Configuring Global Service Entities with CLI

This section provides information to create subscriber (customer) accounts and configure Service Distribution Points (SDPs) using the command line interface.

Topics include:

- Service Model Entities on page 99
- Configuring Customers on page 103
  - → Configuring Multi-Service-Sites on page 105
- Configuring an SDP on page 106
- ETH-CFM Features on page 162
- Service Management Tasks on page 166

Service Model Entities

The Alcatel-Lucent service model uses logical entities to construct a service. The service model contains four main entities to configure a service.

- Subscribers on page 103
- SDPs on page 106
- Services:
  - → ATM VLL (Apipe) Services on page 203
  - → Circuit Emulation Services (Cpipe) on page 209
  - → Ethernet Pipe (Epipe) Services on page 220
  - → Frame Relay VLL (Fpipe) Services on page 228
  - → IP Interworking VLL (Ipipe) Services on page 233
  - → VPLS on page 733
  - → IES on page 1253
  - → VPRN on page 1557

- Service Access Points (SAPs)
  - → Ethernet Pipe (Epipe) Services on page 220
  - → Apipe SAP on page 353
  - → Fpipe SAP on page 376
  - → VPLS SAP on page 750
  - → IES SAP on page 1260
  - → VPRN Interface SAP on page 1579

Basic Configuration

The most basic service configuration must have the following:

- A customer ID
- A service type
- A service ID

  An optional service name can also be configured in addition to the service ID. Service names are optional. All services are required to assign a service ID to initially create a service. However, either the service ID or the service name can be used to identify and reference a given service once it is initially created.
- A SAP identifying a port and encapsulation value
- An interface (where required) identifying an IP address, IP subnet, and broadcast address
- For distributed services: an associated SDP

The following example provides an Epipe service configuration displaying the SDP and Epipe service entities. SDP ID 2 was created with the far-end node 10.10.10.104. Epipe ID 6000 was created for customer ID 6 which uses the SDP ID 2.

```
A:ALA-B>config>service# info detail
#------------------------------------------
...
        sdp 2 gre create
            description "GRE-10.10.10.104"
            far-end 10.10.10.104
            signaling tldp
            no vlan-vc-etype
            keep-alive
            path-mtu 4462
            keep-alive
                shutdown
                hello-time 10
                hold-down-time 10
                max-drop-count 3
                timeout 5
                no message-length
            exit
            no shutdown
        exit
...
       epipe 6000 customer 6 vpn 6000 create
           service-name “customer-ABC-NW” (R8.0)
           service-mtu 1514
           sap 1/1/2:0 create
               no multi-service-site
               ingress
                   no scheduler-policy
                   qos 1
               exit
               egress
```

```
                   no scheduler-policy
                   qos 1
               exit
               no collect-stats
               no accounting-policy
               no shutdown
           exit
           spoke-sdp 2:6111 create
               ingress
                   no vc-label
                   no filter
               exit
               egress
                   no vc-label
                   no filter
               exit
               no shutdown
           exit
           no shutdown
       exit
...
#------------------------------------------
A:ALA-B>config>service#
```

Common Configuration Tasks

This section provides a brief overview of the tasks that must be performed to configure a customer account and an SDP.

Configuring Customers

The most basic customer account *must* have a customer ID. Optional parameters include:

- Description
- Contact name
- Telephone number
- Multi-service site

Customer Information

Use the following CLI syntax to create and input customer information:

CLI Syntax:

```
config>service# customer customer-id create
  contact contact-information
  description description-string
  multi-service-site customer-site-name [create]
     assignment {port port-id | card slot}
     description description-string
     egress
        scheduler-override
           scheduler scheduler-name
              rate pir-rate [cir cir-rate]
        scheduler-policy scheduler-policy-name
     ingress
        scheduler-override
           scheduler scheduler-name
              rate pir-rate [cir cir-rate]
        scheduler-policy scheduler-policy-name
  phone phone-number
```

The following displays a basic customer account configuration.

```
A:ALA-12>config>service# info
-------------------------------------------
...
       customer 5 create
           description "Alcatel Customer"
           contact "Technical Support"
           phone "650 555-5100"
       exit
...
-------------------------------------------
A:A:ALA-12>config>service#
```

Configuring Multi-Service-Sites

Multi-service sites create a virtual scheduler hierarchy and making it available to queues on multiple Service Access Points (SAPs). The **ingress** and **egress scheduler-policy** commands on the SAP are mutually exclusive with the SAP **multi-service-site** command. The multi-service customer site association must be removed from the SAP before local scheduler polices may be applied.

After a multi-service site is created, it must be assigned to a chassis slot or port. Note that the 7750 SR-1 model multi-service site `assignment` configuration defaults to slot 1.

Use the following CLI syntax to configure customer multi-service sites.

CLI Syntax:

```
config>service# customer customer-id
 multi-service-site customer-site-name
    assignment {port port-id | card slot}
    description description-string
    egress
       agg-rate-limit agg-rate
       scheduler-policy scheduler-policy-name
    ingress
       scheduler-policy scheduler-policy-name
    tod-suite tod-suite-name
```

The following displays a customer’s multi-service-site configuration.

```
A:ALA-12>config>service# info
-------------------------------------------
..
       customer 5 create
           multi-service-site "EastCoast" create
               assignment card 4
               ingress
                   scheduler-policy "alpha1"
               exit
           exit
           multi-service-site "WestCoast" create
               assignment card 3
               egress
                   scheduler-policy "SLA1"
               exit
           exit
           description "Alcatel Customer"
           contact "Technical Support"
           phone "650 555-5100"
       exit
...
-------------------------------------------
A:ALA-12>config>service#
```

Configuring an SDP

The most basic SDP must have the following:

- A locally unique SDP identification (ID) number.
- The system IP address of the originating and far-end routers.
- An SDP encapsulation type, either GRE or MPLS.

SDP Configuration Tasks

This section provides a brief overview of the tasks that must be performed to configure SDPs and provides the CLI commands.

Consider the following SDP characteristics:

- SDPs can be created as either GRE or MPLS.
- Each distributed service must have an SDP defined for every remote router to provide VLL, VPLS, and VPRN services.
- A distributed service must be bound to an SDP. By default, no SDP is associated with a service. Once an SDP is created, services can be associated to that SDP.
- An SDP is not specific or exclusive to any one service or any type of service. An SDP can have more than one service bound to it.
- The SDP IP address must be a 7750 SR-Series system IP address.
- In order to configure an MPLS SDP, LSPs must be configured first and then the LSP-to-SDP association must be explicitly created.
- In the SDP configuration, automatic ingress and egress labeling (targeted LDP) is enabled by default. Ingress and egress VC labels are signaled over a TLDP connection between two 7750 SR-Series routers.

  Note that if signaling is disabled for an SDP, then services using that SDP must configure ingress and egress vc-labels manually.

To configure a basic SDP, perform the following steps:

1. Specify an originating node.
2. Create an SDP ID.
3. Specify an encapsulation type.
4. Specify a far-end node.

Configuring an SDP

Use the following CLI syntax to create an SDP and select an encapsulation type. If you do not specify GRE or MPLS, the default encapsulation type is GRE.

NOTE: When you specify the far-end ip address, you are creating the tunnel. In essence, you are creating the path from Point A to Point B. When you configure a distributed service, you must identify an SDP ID. Use the show service sdp command to display the qualifying SDPs.

When specifying MPLS SDP parameters, you must specify an LSP or enable LDP. There cannot be two methods of transport in a single SDP except if the mixed-lsp option is selected. If an LSP name is specified, then RSVP is used for dynamic signaling within the LSP.

LSPs are configured in the **config>router>mpls** context. See the OS MPLS Guide for configuration and command information.

Use the following CLI syntax to create a GRE SDP or an MPLS SDP:

CLI Syntax:

```
config>service>sdp sdp-id [gre | mpls] create
  adv-mtu-override
  description description-string
  far-end ip-address
  keep-alive
     hello-time seconds
     hold-down-time seconds
     max-drop-count count
     message-length octets
     timeout timeout
     no shutdown
           ldp                 (only for MPLS SDPs)
           lsp lsp-name [lsp-name](only for MPLS SDPs)
  path-mtu octets
  signaling {off | tldp}
  no shutdown
```

The following displays a GRE SDP, an LSP-signalled MPLS SDP, and an LDP-signalled MPLS SDP configuration.

```
A:ALA-12>config>service# info
-------------------------------------------
...
        sdp 2  create
            description "GRE-10.10.10.104"
            far-end 10.10.10.104
            keep-alive
                shutdown
            exit
            no shutdown
        exit
        sdp 8 mpls create
            description "MPLS-10.10.10.104"
            far-end 10.10.10.104
            lsp "to-104"
            keep-alive
                shutdown
            exit
            no shutdown
        exit
        sdp 104 mpls create
            description "MPLS-10.10.10.94"
            far-end 10.10.10.94
            ldp
            keep-alive
                shutdown
            exit
            no shutdown
        exit
...
-----------------------------------------
A:ALA-12>config>service#
```

Configuring a Mixed-LSP SDP

Use the following command to configure an SDP with mixed-LSP mode of operation:

config>service>sdp mpls>mixed-lsp-mode

The primary is backed up by the secondary. Two combinations are possible: primary of RSVP is backed up by LDP and primary of LDP is backed up by 3107 BGP.

The **no** form of this command disables the mixed-LSP mode of operation. The user first has to remove one of the LSP types from the SDP configuration or the command will fail.

The user can also configure how long the service manager must wait before it must revert the SDP to a higher priority LSP type when one becomes available by using the following command:

config>service>sdp mpls>mixed-lsp-mode>sdp-revert-time *seconds*

A special value of the timer dictates that the SDP must never revert to another higher priority LSP type unless the currently active LSP type is down:

config>service>sdp mpls>mixed-lsp-mode>sdp-revert-time infinite

The BGP LSP type is allowed. The **bgp-tunnel** command can be configured under the SDP with the **lsp** or **ldp** commands.

Mixed-LSP Mode of Operation

The mixed LSP SDP allows for a maximum of two LSP types to be configured within an SDP. A primary LSP type and a backup LSP type. An RSVP primary LSP type can be backed up by an LDP LSP type.

An LDP LSP can be configured as a primary LSP type which can then be backed up by a BGP LSP type.

At any given time, the service manager programs only one type of LSP in the linecard that will activate it to forward service packets according to the following priority order:

1. RSVP LSP type. Up to 16 RSVP LSPs can be entered by the user and programmed by the service manager in ingress linecard to load balance service packets. This is the highest priority LSP type.
2. LDP LSP type. One LDP FEC programmed by service manager but ingress linecard can use up to 16 LDP ECMP paths for the FEC to load balance service packets when ECMP is enabled on the node.
3. BGP LSP type. One RFC 3107-labeled BGP prefix programmed by the service manager. The ingress linecard can use more than one next-hop for the prefix.

In the case of the RSVP/LDP SDP, the service manager will program the NHLFE(s) for the active LSP type preferring the RSVP LSP type over the LDP LSP type. If no RSVP LSP is configured or all configured RSVP LSPs go down, the service manager will re-program the linecard with the LDP LSP if available. If not, the SDP goes operationally down.

When a higher priority type LSP becomes available, the service manager reverts back to this LSP at the expiry of the **sdp-revert-time** timer or the failure of the currently active LSP, whichever comes first. The service manager then re-programs the linecard accordingly. If the **infinite** value is configured, then the SDP reverts to the highest priority type LSP only if the currently active LSP failed.

Note however, that LDP uses a tunnel down damp timer which is set to three seconds by default. When the LDP LSP fails, the SDP will revert to the RSVP LSP type after the expiry of this timer. For an immediate switchover this timer must be set to zero. Use the **configure>router>ldp>tunnel-down-damp-time** command.

If the value of the **sdp-revert-time** timer is changed, it will take effect only at the next use of the timer. Any timer which is outstanding at the time of the change will be restarted with the new value.

If class based forwarding is enabled for this SDP, the forwarding of the packets over the RSVP LSPs will be based on the FC of the packet as in current implementation. When the SDP activates the LDP LSP type, then packets are forwarded over the LDP ECMP paths using the regular hash routine.

In the case of the LDP/BGP SDP, the service manager will prefer the LDP LSP type over the BGP LSP type. The service manager will re-program the linecard with the BGP LSP if available otherwise it brings down the SDP operationally.

Also note the following difference in behavior of the LDP/BGP SDP compared to that of an RSVP/LDP SDP. For a given /32 prefix, only a single route will exist in the routing table: the IGP route or the BGP route. Thus, either the LDP FEC or the BGP label route is active at any given time. The impact of this is that the tunnel table needs to be re-programmed each time a route is deactivated and the other is activated. Furthermore, the SDP revert-time cannot be used since there is no situation where both LSP types are active for the same /32 prefix.

Ethernet Connectivity Fault Management (ETH-CFM)

Ethernet Connectivity Fault Management (ETH-CFM) is defined in two similar standards: IEEE 802.1ag and ITU-T Y.1731. They both specify protocols, procedures, and managed objects to support transport fault management, including discovery and verification of the path, detection and isolation of a connectivity fault for each Ethernet service instance. CFM functionalities are supported on SR and ESS platforms.

The configuration is split into multiple areas. There is the base ETH-CFM configuration which defines the different Management constructs and administrative elements. This is performed in the ETH-CFM context. The individual management points are configure within the specific service contexts in which they are applied.

The OS Services Guide will provide the basic service applicable material to build the service specific management points, MEPs and MIPs.

The different service types support a subset of the features from the complete ETH-CFM suite.

ETH-CC used for continuity is available to all MEPs configured within a service and all facility MEPs.

The troubleshooting tools ETH-LBM/LBR, LTM/LTR ETH-TST defined by the IEEE 802.1ag specification and the ITU-T Y.1731 recommendation are applicable to all MEPs (MIPs where appropriate).

The advanced notification function AIS defined by the ITU-T Y.1731 is supported on Epipe services and may be terminated by a MEP on a Layer 3 service interface.

The advanced performance functions, 1DM, DMM/DMR and SLM/SLR are supported on all service MEPs, not on facility MEPs.

For a description of the individual features and functions that are supported refer to the applicable OAM Diagnostics Guide.

Acronym	Callout
1DM	One way Delay Measurement (Y.1731)
AIS	Alarm Indication Signal
CCM	Continuity check message
CFM	Connectivity fault management
DMM	Delay Measurement Message (Y.1731)
DMR	Delay Measurement Reply (Y.1731)

Acronym	Callout (Continued)
LBM	Loopback message
LBR	Loopback reply
LTM	Linktrace message
LTR	Linktrace reply
ME	Maintenance entity
MA	Maintenance association
MA-ID	Maintenance association identifier
MD	Maintenance domain
MEP	Maintenance association end point
MEP-ID	Maintenance association end point identifier
MHF	MIP half function
MIP	Maintenance domain intermediate point
OpCode	Operational Code
RDI	Remote Defect Indication
TST	Ethernet Test (Y.1731)
SLM	Synthetic Loss Message (Y.1731)
SLR	Synthetic Loss Reply (Y.1731)

ETH-CFM capabilities may be deployed in many different Ethernet service architectures. The Ethernet based SAPs and SDP bindings provide the endpoint on which the management points may be created. The basic functions can be used in different services, VPLS, Ipipe, Epipe and even in IES, VPRN and the base router instance interfaces. Of course, Layer 3 services are boundaries for Layer 2 ETH-OAM functions. The ETH-CFM functionality is also applicable to broadband access networks. Two models of broadband access are shown below to illustrate how ETH-CFM could be deployed in these cases. (Figure 30 and Figure 31).

