IEEE 1904.2 Universal Management Tunnel

Principles and Design

TiBiT communications inc

Vocabulary

- Universal Management Tunnel (UMT) –
  - The protocol defined by IEEE 1904.2;
  - also refers to an instance of the protocol operating between two implementations of the IEEE 1904.2 protocol
- UMT Protocol Data Unit (UMTPDU) – The unit of UMT data sent across the network
- L2 Domain – A term describing a network that forwards Ethernet frames based solely on MAC address
- Service Data Unit (SDU) – The unit of data carried as payload in service-providing protocol (inferior layer in a stack) for a client protocol (superior layer in a stack)
- Protocol Data Unit (PDU) – The unit of data for a service-providing protocol
- Intermediate Node – A layer-2 switch or bridge that lies in the transit path between two UMT peer
- UMT Peer – A station that implements the UMT stack
- Station – As defined in IEEE 802 (end station) is a source and/or destination of link layer traffic

TiBiT communications inc

Need for the Protocol

- Provide a mechanism to transmit SDUs across a layer-2 domain for protocols that would otherwise not be forwarded due to addressing conflicts or other factors.
  - Scope-Limited Protocols: OAM
  - Non-Native or Virtualized Protocols: OMCI, IGMP, ARP
- And to transport these protocols in the absence of, and without the overhead associated with, Layer-3 protocols and tunnels.

Example: Scope Limited Protocols

- Ethernet slow protocols are specified for transport across a single link: OAM, CDP, STP, UDLD, LACP

Original Use Case: EPOC

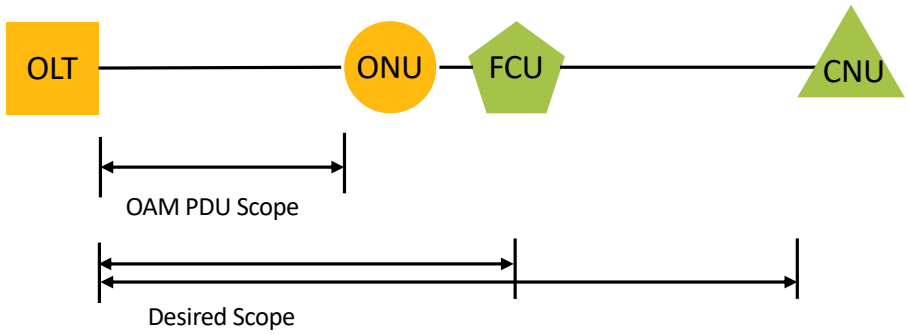

Emerging Use Case: NFV for EPON

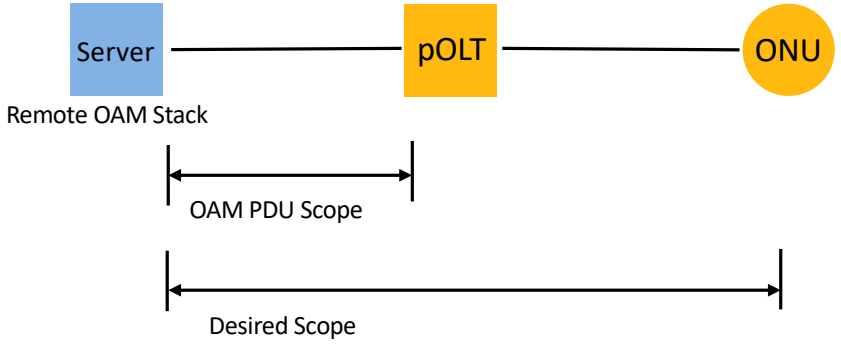

TiBiT communications inc

Example: Non-Native Protocols

Use Case: NFV for GPON

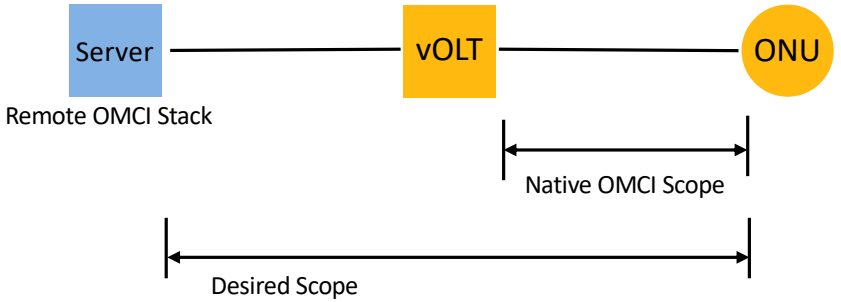

- OMCI operation is defined between an ITU PON OLT and ONU over the OMCC
- OMCI is not a native Ethernet protocol and, like eOAM/SIEPON, is specified for a single link
- OMCI encapsulation in Ethernet frames is defined (to use OMCI in EPON), but transport across an L2 network is not defined

