering from the detected ~~failure~~ faults by switching to an alternate available ~~node~~ bridge and/or ~~link (e.g., node redundancy or link redundancy)~~link.

The scope of this clause is limited to the first three aspects of the Connectivity Fault Management, i.e. ~~Fault Detection~~fault detection, ~~Fault Verification~~ fault verification and ~~Fault Isolation~~fault isolation, in the context of Provider Bridged networks.~~40~~

19.3.1 Fault Detection

Continuity Check (CC) provides a means to detect both hard ~~failures~~ and soft ~~failures such as software failure, memory corruption, or mis-configuration~~faults. ~~The failure~~ Fault detection is achieved by each ~~Maintenance End Point (MEP)~~ MEP transmitting a ~~heartbeat message periodically for each customer service instance~~ CC Message (~~e.g., SVLAN)~~CCM) periodically. ~~Therefore~~As a result, ~~each edge Provider Bridge receives a set of heartbeat messages~~ MEP(s) also receive CCMs periodically from other ~~edge Provider Bridges of that service instance~~peer MEPs. ~~Once the~~ When a MEP on a local ~~PB~~ bridge stops receiving the periodic ~~heartbeats~~ CCM from ~~the~~ a peer MEP on a remote ~~PB~~bridge, it ~~assumes~~ can assume that either the remote ~~PB~~ bridge has failed or ~~an unrecoverable~~ a failure ~~on~~ in the continuity of the path has ~~happened~~occured. The ~~PB~~ bridge can subsequently notify the ~~operator of~~ network management application about the failure, using mechanisms that are out of scope of this clause, and initiate ~~the failure~~ fault verification ~~and~~ and/or fault isolation steps either automatically or through operator command.

If a ~~PB~~ bridge is put out of commission, ~~then in order to avoid triggering false failure detection, the out-ofcommissioned PB shall~~ this bridge may indicate its soon to be out-of-~~state~~ commission status to other ~~member PBs~~ peer bridges for each ~~service instance that it participates through a flag~~ Virtual Bridge LAN Service supported across this bridge to avoid triggering false fault detections. This may be done via indications in ~~the CC message~~CCMs. ~~The other member PBs of the service instance~~Other peer bridges, upon receiving this indication, would deactivate the corresponding timer for the ~~heartbeat of that PB~~CCMs. Once ~~PB devices have~~ MEPs has received and processed the ~~CC messages~~CCMs, each ~~PB~~ MEP will have a view of all ~~active PBs~~ other peer MEPs for a given ~~customer service instance~~Virtual Bridge LAN Service.

Upon receiving ~~CC messages, at the receiving~~ a CCM from a remote MEP, a CC validity timer is started at the receiving MEP which is used to determine the loss of ~~CC messages~~CCMs. A ~~CC~~ CCM loss is assumed when the next ~~CC message~~ CCM from a remote MEP is not received within the timeout of this validity timer. If n consecutive ~~CC messages~~ CCMs are lost, a ~~fault for~~ continuity to that remote MEP is assumed to have failed and a fault is detected. Subsequent fault verification and fault isolation procedures can be exercised.

~~The~~ A hard fault may ~~correspond to a hard failure or a soft failure within the network. Also a hard failure may~~ possiblty result in network isolation which leads to loss of ~~CC messages~~ CCMs for many ~~customer service instances~~Virtual Bridge LAN services. If the hard ~~failure~~ fault can be detected and ~~reported~~ notified to ~~the Management entity~~a management application, additional notifications by each MEP may not be needed ~~a~~ e.g., it ~~is~~ may be desirable to have an alarm suppression mechanism for notifications that get generated as the result of ~~CC~~ CCM timeouts~~. Since this message is sent periodically, in order to facilitate the processing and filtering of this message, both the message type and domain level is embedded in the multicast MAC address~~.

A ~~CC messages~~ CCM does not require a response and a multicast CCM requires only o(~~n~~N) ~~message transmission~~ transmissions within its member group, where N is the number of members within the member group. In other words, if a ~~service instance~~ Virtual Bridge LAN Service has N member ~~PBs~~MEPs, only N ~~CC messages~~ CCMs need to be

transmitted periodically ñ - one from each PBMEP. However, if this was to be done by point-to-point Ping messagesping messages were used, then o(N**^2) messages would have been required.

The Maintenance End Points MEPs shall allow the filtering of CC messages CCMs from either entering or exiting its OAM maintenance domain.

19.3.2 Fault Verification

A unicast Loopback message (LBM) is used for fault verification. To verify the connectivity between Maintenance End a MEP and Intermediate pointsits peer MEP or a MIP, the Loopback request message LBM is initiated by a MEP with a DA MAC address set to the MAC address of either a MIP or the peer MEP. The receiving MIP or MEP shall respond to the Loopback Request LBM with a Loopback response upon verification of the messageReply (LBR). The MEPs shall allow the filtering of fault verification messages from either entering or exiting its OAM domain.

MEPs shall allow filtering of LBMs and LBRs from either entering or exiting its maintenance domain.

19.3.3 Fault Isolation

The Linktrace function (LT) mechanism is used to isolate faults visible at the Ethernet MAC layer. Linktrace can be used to isolate a fault associated with a given customer service instanceVirtual Bridge LAN Service. It should be noted that fault isolation in a connection-less connectionless (multi-point) environment is more challenging than a connection-oriented (point-to-point) environment. In case of Ethernet, fault isolation can be even more challenging since the a MAC address of a target node can age out in several minutes (e.g. typically in order of 5 minminutes) when the a fault results in isolating isolates the target nodeMAC address. As a result of this age-out, the occurrence of Consequently a network-isolating fault results in erasure of information leading to the location of needed for locating the fault!

The A Linktrace Message (LTM) uses a well-defined multicast MAC address. The Linktrace Message LTM gets initiated by a source MEP and traverses hop-by-hop and each MIP along the path intercepts the Linktrace Message this LTM and forwards it onto the next hop only after processing itit until it reaches the destination MEP. The processing includes looking at the target destination MEP's MAC address contained in the Linktrace MessageLTM. The originating MEP expects a response to its Linktrace Message. It should be noted that the source MEP sends a single request message LTM to the next hop along the trace path; however, it can receive many responses LTRs (Link Trace Response) from different MIPs along the trace path and the destination MEP as the result of the message LTM traversing hop by hop.

Given that an end-to-end Linktrace flow LTM is different from that of a user data flow (Linktrace Message goes through the control plane of LTMs undergo processing in bridge brain at each hop; whereas, user while data flow doesnitdoes not get processed in bridge brain), there can exist a rare situation in which the fault can is not be detected by the Linktrace flowLT mechanism. Given that the Linktrace flow can identify all the points MIPs and destination MEP along the traced path (based on responses LTRs received at the source maintenance point) one can run MEP, multiple Loopback messages LBMs between the source maintenance point MEP and different intermediate points (MIPs and the peer maintenance point) destination MEP to further isolate the data-plane fault/corruption faults in such rare situations.

As mentioned previously, the age-out of MAC address entries addresses can lead to erasure of information at intermediate nodesMIPs, which where this information is used for the Linktrace mechanism. Possible ways to address this behavior include:

- Launching Carrying out Linktrace mechanism following fault detectiondetection and/isolation or verification such that it gets exercised within the window of age-out.
- Maintaining information about the destination maintenance point MEP at the intermediate points MIPs along the path (Note: this can be facilitated by the CC messages.CCMs)
- Maintaining visibility of path at the source maintenance points MEPs through periodic Linktrace Messages (Note: this periodicity should be larger than the CC periodicity)