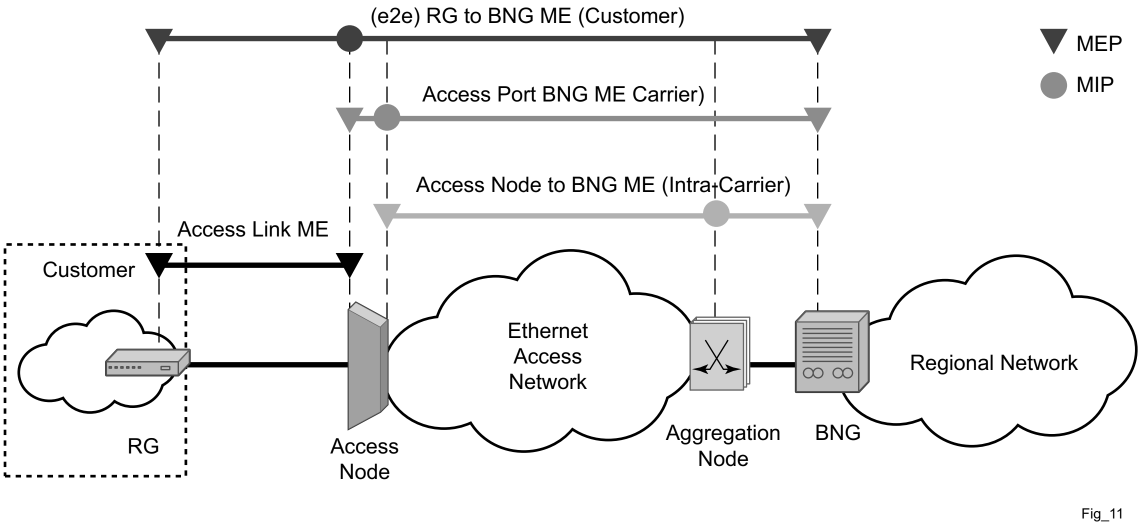

Figure 30: Ethernet OAM Model for Broadband Access - Residential

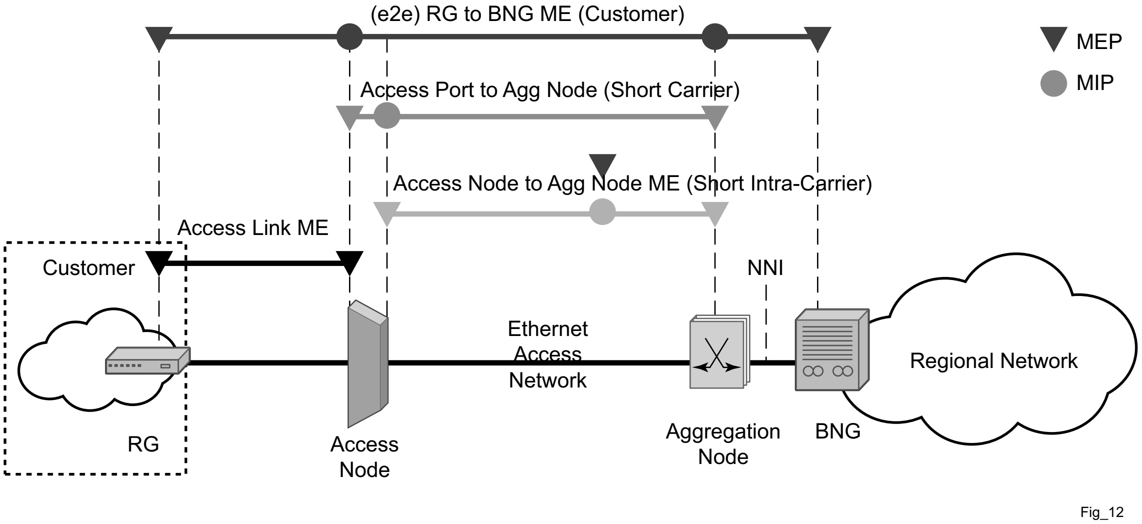
(e2e) RG to BNG ME (Customer)
▼ MEP
● MIP
Access Port to Agg Node (Short Carrier)
Access Node to Agg Node ME (Short Intra-Carrier)
Access Link ME
Customer
NNI
Ethernet Access Network
Regional Network
RG
Access Node
Aggregation Node
BNG
Fig_12

Figure 31: Ethernet OAM Model for Broadband Access - Wholesale

As shown in Figure 30 and Figure 31, the following functions are supported:

- CFM can be enabled or disabled on a SAP or SDP bindings basis.
- The eight ETH-CFM levels are suggested to be broken up numerically between customer 7-5, service provider 4-3 and Operator 2-1. Level 0 is meant to monitor direct connections without any MIPs and should be reserved for port-based facility MEPs. These can be configured, deleted or modified.
- Up and/or down MEP with an MEP-ID on a SAP and SDP binding for each MD level can be configured, modified, or deleted. Each MEP is uniquely identified by the MA-ID, MEP-ID tuple.
  - → MEP creation on a SAP is allowed only for Ethernet ports (with null, q-tags, qinq encapsulations).
- MIP creation on a SAP and SDP binding for each MD level can be enabled and disabled. MIP creation is automatic or manual when it is enabled. When MIP creation is disabled for an MD level, the existing MIP is removed.
  - → MIP creation is not supported on mesh SDP bindings.

Facility MEPs

Facility MEPs have been introduced to improve scalability, reduce operational overhead, and provide fate sharing without requiring service MEPs. This allows for fault notification for Epipe services that share a common transport. Facility MEPs recognize failure based solely on ETH-CFM detection mechanisms.

There are a total of four facility MEPs, as described below:

- Port (physical) — Detects port failure where LoS may be hidden by some intervening network
- LAG (logical) — Validates the connectivity of the LAG entity
- Tunnel (logical) — Enables fate sharing of a MEP configured on a QinQ encapsulated access LAG and outer VLAN-ID.
- Router IP Interface (logical) — Validates the Layer 2 connectivity between IP endpoints (troubleshooting only – no CCM functions)

In general, a Facility MEP detects failure conditions using ETH-CFM at the Ethernet Transport layer. The detection is based solely on the MEP entering a fault state as a result of ETH-CC. Conditions outside the scope of ETH-CFM do not directly influence the state of the MEP. However, these outside influences have indirect influence. For example, upon a failure of a port, CCM messages cannot reach the destination. This condition causes the MEP to enter a fault state after the `3.5*interval` expires, with the only exception being the acceptance of AIS on a Tunnel MEP. AIS received on all other facilities MEPs are discarded silently when normal level matching targets the local facility MEP.

Facility MEPs are supported as part of a down MEP only. Facility MEPs validate the point to point Ethernet transport between two end points. Facility MEPs do not validate switching functions that are not part of the point to point Ethernet transport. Instead, service MEPs validate switching functions that are not part of the point to point Ethernet transport.

A facility MEP allows for the scaling improvements using fate sharing and leveraging OAM mapping. The OAM mapping functions are part of the fault propagation functions and allow ETH-CFM to move from alarms only to network actions. Service based MEPs are not required to generate AIS in reaction to a facility MEP fault. OAM mapping and fault generation, either the R8.0 function or the AIS function as part of a facility MEP) are only available on Epipe services.There is no equivalent AIS generation as part of the facility fault for VPLS, IES, and VPRN. There is no service MEP required to have the SAP transition in the VPLS, IES, and VPRN service context. Normal SAP transition functions does not occur when these services are configured to accept the tunnel fault, or in reaction to a facility fault, where the underlying port or LAG transitions the SAP.

Note: Do not exceed the platform-specific scaling limits. Since a single facility fault may trigger the generation of many service level faults, ensure the specific ETH-CFM processing power of the network element and any configured rate controlling features for the service are not exceed. Exceeding the network element scaling properties may lead to OAM packet loss during processing and result in undesirable behavior.

The implementation of facility MEPs must adhere to all platform-specific specifications. For example, sub-second enabled CCM MEPs are supported on port based MEPs. However, any platform restrictions preventing the sub-second enabled MEPs override this capability and require the operator to configure CCM intervals that are supported for that specific platform.

Facility MEPs are created in the same manner as service MEPs, both related to the ETH-CFM domain and association. However, the association used to build the facility MEP does not include a bridge-identifier. The CLI ensures that a bridge id is not configured when the association is applied to a facility MEP.

Service MEPs and Facility MEPs may communicate with each other, as long as all the matching criteria are met. Since facility MEPs use the standard ETH-CFM packets, there is nothing contained in the packet that would identify an ETH-CFM packet as a facility MEP or Service MEP.

Facility MEPs are not supported on ports that are configured with Eth-Tunnels (G.8031), and only facility MEPs of 1 second and above are supported on the ports that are involved in an Eth-Ring (G.8032).

Common Actionable Failures

It is important to note that AIS operates independently from the **low-priority-defect** setting. The **low-priority-defect** setting configuration parameter affects only the ETH-CFM fault propagation and alarming outside the scope of AIS. Any fault in the MEP state machine generates AIS when it is configured. Table 3 illustrates the ETH-CC defect condition groups, configured low-priority-defect setting, priority and defect as it applies to fault propagation.

Table 3: Defect Conditions and Priority Settings

Defect	Low Priority Defect	Description	Causes	Priority
DefNone	n/a	No faults in the association	Normal operations	n/a
DefRDICCM	allDef	Remote Defect Indication	Feedback mechanism to inform unidirectional faults exist. It provides the feedback loop to the node with the unidirectional failure conditions	1
DefMACStatus (default)	macRemErrXcon	MAC Layer	Remote MEP is indicating a remote port or interface not operational.	2
DefRemoteCCM	remErrXon	No communication from remote peer.	MEP is not receiving CCM from a configured peer. The timeout of CCM occurs at 3.5x the local CC interval. As per the specification, this value is not configurable.	3
DefErrorCCM	errXcon	Remote and local configures do not match required parameters.	Caused by different interval timer, domain level issues (lower value arriving at a MEP configured with a higher value), MEP receiving CCM with its MEPID	4
DefXconn	Xcon	Cross Connected Service	The service is receiving CCM packets from a different association. This could indicate that two services have merged or there is a configuration error on one of the SAP or bindings of the service, incorrect association identification.	5

A facility MEP may trigger two distinct actions as a result of fault. Epipe services generate AIS that have been configured to do so as a result of a failure. The level of the AIS is derived from the facility MEP. Multiple **client-meg-levels** can be configured under the facility MEP to allow for operational efficiency in the event a change is required. However, only the lowest AIS level is generated for all the linked and applicable services. VPLS, IES and VRPN SAPs transition the SAP state that are configured to react to the facility MEP state. In addition, Epipe services may also take advantages of OAM and mapping functions.

Before implementing facility MEPs, it is important to understand the behavior of AIS and Fault propagation. Alcatel-Lucent advises that you strongly considered the following recommendations listed below before enabling or altering the configuration of any facility MEP. These steps must be tested on each individual network prior to building a maintenance operational procedure (MOP).

- Do not configure AIS on the facility MEP until the ETH-CCM has been verified. For instance, when a local MEP is configured with AIS prior to the completion of the remote MEP, the AIS is immediately generated when the MEP enters a fault state for all services linked to that facility MEP.
- Disable the **client-meg-level** configuration parameter when changes are being made to existing functional facility MEPs for AIS. Doing this stops the transmit function but maintains the ability to receive and understand AIS conditions from the network.
- Set the **low-priority-defect** parameter to **noXconn** in order to prevent the MEP from entering a defect state, triggering SAP transitions and OAM mapping reactions.

It is important to consider and select what types of fault conditions causes the MEP to enter a faulty state when using fault propagation functions.

The **ccm-hold-timers** supported on port-based MEPs configured with a sub-second interval. The **ccm-hold-timers** prevents the MEP from entering a failed state for 3.5 times the CCM interval plus the additional hold timer.

General Detection, Processing and Reaction

All Facility MEPs that support CCM functions must only have one remote MEP peer. Facilities MEPs validate point-to-point logical or physical Ethernet transports. Configure service MEPs if multipoint-service validation is required.

There are three distinct functions for a Facility MEP:

- General Detection: Determines that a fault has occurred. In this case, the MEP performs its normal functions such as: recognizing the fault condition, maintaining the local errors and reporting based on low-priority-setting, and taking no further action. This is the default.
- Fault Processing: By default, there is no action taken as a result of a MEP state machine transition beyond alarming. In order to take action which may include a SAP operational state change, generation of AIS, or fault propagation and mapping, the appropriate facility fault configuration parameter must be configured and enabled. The general reaction to a fault is described below. More details are including the section describing the functions of the individual facility MEPs.
    - → Port—Affects link operational status of the port. Facility failure changes the operational state to `Link Up`. This indicates that the port has been brought down as a result of OAM MEP Fault. This operational state has the equivalent function to port down condition.
    - → LAG—Affects link operational status of the LAG. Facility failure changes the operational state of the LAG to `DOWN`. This indicates that the LAG has be brought down as a result of OAM MEP Fault.
    - → Tunnel MEP—Enters faulty state and will further impact the operational state of the SAPs linked to the tunnel MEP state.
        - – Epipe SAP remains operationally up, SAP's flag set to **OamTunnelMEPFault**
        - – Ipipe SAP remains operationally up, SAP's flag set to **OamTunnelMEPFault**
        - – VPLS, IES and VPLS SAPs transition to operationally **down**, the SAP's flag is set to **OamTunnelMEPFault**

          SAP operational states and flags are affect only by the **tunnel-fault** configuration option.
    - → Router IP Interface— Affects operational status of the IP Interface.

- Propagation: Services appropriately linked to the Facility MEP take the following service specific actions:
  - → Epipe generates AIS or use Fault Propagation and OAM mappings.
  - → VPLS does not propagate fault using AIS unless service-based MEPs are configured and contain MEP-specific AIS configuration. SAP transitions will occur when the facility MEP failure is recognized by the service.
  - → IES and VPRN, as Layer 3 functions, act as boundaries for Layer 2 fault processing. No propagation functions occur beyond what is currently available as part of fault propagation, SAP down.
- AIS-enable configuration options: Epipe services support the `ais-enable` configuration option under the SAP hierarchy level. This structure, outside of the MEP context, creates a special link between the Epipe service SAP and the facility MEP. If a facility MEP enters a fault state, all Epipe service SAPs with this configuration generate lowest-level AIS at the level configured under the facility MEP. As with fault propagation, AIS generation is restricted to Epipe services only. The actions taken by the other services is described in more detail in the relevant facility MEP sections.

  NOTE: Facility MEPs do not support the generation of AIS to an explicitly configured endpoint. An explicitly configure endpoint abstracts multiple endpoints within its context, for example, pseudowire (PW) redundancy. Although the linkage of a facility MEP to an Epipe and AIS generation triggered as a result of the facility MEP failure can be configured AIS generation is not supported and will be unpredictable. When an explicit endpoint is configured service based MEPs are required when AIS generation is the desired behavior.

Port-Based MEP

There is an increase in services that share the same facilities, and that service-based ETH-CFM, although very granular, comes at an operational and scalability cost. Configuring a MEP on a physical port allows ETH-CFM to detect Ethernet transport failures, raise a facility alarm, and perform local fault processing. A facility event is coordinated to the services or functions using the affected port.

Port-based facility MEPs are able to run all supported on-demand and SAA, 802.1ag and ITU-T Y.1731 ETH-CFM functions.

The port-based MEP is intended to validate physical connectivity to the peer MEP, provide on-demand and scheduled troubleshooting, and performance management functions.

Port facility MEPs are advantageous in cases where port-to-port connectivity issues are obscured., similar to the deployment use cases for *IEEE 802.3 Clause 57 – Operation, Administration and Maintenance* (formerly 802.3ah). *Clause 57* specification limits the transmit rate to 10pps, or a send rate of 100ms. In order to detect port failure conditions between two peers faster, a port-based facility MEP may be configured to utilize the supported sub-second CCM intervals. Also, 1 second and above timers are available for configuration for cases where aggressive timers are not necessary. Note that all platform-specific requirements must be met for the desired interval. Since both ETH-CFM and IEEE 802.3 Clause 57 attempt to control the port state in event of protocol failure, these two functions are mutually exclusive and can not be configured on the same port.

Port-level ETH-CFM PDUs are sent untagged because they are not specific to any service or VLAN. The ETH-CFM packets generated from a port-based facility MEP must use an ETH-CFM level of 0 or 1. Any ETH-CFM PDU that arrives untagged on a port matching the level for the port based facility MEP will be terminated and processed by the port based MEP.