TiBiT communications inc

Example: Transport for Virtualized Protocols

Use Case: NFV for IGMP snooping

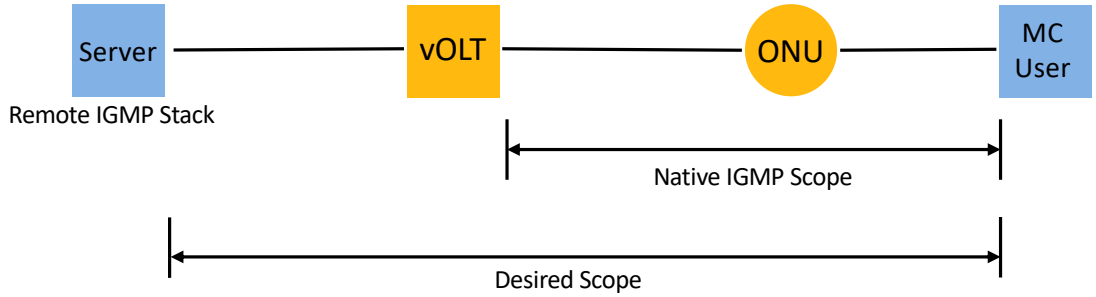

- IGMP Snooping is traditionally performed in the local switch
- Virtualized switches move the IGMP processing to a remote system
- In the SDN/NFV use case, IGMP needs to be transported to the remote system

TiBiT communications inc

What are the requirements?

- Bridge/Switch Traversal:
  - Must transparently transit a Layer-2 bridge/switch
- Multicasting and broadcasting:
  - Enable client protocols to communicate with multiple peers simultaneously
- Extensible:
  - The protocol must be able to adapt to carry new client protocols without significant redesign
  - Should align with IEEE SA policy for Ethertype assignments
- Lightweight:
  - Must not carry lots of overhead as a requirement; stateless; peer-to-peer
  - Must not preclude additional functionality – peer discovery and automatic configuration
- Must operate without L3:
  - Many of the target devices do not support IP
- Require zero modification of existing bridges and switches that are in the transit path
  - Enhanced operations can be enabled if the transit/intermediate bridges/switches are UMT-aware (optional)
- Must be able to emulate a point-to-point link
- Must be independent of the PHY
- Must support encapsulation in a VLAN

TiBiT communications inc

Solutions with Existing Protocols?

- MAC-in-MAC
  - Could work for native Ethernet protocols like eOAM or even IGMP
  - Doesn't address non-Ethernet protocols like OMCI, IGMP
    - would require a new Ethertype and PDU description

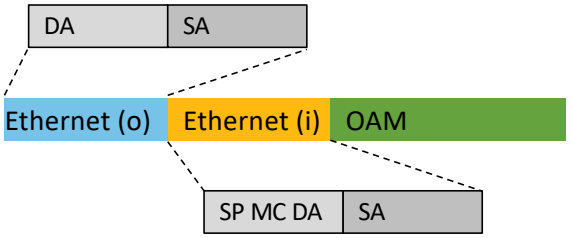

- Layer-2 Protocol Tunneling
  - Designed for Layer-2 Protocols that use special DA
    - like OAM, CDP, VTP, LACP, UDLD, STP, etc
  - Relies on flooding and proprietary multicast address
  - Doesn't address non-Ethernet protocols like OMCI, IGMP
  - Not standardized (originally defined by Cisco), interop not defined