Do not use MEPs configured with level 0 to validate logical transport or services. Consideration should be given to blocking all non-customer (5-7) levels at the entry point of the network.

It is not expected that faults from other parts of the network will be propagate and terminated on a port-based facility MEP. This type of facility MEP provides a one-to-one validation with a single remote MEP across on a physical port, allowing locally detected faults to be propagated to the endpoints of the network.

A physical port may only have a single port based facility MEP. Since the purpose of the MEP is to control the port state, more than one is not required per port. The MEP must be configured with the **direction-down** option.

Port based MEPs are supported in both the IEEE 802.1ag and ITU-T Y.1731 contexts. Therefore, the Y.1731 context must be configured in order to run functions beyond those that are described as part of the IEEE 802.1ag specification.

When a port enters the `link up` operational state due to ETH-CFM, the MEP continues to transmit and received in order to properly clear the condition. However, when the port fails for reasons that are not specific to ETH-CFM, it stops transmit and receive functions until the condition is cleared. This is different than the behavior of a service MEP, because facility MEPs only supports `Down` MEPs, while some service-based MEPs support `UP` and `Down` MEPs. In the case of `UP` MEPs, a single port failure may not prevent all the CCMs from egressing the node. So the operational method for service-based MEPs remains the same: continuing to increase the counter for CCM transmit in the event of port failure, regardless of the reason. The transmit ETH-CCM counters do not apply to sub-second CCM-enabled MEPs.

There are two types of port in the context of port-based facility MEPs. The first type are ports that are not part of a LAG, referred to as non-member ports. The second type of ports are ports that are part of a LAG, referred to member ports, and have slightly different reactions to fault. MEPs configured directly on either type of port will act the same. However, a MEP configured on a non-member port and a MEP configured on a member port handle fault propagation differently.

When a port-based facility MEP causes the port to enter the operational state `Link Up`, normal processing occurs for all higher level functions. If the port is a member port, unless the entire LAG enters a non-operational state, the SAP configured on the LAG remains operational. A facility MEP on a member port has no direct influence on the SAP. The purpose of a facility MEP on a member port is to provide feedback to the LAG. The LAG performs the normal computations in response to a port down condition. A facility MEP configured on a non-member port does have direct control over the SAPs configured on the port. Therefore, when a port fails, all the SAPs transitions to the operation state `down`. When this occurs, fault may be propagated using AIS for those Epipe services that are AIS-enabled under the SAO. For the services that have MEPs configured on the SAP or the binding, fault propagation occurs. For VPLS, IES and VPRN services, normal reaction to a SAP entering a `down` state occurs.

When a LAG is administratively shutdown, the member ports are shutdown automatically. As a result, packet reception is interrupted, causing ETH-CFM functions running on physical member ports to lose connectivity. Therefore, the CFM functions on member ports are somewhat tied to the LAG admin status in this case.

It is important to note that LAG convergence time is not affected by a facility MEP on a member port once the port has entered the `link up` operational state. The ETH-CFM failure of a port-based MEP acts as the trigger to transition the port.

Figure 32 on page 123 provides an example of how an ETH-CFM failure reacts with the various services that share that port. The green Epipe service generates AIS as a result of the port failure using the **`client-meg-level`** command configured on the port facility MEP. The multipoint service takes location configured action when the SAP transitions to the `down` operational state. The blue Epipe service is not affected by the port `link up` state as a result of ETH-CFM fault.

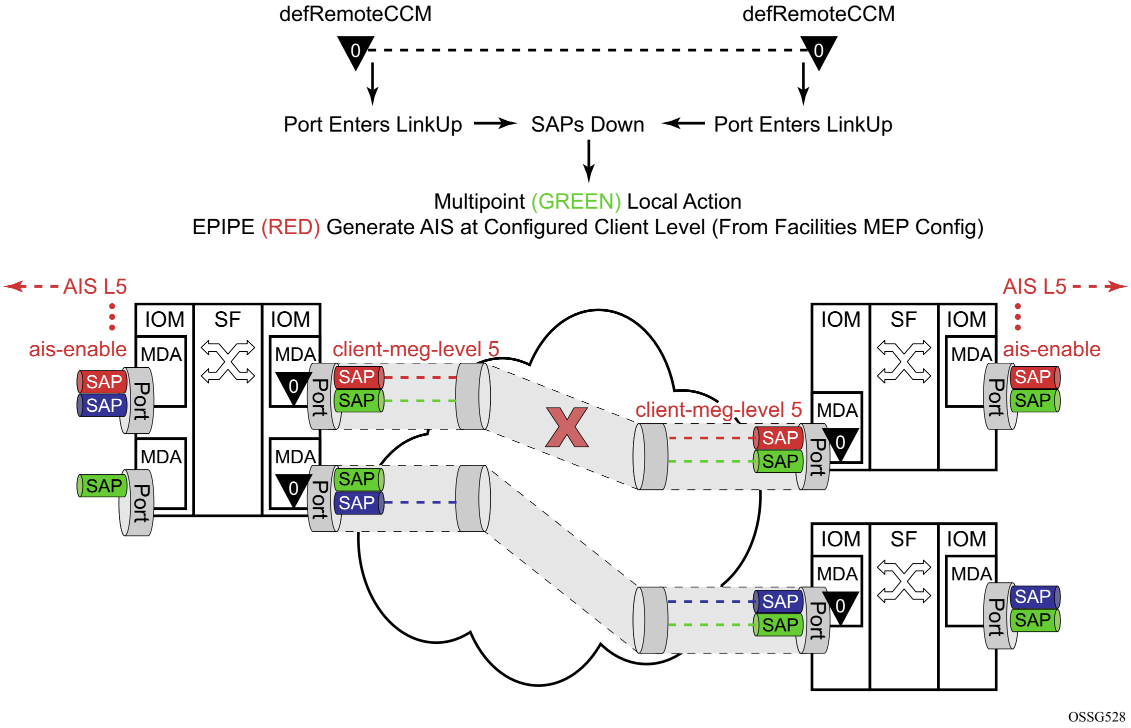

Figure 32: Fault Handling Non-Member Port

A debounce function has been implemented to prevent notifying every port state change if a port bounces multiple times within a window.   Up to four notifications will be accepted in a three second window. If the third port state is a down state change the fourth will be ignored. If the fourth port state change is a down state change it will be processed. After that no further state changes will be accepted for the duration of the three second timer. This helps ensure that the port is not artificially held in the UP state when it is not operation. Following the processing of that last port state change, the third or fourth, the latest state change will be held and processed at the expiration of the three second hold timer.

Port based facility MEPs are not allowed on a port that is configured with G.8031 Ethernet Tunnels.

Example: Port-Based MEP Configuration

The following illustration, Figure 33, provides an example of how port-based MEPs and defect conditions translate into service awareness without service-based MEPs. From the two nodes perspective, they are aware they are directly connected at the port. The two nodes are unaware of any of the cross connections that allow this to occur.

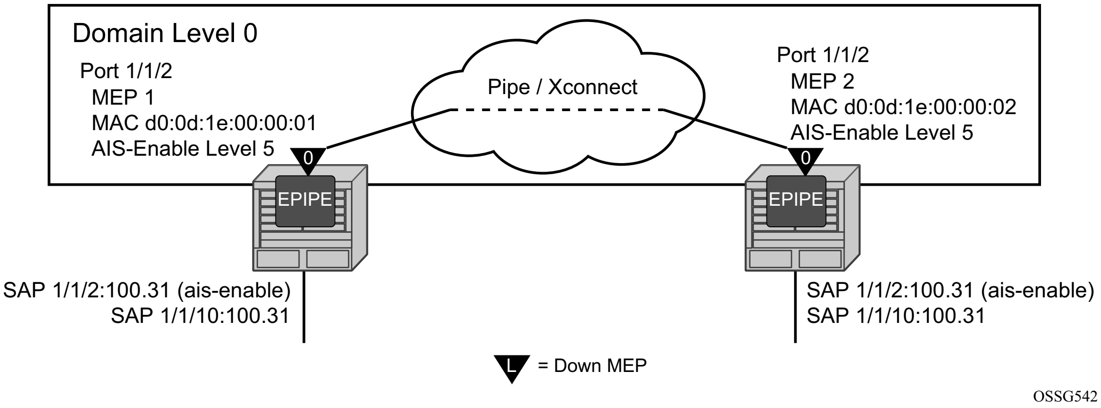

Figure 33: Port-Based MEP Example

Configure port-based MEPs with the **facility-fault** option and **ais-enable client-meg-level** command. When the MEP enters any defect state, an AIS is generated to any Epipe service that has the `ais-enable` configured under the `sap>eth-cfm` hierarchy.

```
NODE1
config>eth-cfm# info
----------------------------------------------
        domain 10 format none level 0
            association 1 format icc-based name "FacilityPort0"
                ccm-interval 1
                remote-mepid 2
            exit
        exit
----------------------------------------------

config>port# info
----------------------------------------------
        ethernet
            mode access
            encap-type qinq
            eth-cfm
                mep 1 domain 10 association 1
                    ais-enable
                        client-meg-level 5
                    exit
                    facility-fault
                  ccm-enable
                    mac-address d0:0d:1e:00:00:01
                    no shutdown
                exit
            exit
        exit
        no shutdown
----------------------------------------------

config>service>epipe# info
----------------------------------------------
            sap 1/1/2:100.31 create
```

```
                eth-cfm
                    ais-enable
                exit
            exit
            sap 1/1/10:100.31 create
            exit
            no shutdown
----------------------------------------------

NODE2
config>eth-cfm# info
----------------------------------------------
        domain 10 format none level 0
            association 1 format icc-based name "FacilityPort0"
                ccm-interval 1
                remote-mepid 1
            exit
        exit
----------------------------------------------

config>port# info
----------------------------------------------
        ethernet
            mode access
            encap-type qinq
            eth-cfm
                mep 2 domain 10 association 1
                    ais-enable
                        client-meg-level 5
                    exit
                    facility-fault
                    ccm-enable
                    mac-address d0:0d:1e:00:00:02
                    no shutdown
                exit
            exit
        exit
        no shutdown
----------------------------------------------

config>service>epipe# info
----------------------------------------------
            sap 1/1/2:100.31 create
                eth-cfm
                    ais-enable
                exit
            exit
            sap 1/1/10:100.31 create
            exit
            no shutdown
----------------------------------------------
```

There are two different levels of fault to consider: Port State / Operational State driven by the low-priority-defect setting and the generation of AIS driven by any defect state for the MEP, regardless of low-priority-defect.

If the low-priority-defect is left at the default `macRemErrXcon` setting, then port state may not match on both nodes. If an unidirectional failure is introduced for port-based MEPs, then RDI is received on one of the nodes and the other node would report and react to RemoteCCM (timeout). The RDI defect is below the default `low-priority-defect` in priority, and the port would remain operationally UP and the port state would remain UP. The MEP that has timed out the peer MEP takes port level action because this defect is higher in priority than the default low-priority-defect. The port state is recorded as `Link Up` and the Port is operationally down with a `Reason Down : ethCfmFault`. To avoid this inconsistency, set the **low-priority-defect** setting to detection unidirectional failures using the `allDef` option.

The following show commands reveal the condition mentioned above within the network. Node 1 is receiving RDI and Node 2 has timed out its peer MEP.

```
NODE1
#show port
===============================================================================
Ports on Slot 1
===============================================================================
Port        Admin Link Port    Cfg  Oper LAG/ Port Port Port   C/QS/S/XFP/
Id          State      State   MTU  MTU  Bndl Mode Encp Type   MDIMDX
-------------------------------------------------------------------------------
…snip..
1/1/2       Up    Yes  Up      1522 1522    - accs qinq xcme
…snip..


#show port 1/1/2
===============================================================================
Ethernet Interface
===============================================================================
Description        : 10/100/Gig Ethernet SFP
Interface          : 1/1/2                      Oper Speed       : 1 Gbps
Link-level         : Ethernet                   Config Speed     : 1 Gbps
Admin State        : up                         Oper Duplex      : full
Oper State         : up                         Config Duplex    : full
Physical Link      : Yes                        MTU              : 1522
…snip…

#show eth-cfm mep 1 domain 10 association 1
===============================================================================
Eth-Cfm MEP Configuration Information
===============================================================================
Md-index           : 10                       Direction         : Down
Ma-index           : 1                        Admin             : Enabled
MepId              : 1                        CCM-Enable        : Disabled
Port               : 1/1/2                    VLAN              : 0
Description        : (Not Specified)
FngState           : fngReset                 ControlMep        : False
LowestDefectPri    : macRemErrXcon            HighestDefect     : none
Defect Flags       : bDefRDICCM
Mac Address        : d0:0d:1e:00:00:01        ControlMep        : False
CcmLtmPriority     : 7
CcmTx              : 1481                     CcmSequenceErr    : 0
Fault Propagation  : disabled                 FacilityFault     : Notify
MA-CcmInterval     : 1                        MA-CcmHoldTime    : 0ms
Eth-1Dm Threshold  : 3(sec)                   MD-Level          : 0
```

```
Eth-Ais:           : Enabled                  Eth-Ais Rx Ais:   : No
Eth-Ais Tx Priorit*: 7                        Eth-Ais Rx Interv*: 1
Eth-Ais Tx Interva*: 1                        Eth-Ais Tx Counte*: 3019
Eth-Ais Tx Levels  : 5
Eth-Tst:           : Disabled
…snip…

# show service sap-using eth-cfm facility
===============================================================================
Service ETH-CFM Facility Information
===============================================================================
SapId             SvcId                          SAP AIS  SAP Tunnel SVC Tunnel
                                                          Fault      Fault
-------------------------------------------------------------------------------
1/1/2:100.31      100                            Enabled  Accept     Ignore
-------------------------------------------------------------------------------
No. of Facility SAPs: 1
===============================================================================

NODE2
# show port
===============================================================================
Ports on Slot 1
===============================================================================
Port        Admin Link Port    Cfg  Oper LAG/ Port Port Port   C/QS/S/XFP/
Id          State      State   MTU  MTU  Bndl Mode Encp Type   MDIMDX
-------------------------------------------------------------------------------
…snip..
1/1/2       Up    Yes  Link Up 1522 1522    - accs qinq xcme
…snip..

# show port 1/1/2
===============================================================================
Ethernet Interface
===============================================================================
Description        : 10/100/Gig Ethernet SFP
Interface          : 1/1/2                      Oper Speed       : N/A
Link-level         : Ethernet                   Config Speed     : 1 Gbps
Admin State        : up                         Oper Duplex      : N/A
Oper State         : down                       Config Duplex    : full
Reason Down        : ethCfmFault
Physical Link      : Yes                        MTU              : 1522
…snip…

# show eth-cfm mep 2 domain 10 association 1
===============================================================================
Eth-Cfm MEP Configuration Information
===============================================================================
Md-index           : 10                       Direction         : Down
Ma-index           : 1                        Admin             : Enabled
MepId              : 2                        CCM-Enable        : Enabled
Port               : 1/1/2                    VLAN              : 0
Description        : (Not Specified)
FngState           : fngDefectReported        ControlMep        : False
LowestDefectPri    : macRemErrXcon            HighestDefect     : defRemoteCCM
Defect Flags       : bDefRemoteCCM
Mac Address        : d0:0d:1e:00:00:02        ControlMep        : False
CcmLtmPriority     : 7
CcmTx              : 5336                     CcmSequenceErr    : 0
```

```
Fault Propagation  : disabled                 FacilityFault     : Notify
MA-CcmInterval     : 1                        MA-CcmHoldTime    : 0ms
Eth-1Dm Threshold  : 3(sec)                   MD-Level          : 0
Eth-Ais:           : Enabled                  Eth-Ais Rx Ais:   : No
Eth-Ais Tx Priorit*: 7                        Eth-Ais Rx Interv*: 1
Eth-Ais Tx Interva*: 1                        Eth-Ais Tx Counte*: 3515
Eth-Ais Tx Levels  : 5
Eth-Tst:           : Disabled
…snip…

# show service sap-using eth-cfm facility
===============================================================================
Service ETH-CFM Facility Information
===============================================================================
SapId             SvcId                          SAP AIS  SAP Tunnel SVC Tunnel
                                                          Fault      Fault
-------------------------------------------------------------------------------
1/1/2:100.31      100                            Enabled  Accept     Ignore
-------------------------------------------------------------------------------
No. of Facility SAPs: 1
===============================================================================
```

LAG Based MEP

LAG bundled ports provide both protection and scalability. Down MEPs configured on a LAG validates the connectivity of the LAG. Failure of this MEP causes the LAG to enter an operational `down` state. SAPs connected to the operationally `down` LAG transitions to operationally `down`. This triggers the configured reaction and processing similar to that of the port-based facility MEP. AIS is generated for those Epipe services with AIS enabled under the SAP. Local processing occurs for VPLS, IES and VPRN services that have experienced the SAP failure as a result of the LAG based SAP. Furthermore, fault propagation is invoked for any SAP with fault propagation operations enabled as a result of the failed LAG based SAP. LAG-based MEPs must be configured with a direction `down`.

LAG ETH-CFM PDUs are sent untagged because they are not specific to any service or VLAN. When running the combination of LAG-based MEPs and port-based MEPs, domain-level nesting rules must be adhered to for proper implementation, and is enforced by the CLI on the local node. As stated earlier, do not configure logical non-port-based MEPs, including service-based MEPs, to use level `0` for the ETH-CFM packets.

LAG-based MEPs are supported in both the IEEE 802.1ag and ITU-T Y.1731 contexts. Therefore, the Y.1731 context must be configured in order to run functions beyond those that are described as part of the IEEE 802.1ag standard. Since the recognition of fault is determined entirely by the ETH-CFM function, timeout conditions for the MEP occurs in 3.5 times the CCM interval. The LAG admin state or other failures that causes the LAG to completely fail, does not directly influence the MEP. The state of the MEP can only be influenced by the ETH-CFM function, specifically ETH-CC.

Since the LAG-based MEP selects a single member port to forward ETH-CFM packets, port-based facilities MEPs must be deployed to validate the individual member ports. Functional tests that require the ability to test individual member ports need to be performed from the port-based MEPs. The LAG-based MEPs validate only the LAG entity.

Figure 34 on page 130, provides an example how an ETH-CFM failure reacts with the various services that share that LAG. There is only one way the LAG state can trigger the propagation of failure, and that is using ETH-AIS. The carrier must enable CCM at the LAG level and a ETH-CCM defect condition exists. The red Epipe service generates AIS as a result of the LAG failure using the **client-meg-level** parameter configured on the LAG facility MEP. The green multipoint service takes location-configured action when the SAP transitions to the down operational state.

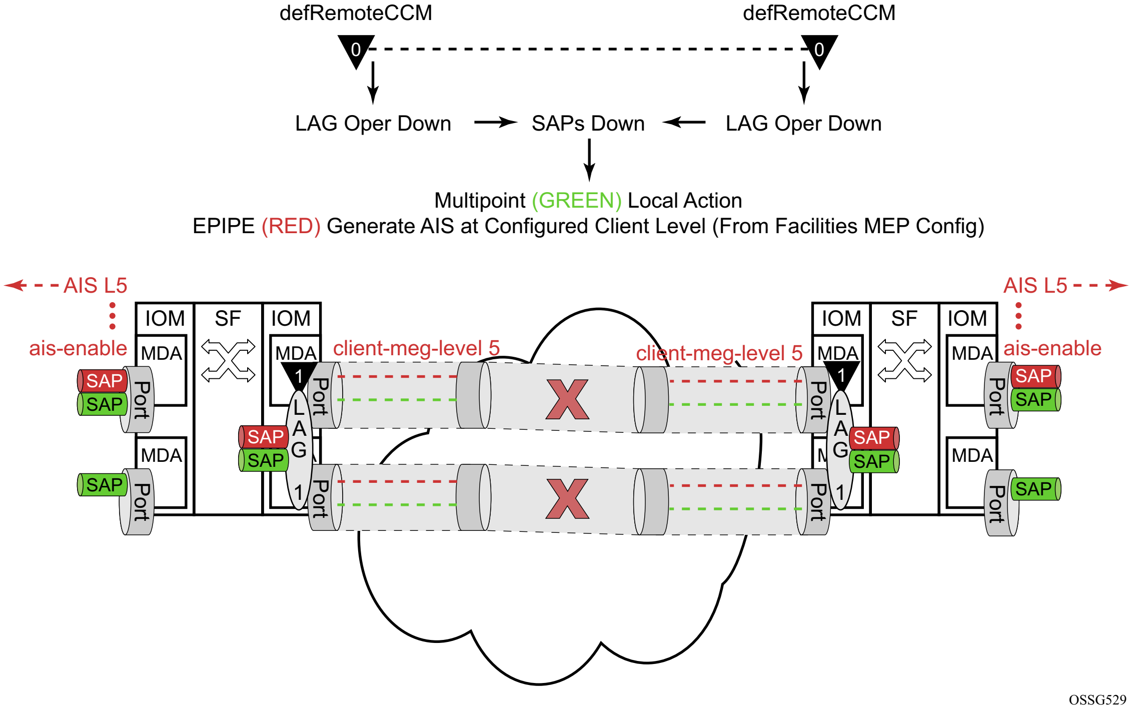

Figure 34: Fault Handling LAG MEP

LAG-based MEP are supported for MultiChassis LAG (MC-LAG) configurations.

A LAG facility MEP must not be configured with **facility-fault** when it is applied to an MC-LAG. Traffic will black hole when the LAG Facility MEP enters a defect state. The LAG enters an operational `down` state but the MC-LAG does not switch over to the peer node. This restriction does not include Tunnel Facility MEPs which are applied to a LAG with an outer VLAN. Tunnel facility MEPs do not control the operational state of the LAG because they are outer VLAN specific.

Example: LAG MEP Configuration

The following illustration, Figure 35, uses a port-based MEP to validate port-to-port connectivity.

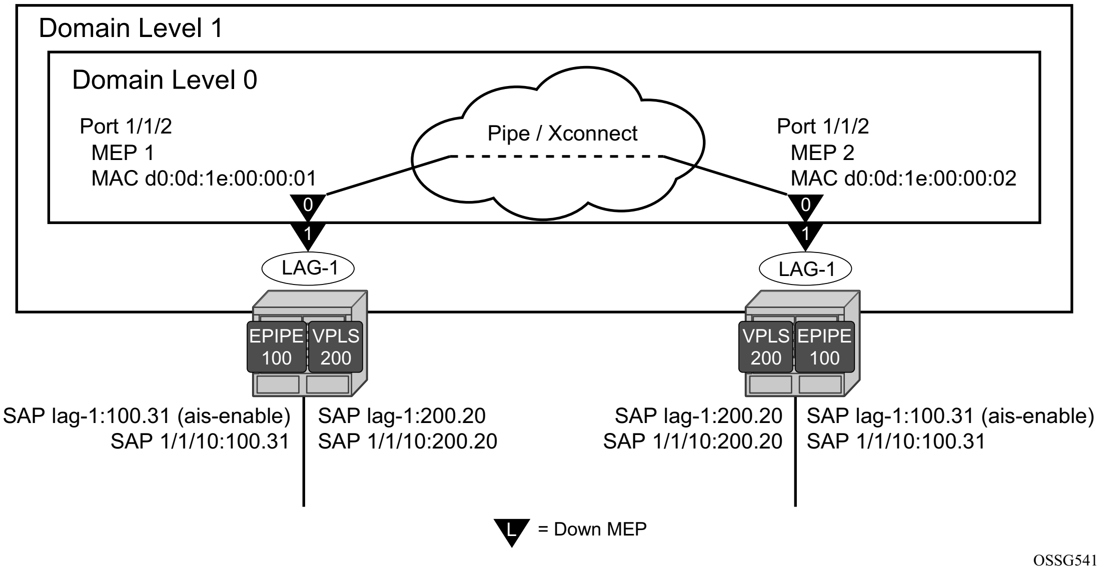

Figure 35: LAG MEP Example

With the introduction of the LAG, the port no longer has direct control over the services SAPs. The ais-enable command has been disabled from the port for this reason. The low-priority-defect condition has been modified to react to all defect conditions “allDef”, avoiding the unidirectional issue demonstrated in the previous port-based MEP example. A LAG MEP is built on top the LAG with the “facility-fault” option and ais-enable command with the associated client-meg-level. This allows the Epipe services to generate AIS when the LAG MEP enters any defect condition. This example introduce the use of a VPLS service. VPLS, IES and VPRN services do not support the generation of AIS as a result of a facility MEP failure. However, all service SAPs which correspond to the failed facility will transition to a down state. Epipe service also generates AIS in this example.

```
NODE1
config>eth-cfm# info
----------------------------------------------
        domain 1 format none level 1
            association 1 format icc-based name "FacilityLag01"
                ccm-interval 1
                remote-mepid 22
            exit
        exit
        domain 10 format none level 0
            association 1 format icc-based name "FacilityPort0"
                ccm-interval 1
                remote-mepid 2
            exit
        exit
----------------------------------------------

config>port# info
----------------------------------------------
        ethernet
            mode access
            encap-type qinq
            eth-cfm
                mep 1 domain 10 association 1
                    facility-fault
                    ccm-enable
                    low-priority-defect allDef
                    mac-address d0:0d:1e:00:00:01
                    no shutdown
                exit
            exit
            autonegotiate limited
        exit
        no shutdown
----------------------------------------------

config>lag# info
----------------------------------------------
        mode access
        encap-type qinq
        eth-cfm
            mep 11 domain 1 association 1
                ais-enable
                    client-meg-level 5
                exit
          ccm-enable
                facility-fault
                low-priority-defect allDef
                no shutdown
            exit
        exit
        port 1/1/2
        no shutdown
----------------------------------------------

config>service# info
----------------------------------------------
        customer 1 create
            description "Default customer"
        exit
        epipe 100 customer 1 create
            sap 1/1/10:100.31 create
            exit
            sap lag-1:100.31 create
                eth-cfm
                    ais-enable
                exit
            exit
            no shutdown
        exit
        vpls 200 customer 1 create
```

```
            stp
                shutdown
            exit
            sap 1/1/10:200.20 create
            exit
            sap lag-1:200.20 create
            exit
            no shutdown
        exit
----------------------------------------------

NODE2
config>eth-cfm# info
----------------------------------------------
        domain 1 format none level 1
            association 1 format icc-based name "FacilityLag01"
                ccm-interval 1
                remote-mepid 11
            exit
        exit
        domain 10 format none level 0
            association 1 format icc-based name "FacilityPort0"
                ccm-interval 1
                remote-mepid 1
            exit
        exit
----------------------------------------------

config>port# info
----------------------------------------------
        ethernet
            mode access
            encap-type qinq
            eth-cfm
                mep 2 domain 10 association 1
                    facility-fault
                    ccm-enable
                    low-priority-defect allDef
                    mac-address d0:0d:1e:00:00:02
                    no shutdown
                exit
            exit
            autonegotiate limited
        exit
        no shutdown
----------------------------------------------

config>lag# info
----------------------------------------------
        mode access
        encap-type qinq
        eth-cfm
            mep 22 domain 1 association 1
                ais-enable
                    client-meg-level 5
                exit
                facility-fault
                ccm-enable
                low-priority-defect allDef
                no shutdown
            exit
        exit
        port 1/1/2
        no shutdown
----------------------------------------------

config>service# info
----------------------------------------------
        customer 1 create
            description "Default customer"
        exit
        epipe 100 customer 1 create
            sap 1/1/10:100.31 create
            exit
            sap lag-1:100.31 create
                eth-cfm
                    ais-enable
                exit
            exit
            no shutdown
        exit
        vpls 200 customer 1 create
            stp
                shutdown
            exit
            sap 1/1/10:200.20 create
            exit
            sap lag-1:200.20 create
            exit
            no shutdown
        exit
----------------------------------------------
```

A fault is introduced that only affects the LAG MEP. The port MEP continues to validate the port, meaning that the port remains operationally up and the lag transitions to operation down. The LAG transition causes all the SAPs tied to the LAG to transition to down. The VPLS service reacts normally with the configured behavior as a result of a SAP down condition. The Epipe SAP also transitions to down, causing the operational state of the Epipe service to transition to down. In this case, AIS is enabled under the SAP in the service those AIS packets will still be generated out the mate SAP.

Output from one of the nodes is included below. Since both react in the same manner, output from both nodes is not shown.

```
NODE1
#show port
===============================================================================
Ports on Slot 1
===============================================================================
Port        Admin Link Port    Cfg  Oper LAG/ Port Port Port   C/QS/S/XFP/
Id          State      State   MTU  MTU  Bndl Mode Encp Type   MDIMDX
-------------------------------------------------------------------------------
…snip..
1/1/2       Up    Yes  Up      1522 1522    - accs qinq xcme
…snip..
```

```
show eth-cfm mep 11 domain 1 association 1
===============================================================================
Eth-Cfm MEP Configuration Information
===============================================================================
Md-index           : 1                        Direction         : Down
Ma-index           : 1                        Admin             : Enabled
MepId              : 11                       CCM-Enable        : Disabled
Port               : lag-1                    VLAN              : 0
Description        : (Not Specified)
FngState           : fngDefectReported        ControlMep        : False
LowestDefectPri    : allDef                   HighestDefect     : defRDICCM
Defect Flags       : bDefRDICCM
Mac Address        : 90:f3:ff:00:01:41        ControlMep        : False
CcmLtmPriority     : 7
CcmTx              : 4428                     CcmSequenceErr    : 0
Fault Propagation  : disabled                 FacilityFault     : Notify
MA-CcmInterval     : 1                        MA-CcmHoldTime    : 0ms
Eth-1Dm Threshold  : 3(sec)                   MD-Level          : 1
Eth-Ais:           : Enabled                  Eth-Ais Rx Ais:   : No
Eth-Ais Tx Priorit*: 7                        Eth-Ais Rx Interv*: 1
Eth-Ais Tx Interva*: 1                        Eth-Ais Tx Counte*: 1085
Eth-Ais Tx Levels  : 5
Eth-Tst:           : Disabled
…snip…

# show service sap-using eth-cfm facility
===============================================================================
Service ETH-CFM Facility Information
===============================================================================
SapId             SvcId                          SAP AIS  SAP Tunnel SVC Tunnel
                                                          Fault      Fault
-------------------------------------------------------------------------------
lag-1:100.31      100                            Enabled  Accept     Ignore
lag-1:200.20      200                            Disabled Accept     Ignore
-------------------------------------------------------------------------------
No. of Facility SAPs: 2
===============================================================================

# show eth-cfm cfm-stack-table facility
===============================================================================
CFM Stack Table Defect Legend:
R = Rdi, M = MacStatus, C = RemoteCCM, E = ErrorCCM, X = XconCCM, A = AisRx

===============================================================================
CFM Facility Port Stack Table
===============================================================================
Port     Tunnel    Lvl Dir  Md-index   Ma-index   MepId  Mac-address     Defect
-------------------------------------------------------------------------------
1/1/2     0          0 Down         10          1    1 d0:0d:1e:00:00:01 ------
===============================================================================

===============================================================================
CFM Facility LAG Stack Table
===============================================================================
Lag      Tunnel    Lvl Dir  Md-index   Ma-index   MepId  Mac-address     Defect
-------------------------------------------------------------------------------
lag-1     0          1 Down          1          1   11 90:f3:ff:00:01:41 R-----
```

```
===============================================================================

A:Dut-C# show service id 1 sap 1/1/1 base
===============================================================================
Service Access Points(SAP)
===============================================================================
Service Id         : 1
SAP                : 1/1/1                    Encap             : null
Description        : (Not Specified)
Admin State        : Up                       Oper State        : Up
Flags              : None
Multi Svc Site     : None
Last Status Change : 02/24/2012 11:37:55
Last Mgmt Change   : 02/24/2012 11:31:32
Sub Type           : regular
Dot1Q Ethertype    : 0x8100                   QinQ Ethertype    : 0x8100
Split Horizon Group: (Not Specified)

Max Nbr of MAC Addr: No Limit                 Total MAC Addr    : 0
Learned MAC Addr   : 0                        Static MAC Addr   : 0
Admin MTU          : 1514                     Oper MTU          : 1514
Ingr IP Fltr-Id    : n/a                      Egr IP Fltr-Id    : n/a
Ingr Mac Fltr-Id   : n/a                      Egr Mac Fltr-Id   : n/a
Ingr IPv6 Fltr-Id  : n/a                      Egr IPv6 Fltr-Id  : n/a
tod-suite          : None                     qinq-pbit-marking : both
Ing Agg Rate Limit : max                      Egr Agg Rate Limit: max
Q Frame-Based Acct : Disabled
ARP Reply Agent    : Disabled                 Host Conn Verify  : Disabled
Mac Learning       : Enabled                  Discard Unkwn Srce: Disabled
Mac Aging          : Enabled                  Mac Pinning       : Disabled
BPDU Translation   : Disabled
L2PT Termination   : Disabled
Vlan-translation   : None

Acct. Pol          : None                     Collect Stats     : Disabled

Anti Spoofing      : None                     Dynamic Hosts     : Enabled
Avl Static Hosts   : 0                        Tot Static Hosts  : 0
Calling-Station-Id : n/a
Application Profile: None

Oper Group         : (none)                   Monitor Oper Grp  : (none)
Restr MacProt Src  : Disabled                 Restr MacUnpr Dst : Disabled
Auto Learn Mac Prot: Disabled                 RestProtSrcMacAct : Disable
Time to RetryReset : never                    Retries Left      : 3
Mac Move           : Blockable                Blockable Level   : Tertiary
Egr MCast Grp      :
Auth Policy        : None

-------------------------------------------------------------------------------
ETH-CFM SAP specifics
-------------------------------------------------------------------------------
Tunnel Faults      : n/a                      AIS               : Disabled
MC Prop-Hold-Timer : n/a                      V-MEP Filtering   : Disabled
===============================================================================
A:Dut-C#
```