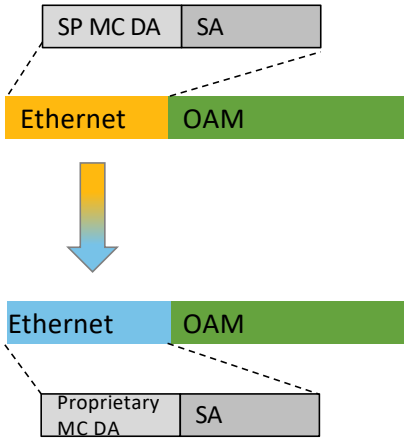

TiBiT communications inc

Generalized Use Cases

- Point-to-Point over L2 Domain with UMT-aware peer
- Point-to-Multipoint over L2 Domain with UMT-aware peer
- PtP or PtMP with non-UMT aware peer(s)

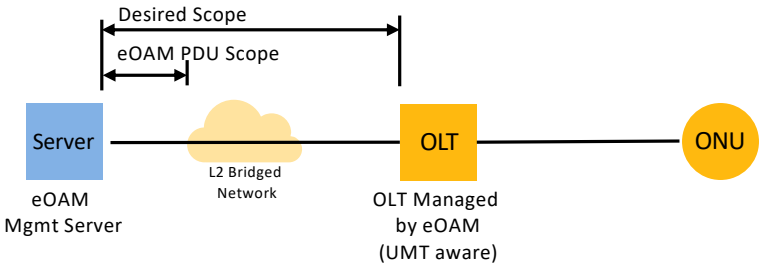

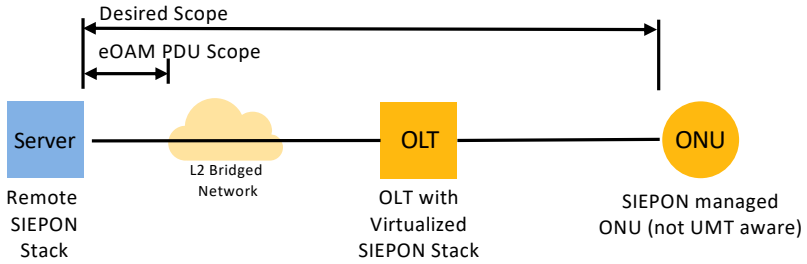

TiBiT communications inc

Another Example Use Case

- Point-to-Multipoint + UMT Client using Multicast

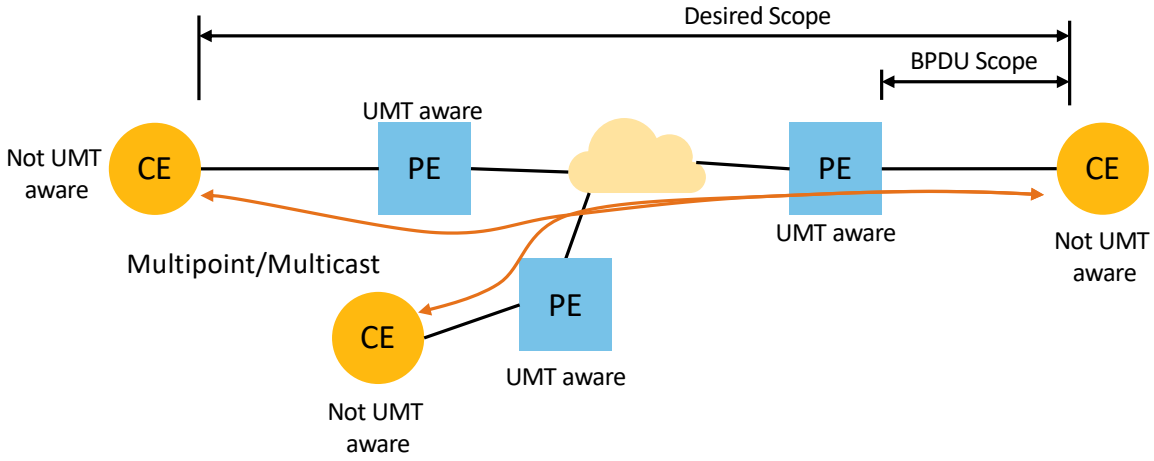

TiBiT communications inc

A specific Use Case

TiBiT communications inc

Let's Consider a Use Case: *OAM over UMT*

- OAM is defined in IEEE 802.3 Clause 57
- Intended to Operate on Point-to-Point Links
- Uses Slow Protocols Multicast DA which is link-local only
  - filtered by 802.1Q compliant switches/bridges