Tunnel Based MEP

The concept of a logical tunnel carrying many unique and individual services has been deployed in many networks on QinQ encapsulated access ports where the outer VLAN represents the common transports and the inner VLAN represents the specific service. Typically, the tunnel transparently passes frames from multiple services through some common network. Tunnel MEPs are logically configured on the Port or LAG and outer VLAN for access ports use QinQ Ethernet encapsulation. Service processing is done after the tunnel MEP. This means that any service-based MEPs are required to be a higher level than that of the tunnel MEP. Tunnel MEPs are only supported on LAGs that are configured with QinQ encapsulation and must specify the outer VLAN.

The Tunnel MEP must validate connectivity between the tunnel end points. As with all facility MEPs, this is a point-to-point relationship between the local MEP and one remote MEP. By default, the MEP configured at the tunnel level performs only alarming functions. Actionable functions such as AIS, SAP transition, and fault propagation requires the operator to enable these functions.

The tunnel MEP must first be configured to take action when the MEP enters a fault state, similar to all other facilities MEPs. In order for the individual services to share the fate of the tunnel, each service must accept the facility MEP state. This is service-dependent and depends on the desired goals. Services share the tunnel fate based on the lag-id and the outer VLAN.

Epipe services support the **ais-enable** configuration option on the SAP. Enabling this option generates AIS in the event the tunnel MEP has entered a fault state as a result of ETH-CC failure, similar to other facility MEPs. However, since the individual SAPs configured within the different services are not directly affected by the tunnel MEP, an additional configuration is necessary to perform local SAP transitions, in the case of VPLS, EIS and VPRN services and OAM mapping functions for Epipe services.

The **tunnel-fault** service-level command configured on an Epipe allows SAP flags to be set and fault propagation and OAM mapping functions between technology. The operational state of the SAP remains `up`. The operator needs to determine if the AIS generation of fault propagation is the best approach in their specific network. It is possible to configure both **ais-enable** and **tunnel-fault** accept within the Epipe service. However, this may generate multiple ETH-CFM packets, or multiple actions as a result of a single failure.

The **tunnel-fault accept** service level option is also available under Epipe, VPLS and IES services hierarchy level within the CLI. This allows for a tunnel fault to share fate with these service SAPs. For the non-Epipe services, the SAP enters an operationally **down** state, and normal processing occurs as a result of the SAP transition. In order to generate any ETH-CC based fault propagation, **suspend-cmm** or **use-int-stat-tlv**, this requires service-based MEPs that are actively running CCM with a peer.

The **tunnel-fault** configuration options occur in two levels of the CLI hierarchy: service level and SAP level. Both of the levels within a service and within the SAP (whose underlying port and outer tag has a tunnel MEP) must be set to accept, in order to have the function enabled. By default the **tunnel-fault** is set to ignore at the service level and accept at the SAP level. This means that a single **tunnel-fault** `accept` at the service level will enable fault operations for all SAPs in the service. The operator is free to enable and disable on specific SAPs by choosing the `ignore` option under the individual SAP. The combination of **accept** at the service level and `ignore` at the SAP level prevents that specific SAP from recognizing fault. AIS generation for Epipe services is not controlled by the **tunnel-fault** configuration options.

Specific to tunnel MEPs, reception of AIS on the tunnel MEP causes AIS to be cut through to all Epipe services that have the `ais-enabled` command configured under the SAP. During a fault condition, it is important that the AIS configuration under the tunnel MEP not be modified. This causes increased network element CPU processing requirements and in scaled environments transitioning this command during a heavily loaded fault condition, where highly scaled SAPs are linked to the fate of the tunnel MEP, may cause the system to spend more than normal processing time to be spent dealing with this artificially induced clear and fault situation. It is not expected that operators perform these types of tasks in production networks. Reception of AIS will not trigger a fault condition or AIS to be cut through when sub second CCM intervals have been configured on the Tunnel MEP.

Service-based MEPs may also be configured as normal for all services. They perform normal processing tasks, including service-based MEP with fault propagation.

As with all other facility MEPs, use only ETH-CFM functions to cause the Tunnel MEP to enter the fault state. Tunnel MEPs support sub second ccm-intervals on selected hardware. Tunnel MEPs must be configured with a direction of down. UP MEPs are not supported as part of the facility MEP concept.

LAG-based MEPs and LAG-based tunnel MEPs cannot be configured on the same LAG. Port-based MEPs may be configured on the LAG member ports of a tunnel MEP as long as they follow the requirements for port-based MEPs on LAG member ports. All those consideration are applicable here, including nesting and port-level control only without propagation.

Port-based MEPs and Port-based tunnel MEPs cannot be configured on the same port.

LAG-based Tunnel MEPs are supported in MultiChassis LAG (MC-LAG) configuration. However, sub second CCM enabled intervals should not be configured when the LAG-based Tunnel MEP utilizes the transport of an MC-LAG. Only one second and above CCM intervals should be used. Not all platforms support sub second CCM enable Tunnel MEPs.

Tunnel MEPs are meant to propagate fault from one segment to the other for Epipe services. Figure 36 on page 139 shows how individual Epipes have SAPs connecting to a legacy network. A MEP is configured at the tunnel level and peers with a single remote peer MEP.

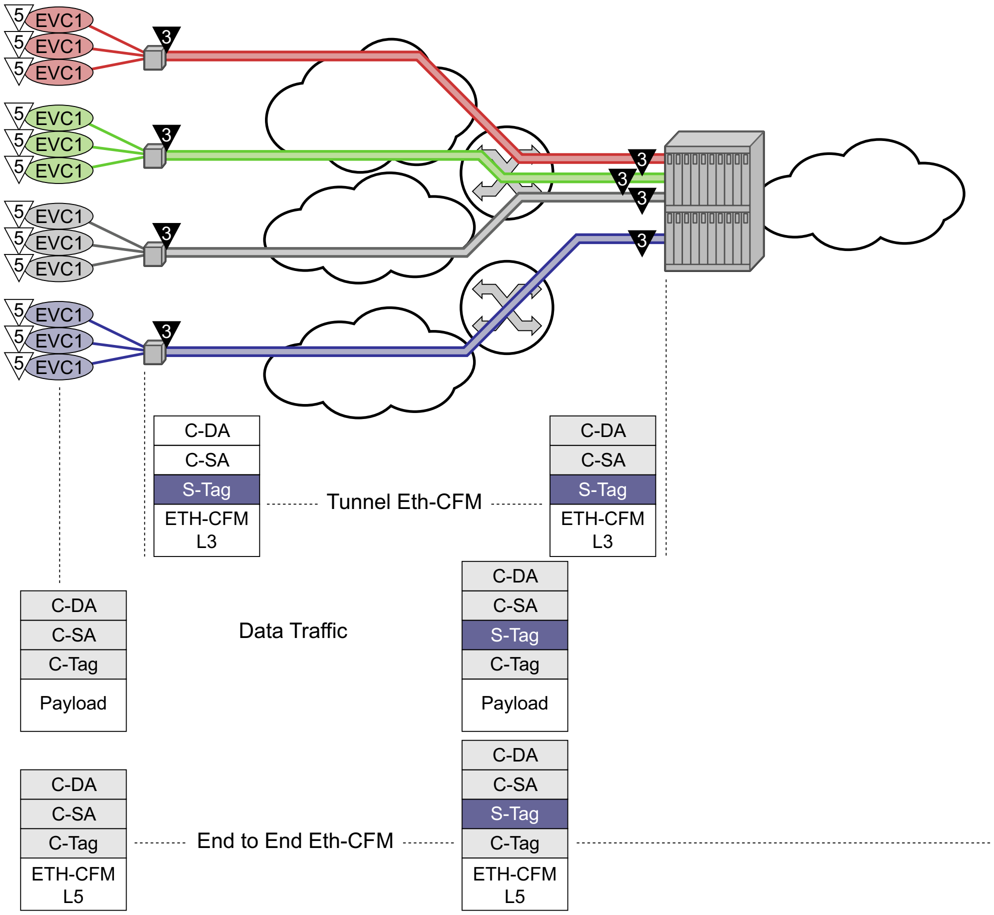

Figure 36: Tunnel Concepts and Encapsulation

This is only one example of a tagged service. The principles of a tunnel MEP may be applied to other service as applicable. Remember that tunnel MEPs are only supported on LAGs that are configured with QinQ encapsulation and must have an outer VLAN.

Individual services can be monitored end-to-end by placing a MEP on the service endpoint at the CPE, denoted by the MEP at level 5 on the individual EVC (customer levels 5-7). The Network Interface Demarcation (NID) typically places a single tag, outer or only, on the customer traffic. This is cross connected to the proper connection in the access network and eventually arrive on the Ethernet Aggregation Switch. The connection between the legacy or access network and the aggregation switch must be either a LAG bundle or MC-LAG in order for tunnel MEPs to be configured.

Since there can be a large number of services transported by a single tunnel, the MEP executing at the tunnel-level reduces network overhead and simplifies the configuration. It is important to note that all services in the tunnel must share a common physical path.

A SAP is needed in order for the Tunnel MEP to extract the tunnel MEP ETH-CFM packets at the appropriate level. No SAP record is created by default. A service must already exist that includes a SAP in the form `lag-id:vid.*` or `lag-id:vid.0` where the vid matches the outer VLAN in which the tunnel is to monitor. Since the ETH-CFM traffic arrives at the Ethernet aggregation node as a single outer tag with no inner tag, the operator may want to consider the ability to configure the `lag-id:vid.0` to accept untagged only frames with the matching outer tag and no inner tag. The global command **configure>system->ethernet>new-qinq-untagged-sap** is available to enable this functionality. By default both the `vid.*` and `vid.0` accepts all packets that match the outer vid and any inner vid. If no SAP record exists for this VLAN, one must be created manually. Manually creating this SAP requires a service context. Alcatel-Lucent recommends that an Epipe service be configured with this single SAP, preventing any flooding of packets. It is possible to use a VPLS instance and combine many tunnel SAP records into a single service instance. However, configuration errors may result in leakage because of the multipoint nature of a VPLS service. Regardless of the service type chosen, it should be in a `shutdown` state. Also, normal ETH-CFM rules apply. ETH-CFM packets arriving on the SAP passes all ETH-CFM packets at and below the tunnel MEP to the ETH-CFM application for processing.

The goal of a Tunnel MEP is to validate an attachment circuit and relate the state to services that share the same LAG and outer VLAN to other services across the network. Tunnel MEPs are not intended for propagating fault between two endpoints that share the same LAG and outer VLAN. For this reason, locally switched circuits that share the same LAG and the same outer tag must not use the **ais-enable** function under those SAPs. As an example, lag-1 may have two SAPs associated with it: lag-1:1.1 and lag-1:1.2. These two SAP represent two different endpoints on the same LAG using the same outer VLAN. In this case, if the `ais-enable` is configured under both SAPs, AIS functionality does not work properly. Normal fault propagation could be used in this case instead. Since the tunnel MEP is validating the common physical path and these two MEPs share the common physical path, there is no reason to propagate fault. Service-based MEPs could be configured on the endpoints in order to validate the connectivity between the two endpoints when this type of model is deployed. However, two SAPs that are connected to different LAGs is a supported configuration. An example of this would be lag-1:1.1 and lag-2:1.1.

Sub second Tunnel MEPs will be monitored for every three seconds to ensure that they are not continuously bouncing and consuming an unfair allocation of ETH-CFM resources. A sub second MEP will only be allowed three operational status changes in a three second window before holding the state for the remaining time in that window. Messages will be paced from Tunnel MEPs. Fault propagation depends on factors such as how busy the node is, or how scaled the node configuration is.

Five percent of the operational/negotiated port speed not physical speed is available for Tunnel MEP control traffic. When applying this to the LAG-based Tunnel MEPs the five percent is derived from the lowest speed of a single member port in the bundle. If this bandwidth percentage required for ETH-CFM is exceeded the ETH-CFM packets will not be able to be sent and failures will occur. As an example, a physical port of 1Gbps that has negotiated an operational speed of 100Mbps with a peer will be allowed to send up to a maximum of 5Mbps of Tunnel MEP control traffic.

Example: Tunnel MEP Configuration

The following illustration, Figure 37, shows how fate can be shared between the Tunnel MEP and the services configured on the same LAG and outer VLAN.

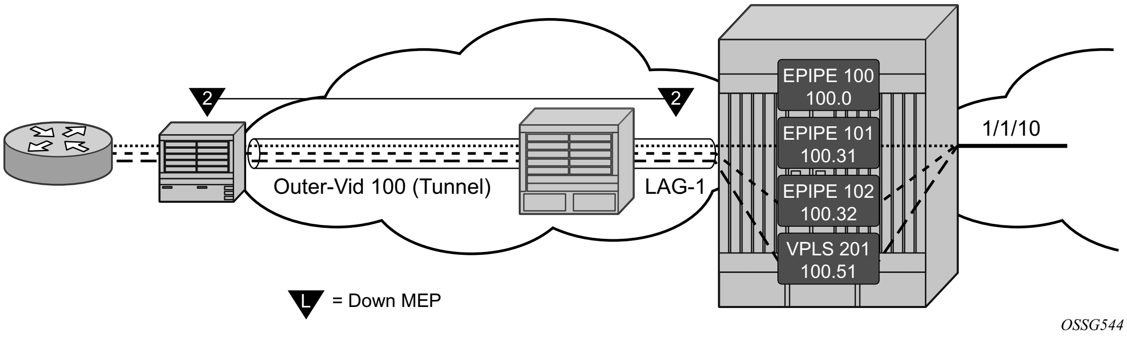

Figure 37: Tunnel MEP Example

In this example, a single Tunnel, LAG-1 outer VLAN 100, carries three services. Epipe 101, Epipe 102 and VPLS 201 are the service extraction points on the aggregation node. Epipe 100 is the extraction point for the Tunnel MEP eth-cfm traffic. This is a single SAP Epipe that is operationally shutdown. One common configuration error when using Tunnel MEPs is the lack extraction on the aggregation node, causing unidirectional failures. The aggregation node is sending eth-cfm traffic to the NID, but is not extracting the eth-cfm traffic that the NID is sending.

Epipe 101 is configured to accept the tunnel MEP fate and generate AIS.

Epipe 102 is configured to accept the tunnel MEP state and apply fault propagation rules. If the network-side mate were an SDP binding, then the applicable setting of the LDP status bits are in the header. Since this example uses an Ethernet SAP as the mate, and only tunnel fault-accept is configured with no ais-enable, only the SAP flag is set to indicate an error.

VPLS 201 also shares the fate of the tunnel MEP. The tunnel-fault accept transitions the SAP to operationally down. Any configured event that occurs because of a SAP down for the VPLS also occur.

Only the configuration for the aggregation node is shown below. The NID configuration is not required to show how this function works.

```
AGGREGATION NODE
config>eth-cfm# info
----------------------------------------------
        domain 2 format none level 2
            association 1 format icc-based name "FacilityTun01"
                ccm-interval 1
                remote-mepid 101
            exit
        exit
----------------------------------------------

config>lag# info
----------------------------------------------
        mode access
        encap-type qinq
        eth-cfm
            mep 100 domain 2 association 1 vlan 100
                description "Tunnel Facility MEP - Do NOT Delete"
                ais-enable
                    client-meg-level 5
                exit
                facility-fault
                ccm-enable
                low-priority-defect allDef
                no shutdown
            exit
        exit
        port 1/1/2
        no shutdown
----------------------------------------------

config>service# info
----------------------------------------------
        customer 1 create
            description "Default customer"
        exit
        epipe 100 customer 1 create
            shutdown
            description "Tunnel Extraction Service"
            sap lag-1:100.0 create
            exit
        exit
        epipe 101 customer 1 create
            description "Customer Service 100.31"
            sap 1/1/10:100.31 create
            exit
            sap lag-1:100.31 create
                eth-cfm
                    ais-enable
                exit
            exit
            no shutdown
        exit
        epipe 102 customer 1 create
            description "Customer Service 100.32"
            eth-cfm
                tunnel-fault accept
            exit
            sap 1/1/10:100.32 create
            exit
            sap lag-1:100.32 create
            exit
            no shutdown
        exit
        vpls 201 customer 1 create
            description "Customer Service 100.51"
            stp
```

```
                shutdown
            exit
            eth-cfm
                tunnel-fault accept
            exit
            sap 1/1/10:100.51 create
            exit
            sap lag-1:100.51 create
            exit
            no shutdown
        exit
----------------------------------------------
```

```
# show eth-cfm mep 100 domain 2 association 1
===============================================================================
Eth-Cfm MEP Configuration Information
===============================================================================
Md-index           : 2                        Direction         : Down
Ma-index           : 1                        Admin             : Enabled
MepId              : 100                      CCM-Enable        : Enabled
Port               : lag-1                    VLAN              : 100
Description        : Tunnel Facility MEP - Do NOT Delete
FngState           : fngReset                 ControlMep        : False
LowestDefectPri    : allDef                   HighestDefect     : none
Defect Flags       : None
Mac Address        : 90:f3:ff:00:01:41        ControlMep        : False
CcmLtmPriority     : 7
CcmTx              : 3958                     CcmSequenceErr    : 0
Fault Propagation  : disabled                 FacilityFault     : Notify
MA-CcmInterval     : 1                        MA-CcmHoldTime    : 0ms
Eth-1Dm Threshold  : 3(sec)                   MD-Level          : 2
Eth-Ais:           : Enabled                  Eth-Ais Rx Ais:   : No
Eth-Ais Tx Priorit*: 7                        Eth-Ais Rx Interv*: 1
Eth-Ais Tx Interva*: 1                        Eth-Ais Tx Counte*: 175
Eth-Ais Tx Levels  : 5
Eth-Tst:           : Disabled

Redundancy:
    MC-LAG State   : n/a

CcmLastFailure Frame:
    None

XconCcmFailure Frame:
    None
===============================================================================

# show eth-cfm cfm-stack-table facility all-tunnel-meps
===============================================================================
CFM Stack Table Defect Legend:
R = Rdi, M = MacStatus, C = RemoteCCM, E = ErrorCCM, X = XconCCM, A = AisRx

===============================================================================
CFM Facility LAG Stack Table
===============================================================================
Lag      Tunnel    Lvl Dir  Md-index   Ma-index   MepId  Mac-address     Defect
-------------------------------------------------------------------------------
lag-1     100        2 Down          2          1  100 90:f3:ff:00:01:41 ------
```

```
===============================================================================

# show service sap-using eth-cfm facility

===============================================================================
Service ETH-CFM Facility Information
===============================================================================
SapId             SvcId                          SAP AIS  SAP Tunnel SVC Tunnel
                                                          Fault      Fault
-------------------------------------------------------------------------------
lag-1:100.0       100                            Disabled Accept     Ignore
lag-1:100.31      101                            Enabled  Accept     Ignore
lag-1:100.32      102                            Disabled Accept     Accept
lag-1:100.51      201                            Disabled Accept     Accept
-------------------------------------------------------------------------------
No. of Facility SAPs: 4
===============================================================================

When the tunnel MEP enters a fault state
•    Epipe 101 will start to generate AIS out the mate sap
•    Epipe 102 SAP flag will be set
•    VPLS 201 SAP will go down

Output from one of the nodes is included below.  Since both will react in the same manner
output from both nodes is not required.

AGGREGATION NODE

# show eth-cfm cfm-stack-table facility all-tunnel-meps
===============================================================================
CFM Stack Table Defect Legend:
R = Rdi, M = MacStatus, C = RemoteCCM, E = ErrorCCM, X = XconCCM, A = AisRx
===============================================================================
CFM Facility LAG Stack Table
===============================================================================
Lag      Tunnel    Lvl Dir  Md-index   Ma-index   MepId  Mac-address     Defect
-------------------------------------------------------------------------------
lag-1     100        2 Down          2          1  100 90:f3:ff:00:01:41 --C---
===============================================================================

# show service sap-using eth-cfm facility tunnel 100
===============================================================================
Service ETH-CFM Facility Information
===============================================================================
SapId             SvcId                          SAP AIS  SAP Tunnel SVC Tunnel
                                                          Fault      Fault
-------------------------------------------------------------------------------
lag-1:100.0       100                            Disabled Accept     Ignore
lag-1:100.31      101                            Enabled  Accept     Ignore
lag-1:100.32      102                            Disabled Accept     Accept
lag-1:100.51      201                            Disabled Accept     Accept
-------------------------------------------------------------------------------
No. of Facility SAPs: 4
===============================================================================

# show eth-cfm mep 100 domain 2 association 1
===============================================================================
Eth-Cfm MEP Configuration Information
===============================================================================
Md-index           : 2                          Direction         : Down
Ma-index           : 1                          Admin             : Enabled
MepId              : 100                        CCM-Enable        : Enabled
Port               : lag-1                      VLAN              : 100
Description        : Tunnel Facility MEP - Do NOT Delete
FngState           : fngDefectReported          ControlMep        : False
LowestDefectPri    : allDef                     HighestDefect     : defRemoteCCM
Defect Flags       : bDefRemoteCCM
Mac Address        : 90:f3:ff:00:01:41          ControlMep        : False
CcmLtmPriority     : 7
CcmTx              : 4211                       CcmSequenceErr    : 0
Fault Propagation  : disabled                   FacilityFault     : Notify
MA-CcmInterval     : 1                          MA-CcmHoldTime    : 0ms
Eth-1Dm Threshold  : 3(sec)                     MD-Level          : 2
Eth-Ais:           : Enabled                    Eth-Ais Rx Ais:   : No
Eth-Ais Tx Priorit*: 7                          Eth-Ais Rx Interv*: 1
Eth-Ais Tx Interva*: 1                          Eth-Ais Tx Counte*: 215
Eth-Ais Tx Levels  : 5
Eth-Tst:           : Disabled

Redundancy:
    MC-LAG State   : n/a

CcmLastFailure Frame:
    None

XconCcmFailure Frame:
    None
===============================================================================

show service id 101 base
===============================================================================
Service Basic Information
===============================================================================
Service Id        : 101                 Vpn Id            : 0
Service Type      : Epipe
Name              : (Not Specified)
Description       : Customer Service 100.31
Customer Id       : 1
Last Status Change: 02/04/2010 15:53:12
Last Mgmt Change  : 02/04/2010 16:31:00
Admin State       : Up                  Oper State        : Up
MTU               : 1514
Vc Switching      : False
SAP Count         : 2                   SDP Bind Count    : 0
Per Svc Hashing   : Disabled
Force QTag Fwd    : Disabled

-------------------------------------------------------------------------------
Service Access & Destination Points
-------------------------------------------------------------------------------
Identifier                               Type         AdmMTU  OprMTU  Adm  Opr
-------------------------------------------------------------------------------
sap:1/1/10:100.31                        qinq         1522    1522    Up   Up
sap:lag-1:100.31                         qinq         1522    1522    Up   Up
===============================================================================

# show service id 102 base
===============================================================================
Service Basic Information
===============================================================================
Service Id        : 102                 Vpn Id            : 0
Service Type      : Epipe
Name              : (Not Specified)
Description       : Customer Service 100.32
Customer Id       : 1
Last Status Change: 02/04/2010 15:45:07
Last Mgmt Change  : 02/04/2010 16:30:43
Admin State       : Up                  Oper State        : Up
MTU               : 1514
Vc Switching      : False
SAP Count         : 2                   SDP Bind Count    : 0
Per Svc Hashing   : Disabled
Force QTag Fwd    : Disabled

-------------------------------------------------------------------------------
Service Access & Destination Points
-------------------------------------------------------------------------------
Identifier                               Type         AdmMTU  OprMTU  Adm  Opr
-------------------------------------------------------------------------------
sap:1/1/10:100.32                        qinq         1522    1522    Up   Up
sap:lag-1:100.32                         qinq         1522    1522    Up   Up
===============================================================================

# show service id 102 sap lag-1:100.32
===============================================================================
Service Access Points(SAP)
===============================================================================
Service Id         : 102
SAP                : lag-1:100.32             Encap             : qinq
QinQ Dot1p         : Default
Description        : (Not Specified)
Admin State        : Up                       Oper State        : Up
Flags              : OamTunnelMEPFault
Multi Svc Site     : None
Last Status Change : 02/04/2010 15:45:07
Last Mgmt Change   : 02/04/2010 15:44:26

-------------------------------------------------------------------------------
ETH-CFM SAP specifics
-------------------------------------------------------------------------------
Tunnel Faults      : accept                   AIS               : Disabled
MC Prop-Hold-Timer : n/a
===============================================================================

# show service id 201 base
===============================================================================
Service Basic Information
===============================================================================
Service Id        : 201                 Vpn Id            : 0
Service Type      : VPLS
Name              : (Not Specified)
Description       : Customer Service 100.51
Customer Id       : 1
Last Status Change: 02/04/2010 15:46:03
Last Mgmt Change  : 02/04/2010 16:30:29
Admin State       : Up                  Oper State        : Up
MTU               : 1514                Def. Mesh VC Id   : 201
```

```
SAP Count         : 2                   SDP Bind Count    : 0
Snd Flush on Fail : Disabled            Host Conn Verify  : Disabled
Propagate MacFlush: Disabled            Per Svc Hashing   : Disabled
Allow IP Intf Bind: Disabled
Def. Gateway IP   : None
Def. Gateway MAC  : None
Temp Flood Time   : Disabled            Temp Flood        : Inactive
Temp Flood Chg Cnt: 0

-------------------------------------------------------------------------------
Service Access & Destination Points
-------------------------------------------------------------------------------
Identifier                               Type         AdmMTU  OprMTU  Adm  Opr
-------------------------------------------------------------------------------
sap:1/1/10:100.51                        qinq         1522    1522    Up   Up
sap:lag-1:100.51                         qinq         1522    1522    Up   Down
===============================================================================
```

Router Interface MEP

MEPs and associated on-demand troubleshooting functions act as router interfaces that are part of the base routing instance. This feature allows the operator to verify Layer 2 transport that connects the Layer 3 interfaces.

Router interfaces MEPs are supported for all router interface instances (null port 1/1/1, dot1q port 1/1/3:vid, null LAG-lag-id and dot1q LAG-lag-id:vid).

Example: Router MEP Configuration

The following illustration, Figure 38, shows how a Router Facility MEP can be configured on a routed interface in the base router instance.

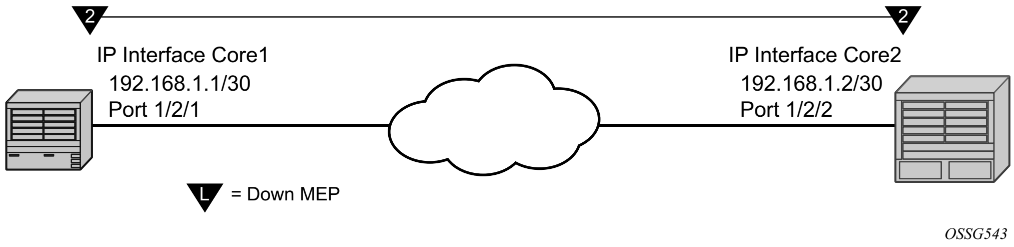

Figure 38: Router MEP Example

ETH-CFM tools for proactive management (ETH-CC), troubleshooting (Loopback, Linktrace, etc…) and profiling (Delay Measurement, etc…) are supported. The configuration and some ETH-CFM test commands are shown for Node1 (left). Following the on-demand test output, the configuration for Node 2 is included for completeness, without repeating the on-demand tests.

```
NODE1
config>port# info
----------------------------------------------
        ethernet
        exit
        no shutdown
----------------------------------------------

config>eth-cfm# info
----------------------------------------------
        domain 2 format none level 2
            association 2 format icc-based name "FacilityRtr01"
            exit
        exit
----------------------------------------------
```

```
config>router# info
----------------------------------------------
#--------------------------------------------------
echo "IP Configuration"
#--------------------------------------------------
        interface "Core1"
            address 192.168.1.1/30
            port 1/2/1
            eth-cfm
                mep 1 domain 2 association 2
                    mac-address d0:0d:1e:00:00:01
                    no shutdown
                exit
            exit
        exit
        interface "system"
        exit
----------------------------------------------

# show eth-cfm cfm-stack-table facility all-router-interfaces
===============================================================================
CFM Stack Table Defect Legend:
R = Rdi, M = MacStatus, C = RemoteCCM, E = ErrorCCM, X = XconCCM, A = AisRx

===============================================================================
CFM Facility Interface Stack Table
===============================================================================
Interface          Lvl Dir  Md-index   Ma-index   MepId  Mac-address     Defect
-------------------------------------------------------------------------------
Core1                2 Down          2          2    1 d0:0d:1e:00:00:01 ------
===============================================================================

# show eth-cfm cfm-stack-table facility all-router-interfaces
===============================================================================
CFM Stack Table Defect Legend:
R = Rdi, M = MacStatus, C = RemoteCCM, E = ErrorCCM, X = XconCCM, A = AisRx

===============================================================================
CFM Facility Interface Stack Table
===============================================================================
Interface          Lvl Dir  Md-index   Ma-index   MepId  Mac-address     Defect
-------------------------------------------------------------------------------
Core1                2 Down          2          2    1 d0:0d:1e:00:00:01 ------
===============================================================================

# oam eth-cfm loopback d0:0d:1e:00:00:02 mep 1 domain 2 association 2
 send-count 5
Eth-Cfm Loopback Test Initiated: Mac-Address: d0:0d:1e:00:00:02, out service: 0
Sent 5 packets, received 5 packets [0 out-of-order, 0 Bad Msdu]

# oam eth-cfm linktrace d0:0d:1e:00:00:02 mep 1 domain 2 association
2
Index Ingress Mac          Egress Mac           Relay      Action
----- -------------------- -------------------- ---------- ----------
1     D0:0D:1E:00:00:02    00:00:00:00:00:00    n/a        terminate
----- -------------------- -------------------- ---------- ----------
No more responses received in the last 6 seconds.
```

```
# oam eth-cfm two-way-delay-test d0:0d:1e:00:00:02 mep 1 domain 2 association 2
Two-Way-Delay-Test Response:
Delay 1130 microseconds        Variation 63 microseconds

# oam eth-cfm two-way-delay-test d0:0d:1e:00:00:02 mep 1 domain 2 association 2
Two-Way-Delay-Test Response:
Delay 1218 microseconds        Variation 88 microseconds
```

```
NODE2
config>port# info
----------------------------------------------
        ethernet
        exit
        no shutdown
----------------------------------------------

config>eth-cfm# info
----------------------------------------------
        domain 2 format none level 2
            association 2 format icc-based name "FacilityRtr01"
            exit
        exit
----------------------------------------------

config>router# info
----------------------------------------------
#--------------------------------------------------
echo "IP Configuration"
#--------------------------------------------------
        interface "Core2"
            address 192.168.1.2/30
            port 1/2/2
            eth-cfm
                mep 2 domain 2 association 2
                    mac-address d0:0d:1e:00:00:02
                    no shutdown
                exit
            exit
        exit
        interface "system"
        exit
----------------------------------------------
```

Hardware Support

All facility MEPs require a minimum of IOM3/IMM. However, only the facility MEP has an IOM-specific requirement. SAPs and ports that are not configured as part of facility MEPs are not restricted to a specific IOM. For example, a Tunnel MEP would be required to meet the minimum IOM requirement, similar to the fated shared service SAPs. However, the mate or egress SAP or binding is not required to meet the facility MEP requirement. Of course, there may be other reasons why a mate SAP or binding requires specific IOM/IMM that are outside that of facility MEPs. Similarly, a LAG MEP requires all port members to meet the IOM/IMM requirements for facility MEPs.