By examining this specific use case, we can

gain insight into the more general use case for UMT

TiBiT communications inc

OAM over UMT: The most general scenario

- OAM Stack is not co-located with the UMT stack

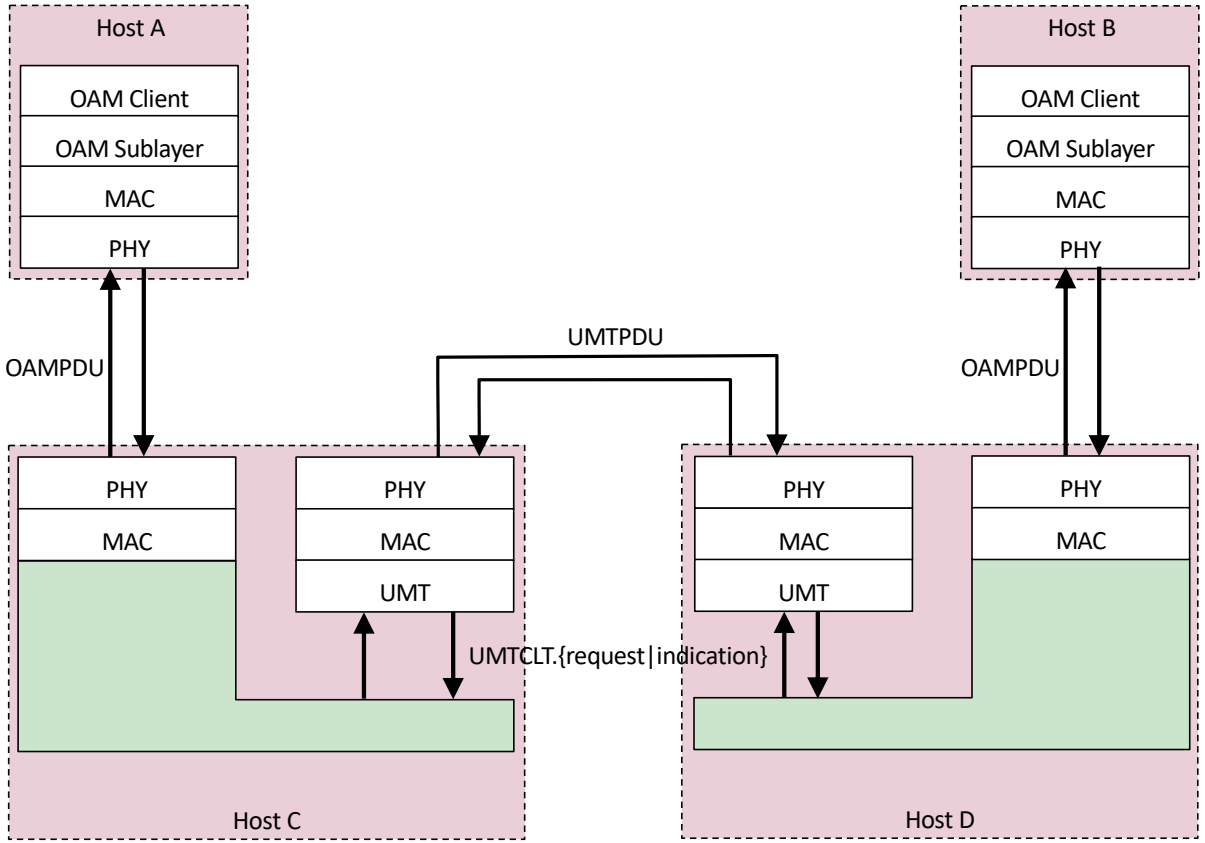

A Day in the Life of the OAMPDU over UMT… (1/2)

1. OAM Peer (A) transmits OAMPDU (SA=host A MAC, DA= 01-80-c2-00-00-02, type=0x8809)
2. OAMPDU is received by UMT host C
3. OAMPDU is relayed inside Host C to the UMT Layer
4. UMT Layer forms the UMTPDU
5. UMT Layer forwards UMTPDU to MAC and on to the PHY
6. UMTPDU transits the Layer-2 network and is received by Host D

TiBiT communications inc

A Day in the Life of the OAMPDU over UMT… (2/2)

7. UMTPDU transits the Layer-2 network and is received by Host D
8. MAC forwards UMTPDU to UMT Layer
9. UMT Layer inspects UMTPDU to determine destination UMT Client
10. OAMSDU is relayed to Host D's MAC
11. OAMPDU is forwarded through Host D's MAC (SA=host A MAC, DA= 01-80-c2-00-00-02, type=0x8809)
12. OAMPDU is received by OAM Peer (B)

What's Missing?