Table 4 provides an overview of Facility MEP support.

Table 4: Facility MEP Support Overview

	Port MEPs	**Tunnel MEPs**		**LAG MEPs**	**Router MEPs**
		Port	**LAG**		
Sub Second	Yes	Yes	Yes	Yes	Yes
Port:					
Hybrid	Dot1q/QinQ	QinQ	QinQ	Dot1q/QinQ	Dot1q/QinQ
Network	Null/Dot1q	no	no	Null/Dot1q	Null/QinQ
Access	Null/Dot1q/ QinQ	QinQ	QinQ	Null/Dot1q/QinQ	N/A
CCM	Yes	Yes	Yes	Yes	Yes
Y.1731 PM Tools	Yes	Yes	Yes	Yes	Yes
AIS Reception	No	Yes	Yes	No	No
Facility Fault	Controls port operational state Failure=Link Up Success=Up	Controls shared fate service SAPs and EPIPE AIS	Controls Shared fate service SAPs and Epipe AIS	Controls LAG operational state Failure=Oper: down, Success=Oper=up	Controls IP interface operational state in reaction to CFM state

Mutually Exclusive

Mutually Exclusive

Sub second CCM enabled MEPs are supported on SR-7/12 and ESS-7/12 platforms only. The following restrictions apply to tunnel MEPs:

- SF/CPM3 deployments of the 7750 SR-7/12 and 7450 ESS-7/12 support sub second CCM intervals for tunnel MEPs for LAG MEPs and router interface MEPs and port MEPs.
- SF/CPM1 and SF/CPM2 deployments of the 7750 SR-7/12 and 7450 ESS-7/12 support sub second CCM intervals for LAG MEPs and router interface MEPs and port MEPs.

ETH-CFM and MC-LAG

By default, ETH-CFM Management Points (MEPs and MIPs) and MC-LAG operate independently. Alcatel-Lucent recommends not enabling fault propagation when the default behavior is in use. A global command is available in order to allow ETH-CFM the ability to track the state of the MC-LAG for MPs that are configured on MC-LAG ports. This feature does not allow MEPs to influence MC-LAG state. Since the MP relies heavily on the underlying MC-LAG construct, consideration must be given for the proper MC-LAG design and deployment. It is important to understand that the state of MC-LAG can be reflected in the state of the MPs which are configured on SAPs that are part MC-LAGs. For example, a SAP on a LAG that is part of an MC-LAG configuration can behave in a manner that more appropriately represents the MC-LAG.

ETH-CFM and MC-LAG Default Behavior

ETH-CFM MPs track the SAPs, bindings and facility independently. Therefore, when an MP is configured on a SAP which is not operationally `up` because of MC-LAG ETH-CFM defect, conditions are raised for what could be considered normal conditions. Figure 39 shows the default behavior for a point-to-point service without regard for MC-LAG. In the case below, the two up MEPs operating at level 4 on the affected SAPs set the **Interface-Status-TLV** bit in the ETH-CC header to represent the **isDown** condition, assuming ETH-CC is executing between the peer MEPs. This is the correct action based on the ETH-CFM perspective, SAPs are operationally **down**.

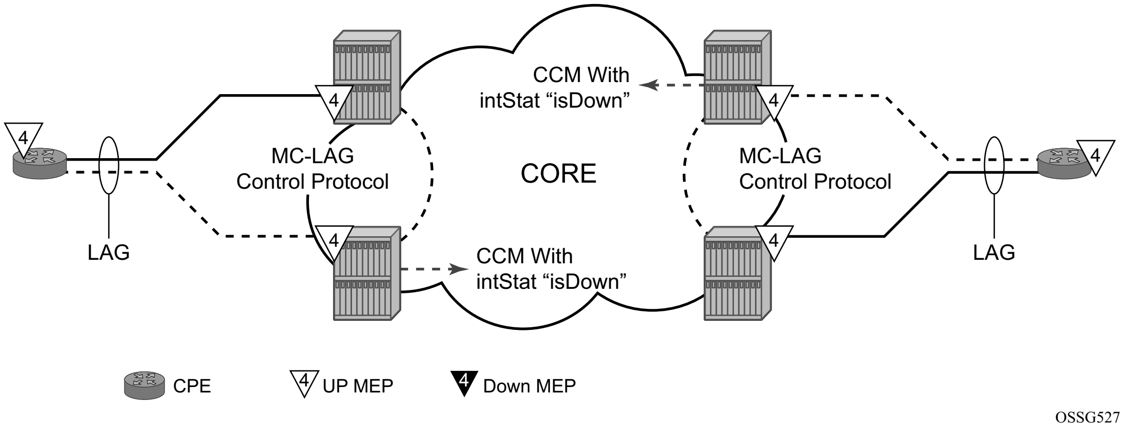

Figure 39: Independent Processing UP MEP Example

A similar condition exists if down MEPs are configured on the SAPs that are operationally `down`. Figure 40 shows how the same service configured with down MEPs would generate AIS, if enabled, toward the remote client at the configured client-meg-level, in the reverse direction of the MEP. This is also the proper behavior from the perspective ETH-CFM.

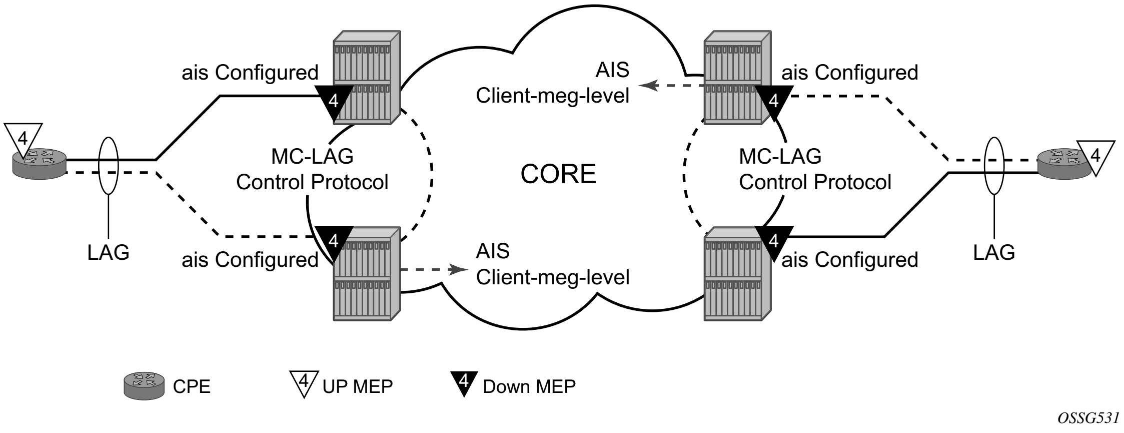

Figure 40: Independent Processing Down MEP Example

Linking ETH-CFM to MC-LAG State

Allowing ETH-CFM to understand the state of MC-LAG and adjust the behavior of the MP (MEP and MIP) according to that state has benefits.

MC-LAG represents the two upstream nodes as a single system to the node terminating a standard LAG. Linking the ETH-CFM MPs to the state of the MC-LAG allows the operator to configure MPs across the two boxes that appear the same. Under the default configuration, this would introduce various defect conditions to be raised and event conditions. However, when ETH-CFM is tracking the state of the MC-LAG, the MPs performs a role that represents the state of the resiliency mechanism. In order to enable this new behavior, configure the system-wide command **standby-mep-shutdown** under the **config>eth-cfm>redundancy>mc-lag** hierarchy.

When a MP is part of the active MC-LAG system, it performs as a normal MP: terminating, generating, responding to, and processing all appropriate ETH-CFM packets. An MP that is on the standby MC-LAG node enters a pseudo-shutdown state. These MPs terminates all ETH-CFM that are part of the regular interception process, but will not process them. They are silently discarded. Also, an MP that exists on a standby MC-LAG system does not generate any ETH-CFM packets. All proactive and on-demand functions are blocked on the standby MC-LAG node. When scheduled tests are executed through SAA these test will attempt to execute. The tests will record failures as a result of the MEP state. These failures are not representative of the network.

This feature relies on the proper configuration, design, and deployment of the MC-LAG protocol. There are numerous optimizations and configuration parameters that are available as part of the

MC-LAG functions. For example, by default, when a currently active MC-LAG port transitions to standby, by any means including manual operator intervention, the remote node terminating the standard LAG sees the LAG transition because all ports in the LAG are down for an instance in time. This is standard LAG behavior does not change as a result of the linkage of MP state to MC-LAG state. This transition causes the propagation of faults for MEPs configured on that node. Normal architectural LAG design must take these types of events into consideration. MC-LAG provides numerous tuning parameters that need to be considered before deploying in the field. These include a **hold-time down** option on the node terminating the standard LAG, as well as other parameters for revertive behavior such as the **hold-time** `up` option. It is important to ensure that the operator's specific environment be taken into consideration when tuning the MC-LAG parameters to avoid the propagation of error conditions during normal recover events. In the case that the resumption of data forwarding exceed the timeout value of a MEP (3.5 times the CCM-Interval), the appropriate defect conditions are raised.

ETH-CFM will register a fault propagation delay timer equal to **propagate-hold-time** under the **config>eth-cfm>redundancy>mc-lag** hierarchy (default of 1s) to delay notification of an event that may be a result of MC-LAG failover. This allows the system time to coordinate events and triggers that together represent the MC-LAG transition from active to standby.

A fixed timer value of 1s will delay an UP MEP from announcing a SAP down condition through CCM Interface-Status-TLV bits, **`isDown`**. ETH-CFM maintains a status of last sent to the UP MEPs peer. When the SAP transitions either to UP or DOWN that fault will be held for the fixed 1s interval and the last Interface-Status-TLV bits will set based on the previous transmission. If the condition, different from the previous sent, still exists at the end of the 1s fixed timer and when the next CCM interval expires, the representative value of the SAP will be sent in the Interface-Status-TLV. These two timers help to smooth out network transitions at the cost of propagation and clearing of faults.

When a node with ETH-CFM linked to MC-LAG is transitioning from standby to active ETH-CFM will assume there are no underlying conditions for any of the SAPs that are now part of the newly activating MC-LAG. The initial notification to an UP MEPs peer will not include any faults. It will assume that the transitioning SAPs are stabilizing as the switchover proceeds. The fixed 1s timer will be starting and a second CCM PDU based on the UP MEPs interval will be sent without any recognition of potential fault on the SAP. However, after the expiration of the fixed timer and on the next CCM-Interval, the Interface-Status-TLV will represent the state of the SAP.

In scaled environments it is important to configure the propagation-hold-time and the CCM intervals to achieve the desired goals. If these timers are set too aggressively, then fault and defect conditions may be generated during times of network stabilization. The use of fault propagation and AIS transmission needs to be carefully considered in environments where MC-LAG protection mechanisms are deployed. Timer values do not guarantee that transitional state will not be propagated to the peer. The propagation of such state may be more taxing and disruptive that allowing the transmission states to complete.

Administrative functions, like **admin down**, are special cases. When the administrative state changes from **up** to **down**, the timer is bypassed and communication from ETH-CFM is immediate.

When an MP is configured in an MC-LAG environment, Alcatel-Lucent recommends that each aspect of the MP be configured the same, including MAC address. Also, although this may be obvious, both nodes participating in the MC-LAG requiring this functionality should include the global command in the **config>eth-cfm>redundancy>mc-lag>standby-mep>shutdown** context to avoid unpredictable behavior.

In summary, a SAP with ETH-CFM tracking the state of the MC-LAG represents the state of the MC-LAG. MPs configured on the standby MC-LAG ports enters a state similar to shutdown. MPs on the MC-LAG ports on the active MC-LAG ports performs all normal processing.

Example: ETH-CFM and MC-LAG Configuration

The following illustration, shows how MEPS can be linked to MC-LAG state. In this example, a service MEP is created on the LAG SAP on NODE1 within service VPLS 100. The MEPs configured on the MC-LAG nodes within service 100 are both configured the same. Both MEPs use the same MEP-ID, the same MAC address.

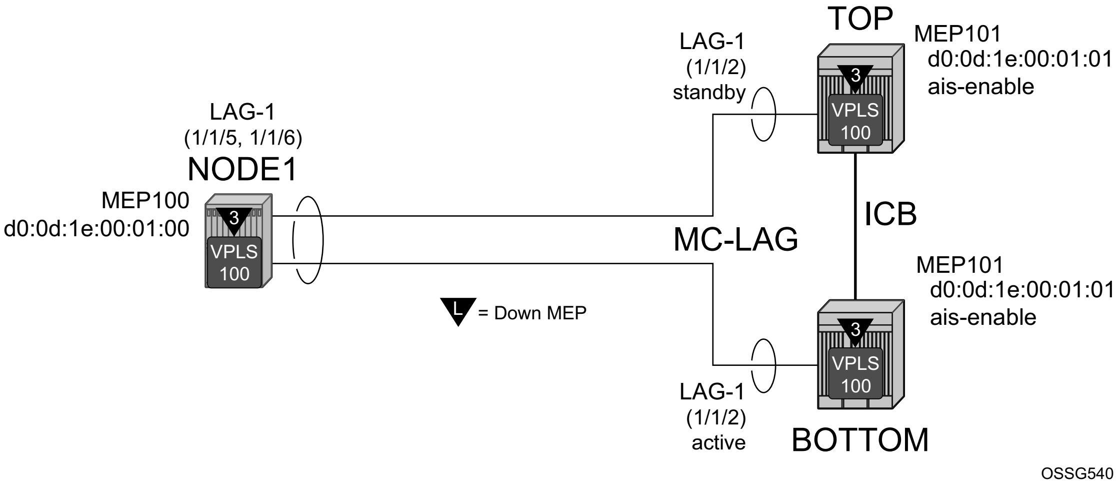

Figure 41: ETH-CFM and MC-LAG Example

Only one of the MEPs on the MC-LAG nodes is active for VPLS service 100. The other MEP is in a shutdown mode, so that even when the MC-LAG is in standby and the port state is **Link Up**, the MEP is in a pseudo shutdown state.

The following configuration example is not meant to provide all possible MC-LAG configuration statement to tune each provider's network. It does provide a base configuration to demonstrate the ETH-CFM feature.

```
NODE1
config>port# info (both ports)
----------------------------------------------
        ethernet
            mode access
            encap-type qinq
            autonegotiate limited
        exit
        no shutdown
----------------------------------------------

config>lag# info
----------------------------------------------
 mode access
        encap-type qinq
        access
            adapt-qos link
        exit
        port 1/1/5
        port 1/1/6
        lacp active administrative-key 32768
        hold-time down 10
        no shutdown
----------------------------------------------

config>eth-cfm# info
----------------------------------------------
        domain 3 format none level 3
            association 1 format icc-based name "03-0000000100"
                bridge-identifier 100
                exit
                ccm-interval 1
                remote-mepid 101
            exit
        exit
----------------------------------------------

config>service>vpls# info
----------------------------------------------
            stp
                shutdown
            exit
            sap 1/1/3:100.100 create
            exit
            sap lag-1:100.100 create
                eth-cfm
                    mep 100 domain 3 association 1 direction down
                        ccm-enable
                        mac-address d0:0d:1e:00:01:00
                        no shutdown
                    exit
                exit
            exit
            no shutdown
----------------------------------------------


TOP (MC-LAG Standby)
config>port# info
```

```
----------------------------------------------
        ethernet
            mode access
            encap-type qinq
            autonegotiate limited
        exit
        no shutdown
----------------------------------------------

config>lag# info
----------------------------------------------
        mode access
        encap-type qinq
        access
            adapt-qos link
        exit
        port 1/1/2
        lacp active administrative-key 32768
        no shutdown
----------------------------------------------

config>router# info
----------------------------------------------
#--------------------------------------------------
echo "IP Configuration"
#--------------------------------------------------
        interface "Core2"
            address 192.168.1.2/30
            port 1/2/2
        exit
        interface "system"
        exit
----------------------------------------------

config>redundancy# info
----------------------------------------------
        multi-chassis
            peer 192.168.1.1 create
                source-address 192.168.1.2
                mc-lag
                    lag 1 lacp-key 1 system-id 00:00:00:00:00:01 system-priority
 100
                    no shutdown
                exit
                no shutdown
            exit
        exit
        synchronize boot-env
----------------------------------------------

config>eth-cfm# info
----------------------------------------------
        domain 3 format none level 3
            association 1 format icc-based name "03-0000000100"
                bridge-identifier 100
                exit
                ccm-interval 1
                remote-mepid 100
            exit
        exit
        redundancy
            mc-lag
                standby-mep-shutdown
            exit
        exit
----------------------------------------------

config>service>vpls# info
----------------------------------------------
            stp
                shutdown
            exit
            sap lag-1:100.100 create
                eth-cfm
                    mep 101 domain 3 association 1 direction down
                        exit
                        ccm-enable
                        mac-address d0:0d:1e:00:01:01
                        no shutdown
                    exit
                exit
            exit
            no shutdown
----------------------------------------------

# show lag 1
===============================================================================
Lag Data
===============================================================================
Lag-id         Adm     Opr     Port-Threshold   Up-Link-Count   MC Act/Stdby
-------------------------------------------------------------------------------
1              up      down    0                0               standby
===============================================================================

# show port
===============================================================================
Ports on Slot 1
===============================================================================
Port        Admin Link Port    Cfg  Oper LAG/ Port Port Port   C/QS/S/XFP/
Id          State      State   MTU  MTU  Bndl Mode Encp Type   MDIMDX
-------------------------------------------------------------------------------
… snip …
1/1/2       Up    Yes  Link Up 1522 1522    1 accs qinq xcme
…snip…
==========================================================================


BOT (MC-LAG Active)
config>port# info
----------------------------------------------
        ethernet
            mode access
            encap-type qinq
            autonegotiate limited
        exit
        no shutdown
----------------------------------------------
```

```
config>lag# info
----------------------------------------------
        mode access
        encap-type qinq
        access
            adapt-qos link
        exit
        port 1/1/2
        lacp active administrative-key 32768
        no shutdown
----------------------------------------------

config>router# info
----------------------------------------------
#--------------------------------------------------
echo "IP Configuration"
#--------------------------------------------------
        interface "Core1"
            address 192.168.1.1/30
            port 1/2/1
        exit
        interface "system"
        exit
----------------------------------------------

config>redundancy# info
----------------------------------------------
        multi-chassis
            peer 192.168.1.2 create
                source-address 192.168.1.1
                mc-lag
                    lag 1 lacp-key 1 system-id 00:00:00:00:00:01 system-priority
 100
                    no shutdown
                exit
                no shutdown
            exit
        exit
        synchronize boot-env
----------------------------------------------

config>eth-cfm# info
----------------------------------------------
        domain 3 format none level 3
            association 1 format icc-based name "03-0000000100"
                bridge-identifier 100
                exit
                ccm-interval 1
                remote-mepid 100
            exit
        exit
        redundancy
            mc-lag
                standby-mep-shutdown
            exit
        exit
----------------------------------------------

config>service>vpls# info
```

```
----------------------------------------------
            stp
                shutdown
            exit
            sap lag-1:100.100 create
                eth-cfm
                    mep 101 domain 3 association 1 direction down
                        exit
                        ccm-enable
                        mac-address d0:0d:1e:00:01:01
                        no shutdown
                    exit
                exit
            exit
            no shutdown
----------------------------------------------

# show lag 1
===============================================================================
Lag Data
===============================================================================
Lag-id         Adm     Opr     Port-Threshold   Up-Link-Count   MC Act/Stdby
-------------------------------------------------------------------------------
1              up      up      0                1               active
===============================================================================

# show port
===============================================================================
Ports on Slot 1
===============================================================================
Port        Admin Link Port    Cfg  Oper LAG/ Port Port Port   C/QS/S/XFP/
Id          State      State   MTU  MTU  Bndl Mode Encp Type   MDIMDX
-------------------------------------------------------------------------------
…snip…
1/1/2       Up    Yes  Up      1522 1522    1 accs qinq xcme
…snip…

===============================================================================
```

ETH-CFM Features

CCM Hold Timers

In some cases the requirement exists to prevent a MEP from entering the defRemoteCCM defect, remote peer timeout, from more time than the standard 3.5 times the CCM-interval. Both the IEEE 802.1ag standard and ITU-T Y.1731 recommendation provide a non-configurable 3.5 times the CCM interval to determine a peer time out. However, when sub second CCM timers (10ms/ 100ms) are enabled the carrier may want to provide additional time for different network segments to converge before declaring a peer lost because of a timeout. In order to maintain compliance with the specifications the `ccm-hold-timer down <delay-down>` option has been introduced to artificially increase the amount of time it takes for a MEP to enter a failed state should the peer time out. This timer is only additive to CCM timeout conditions. All other CCM defect conditions, like defMACStatus, defXconCCM, and so on, will maintain their existing behavior of transitioning the MEP to a failed state and raising the proper defect condition without delay.

When the **ccm-hold-timer down** *delay-down* option is configured the following calculation is used to determine the remote peer time out (3.5 times the CCM-Interval + ccm-hold-timer delay-down).

This command is configured under the association. Only sub second CCM enabled MEPs support this hold timer. Ethernet-Tunnel Paths use a similar but slightly different approach and will continue to utilize the existing method. Ethernet-tunnels will be blocked from using this new hold timer.

It is possible to change this command on the fly without deleting it first. Simply entering the command with the new values will change to values without having to delete the command prior to the change.

It is possible to change the ccm-interval of a MEP on the fly without first deleting it. This means it is possible to change a sub second CCM enabled MEP to 1 second or above. The operator will be prevented from changing an association from a sub second CCM interval to a non-sub second CCM interval when a `ccm-hold-timer` is configured in that association. The `ccm-hold-timer` must be removed using the `no` option prior to allowing the transition from sub second to non-sub second CCM interval.

CCN Interval

This feature is an enhancement that enables slow timers OAM handling of G.8032 enabling G.8032 on the 7750 C4/C12 and ESS6 platforms. G.8032 uses the OAM for Ring Protection messages. This feature enables full G.8032 Ring support on these platforms. In addition, this feature enables Continuity Check messages (CCM) on Ring ports at 1 second intervals for all platforms where G.8032 is supported. With this feature, G.8032 can be configured on additional 7x50 platforms. CCM are optional with G.8032 but normally deployed for higher assurance of protection. The SR 7750 and ESS 7450 additionally support CCM of 100ms and 10ms. CCM is configured on a neighbor node basis so the only requirement is that neighbor switches be configured with same interval or with CCM disabled.

Note[1]: Ethernet-Tunnels and Ethernet-Rings are not configurable under all service types. Any service restrictions for MEP direction or MIP support will override the generic capability of the Ethernet-Tunnel or Ethernet-Ring MPs. Check the applicable user guide for applicability.

A Virtual MEP (vMEP) is a MEP that is configured at the service level rather than on a SAP or SDP binding. A vMEP sends ETH-CFM to all the SAPs and SDP bindings in the VPLS, depending on the type of traffic. If it is multicast traffic, the packets forward out all SAPs and SDP bindings. Unicast traffic is forwarded appropriately based on the type of ETH-CFM packet and the forwarding tables. Packets inbound to a context containing a vMEP performs normal processing and forwarding through the data plane with a copying of the ETH-CFM packet delivered to the local MEP for the appropriate levels. The local MEP will determine whether or not it should process a copied inbound ETH-CFM frame acting in accordance with standard rules.

Configuring a vMEP is similar in concept to placing down MEPs on the individual SAPs and SDP bindings in the associated VPLS. This ensures that packets inbound to the service get redirected to the vMEP for processing. Proper domain nesting must be followed in order to avoid ETH-CFM error conditions.

vMEPs have been expanded to include VPLS, m-VPLS, and I-VPLS contexts.   The original B-VPLS vMEP remains supported within that context and maintain the original restrictions (no MIPs and only in a B-VPLS context). A vMEP in a B-VPLS context should be migrated to support the enhancements by adding the “vmep-extensions” command, if the hardware requirements are met.   The vmep-extensions command is disabled by default for any vMEP configured within a B-VPLS context. This ensures backwards compatibility and does not impose any new hardware requirements for existing vMEPs in B-VPLS contexts. The “vmep-extensions” command is in effect by default and cannot be negated for any other supported VPLS context, meaning these VPLS contexts must meet explicit hardware requirements.

A vMEP in an I-VPLS context can only extract packets inbound on local SAP and SDP bindings. This extraction does not include packets that are mapped to the I-VPLS from associated B-VPLS context. If this type of extraction is required in an I-VPLS context then UP MEPs are required on appropriate SAPs and SDP bindings in the I-VPLS service.

The enhanced functionality and wider scope of this feature requires all the SAPs within the service and every network port on the node to be IOM3 and higher hardware. When the operator attempts to configure a vMEP in an instance that does not meet the hardware requirements the configuration will be rejected. The only exception to this is a vMEP configured within a B-VPLS context. However, if an attempt is made to transition that vMEP using the **vmep-extensions** command the action will be rejected.

As with the original vMEP functionality introduced for B-VPLS contexts, DOWN MEPs are supported on the individual SAPs or SDP bindings as long as domain nesting rules are not violated. Of course, local UP MEPs are only supported at the same level as the vMEP otherwise various CCM defect conditions will be raised, assuming CCM is enabled, and leaking of ETH-CFM packets will occur (lower level ETH-CFM packets arriving on a lower level MEP). Domain nesting must be properly deployed to avoid unexpected defect conditions and leaking between ETH-CFM domains.

The vMEP enhancements increase scalability, allow for MIPs and include an optional **vmep-filter**.

MIPs map be configured on the SAPs and SDP-Spokes at or above level of the vMEP.

An optional **vmep-filter** provides a coarse means of silently dropping all ETH-CFM packets that would normally be redirected to the CPU following egress processing. These includes any ETH-CFM level equal to or lower than the vMEP and any level equal to and lower than any other Management Points on the same SAP or SDP binding that includes the **vmep-filter**. MIPs will automatically be deleted when they coexist on the same SAP or spoke-sdp as the **vmep-filter**. Since DOWN MEPs are ingress processed they are supported in combination with a vMEP and operate normally regardless of any **vmep-filter**. Domain nesting rules must be adhered to.

If the operator requires an MP on the SAP or SDP binding an UP MEP may be created at the same level as the vMEP on the appropriate SAP or SDP binding to perform the same function as the filter but at the specific level of the MEP. Scalability needs to be clearly understood because this will redirect the ETH-CFM packets to the CPU (consider using CPU protection introduced in release 8.0r5). Consideration must also be given to the impact this approach could have on the total number of MEPs required. There are a number of other approaches that may lend themselves to the specific network architecture.

vMEP filtering is not supported within the a PBB VPLS since it already provides separation between B-components (typically the core) and I-components (typically the customer)

vMEPs do not support any ETH-AIS functionality and do not support fault propagation functions.

Below is a sample configuration that shows how to configure a vMEP in a VPLS context.

```
config>service# vpls 100 customer 1 create

config>service>vpls$ info
```

```
----------------------------------------------
  stp
 shutdown
  exit
    eth-cfm
     mep 100 domain 3 association 1
         mac-address d0:0d:1e:00:01:11
   ccm-enable
         no shutdown
      exit
  exit
  no shutdown
----------------------------------------------
```

Different service types support different ETH-CFM functionality. This is explained in the applicable service sections throughout this manual.

Note: To use any Y.1731 specific function, the domain must be configured with a domain format of “none”. This includes the MEPs that are created as part of the G.8031 and G.8032 protection scheme. That is because they use ETH-APS as defined in the ITU-T Y.1731 recommendation and are not part of the IEEE 802.1ag specification.

Service Management Tasks

This section discusses the following service management tasks:

- Modifying Customer Accounts on page 166
- Deleting Customers on page 168
- Modifying SDPs on page 169
- Deleting SDPs on page 170

Modifying Customer Accounts

To access a specific customer account, you must specify the customer ID.
To display a list of customer IDs, use the show service customer command.
Enter the parameter (description, contact, phone) and then enter the new information.

CLI Syntax:

```
config>service# customer customer-id create
 [no] contact contact-information
 [no] description description-string
 [no] multi-service-site customer-site-name [create]
    assignment {port port-id | card slot}
    no assignment
    [no] description description-string
    egress
       agg-rate-limit agg-rate [queue-frame-based-accounting]
       no agg-rate-limit
       policer-control-policy name
       no policer-control-policy
       scheduler-policy scheduler-policy-name
       no scheduler-policy
       [no] scheduler-override
          scheduler scheduler-name [create]
          no scheduler scheduler-name
    ingress
       policer-control-policy name
       no policer-control-policy
       scheduler-policy scheduler-policy-name
       no scheduler-policy
       [no] scheduler-override
          scheduler scheduler-name [create]
          no scheduler scheduler-name
    tod-suite tod-suite-name
 [no] phone phone-number
```

Example:

```
config>service#  customer 27 create
config>service>customer$ description “Western Division”
config>service>customer# contact “John Dough”
config>service>customer# no phone “(650) 237-5102”
```

Deleting Customers

The no form of the customer command removes a customer ID and all associated information. All service references to the customer must be shut down and deleted before a customer account can be deleted.

CLI Syntax: `config>service# no customer customer-id`

Example:

```
config>service# epipe 5 customer 27 shutdown
config>service# epipe 9 customer 27 shutdown
config>service# no epipe 5
config>service# no epipe 9
config>service# no customer 27
```

Modifying SDPs

To access a specific SDP, you must specify the SDP ID. To display a list of SDPs, use the show service sdp command. Enter the parameter, such as description, far-end, and lsp, and then enter the new information.

NOTE: Once created, you cannot modify the SDP encapsulation type.

CLI Syntax: `config>service# sdp sdp-id`

Example:

```
config>service# sdp 79
config>service>sdp# description “Path-to-107”
config>service>sdp# shutdown
config>service>sdp# far-end “10.10.10.107”
config>service>sdp# path-mtu 1503
config>service>sdp# no shutdown
```

Deleting SDPs

The no form of the **sdp** command removes an SDP ID and all associated information. Before an SDP can be deleted, the SDP must be shutdown and removed (unbound) from all customer services where it is applied.

CLI Syntax:

```
config>service# no sdp 79
```

Example:

```
config>service# epipe 5 spoke-sdp 79:5
config>service>epipe>sdp# shutdown
config>service>epipe>sdp# exit
config>service>epipe# exit
config>service# no sdp 79
```