- What is the format of a UMTPDU?
- How is the OAM encapsulated in a UMTPDU?

- How does UMT Peer C know about UMT Peer D?
- How do UMT Peer C and UMT Peer D know that OAM is to be transported?
- Can there be multiple tunnels on a UMT Peer?
- Can there be multiple tunnels between the same two UMT Peers?
- What defines a tunnel? Is it simply the UMT peer-to-peer, or is each UMT client carried in its own tunnel?

- What, in the UMT peers, relays the OAMPDU to the UMT layer?
- What if there are multiple tunnels on host C? How does it know which tunnel to send the OAM through?

UMT Stack and UMTPDU Format

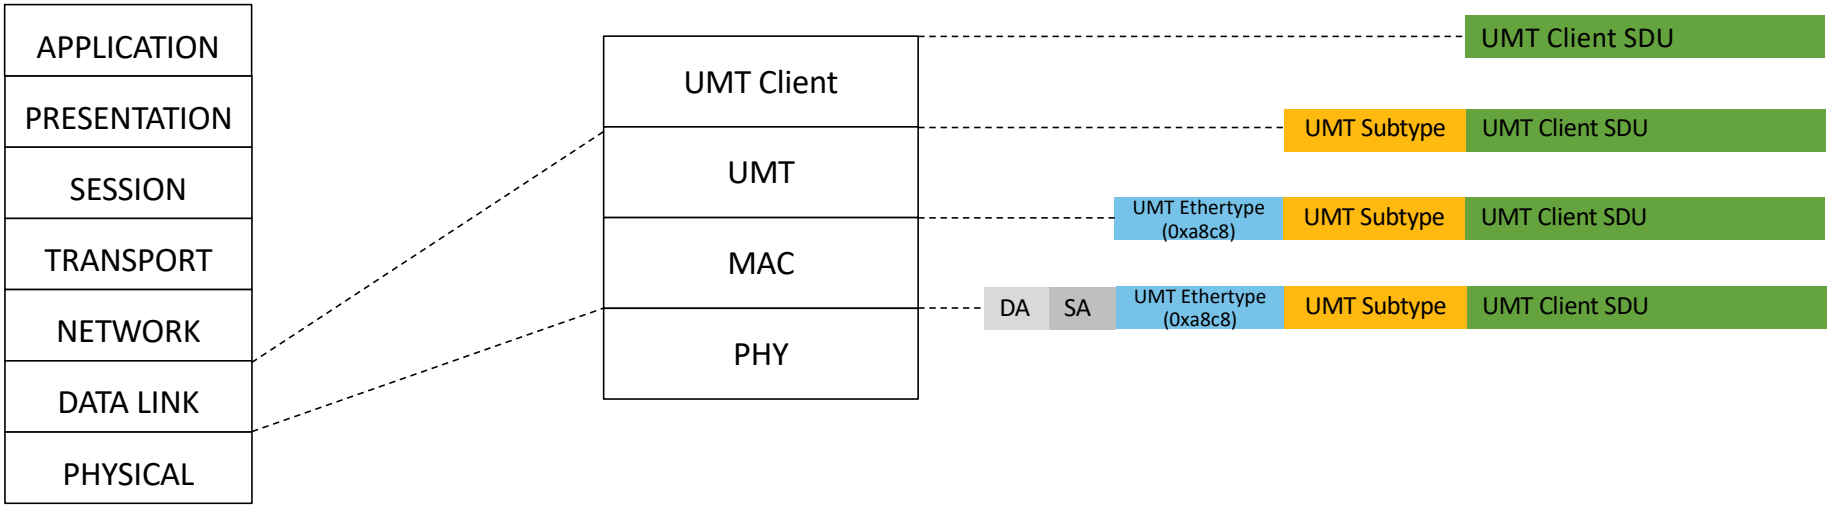

- UMT sits just above the MAC
  - Enables UMT to be visible to many client protocols
- UMT applies a new Ethertype (0xa8c8) to identify it at the MAC layer

- UMT defines a subtype
  - Aligns with IEEE SA policy for extensibility
  - Do we need a version number?
  - Subtype identifies the client protocol

TiBiT communications inc

OAM in a UMTPDU

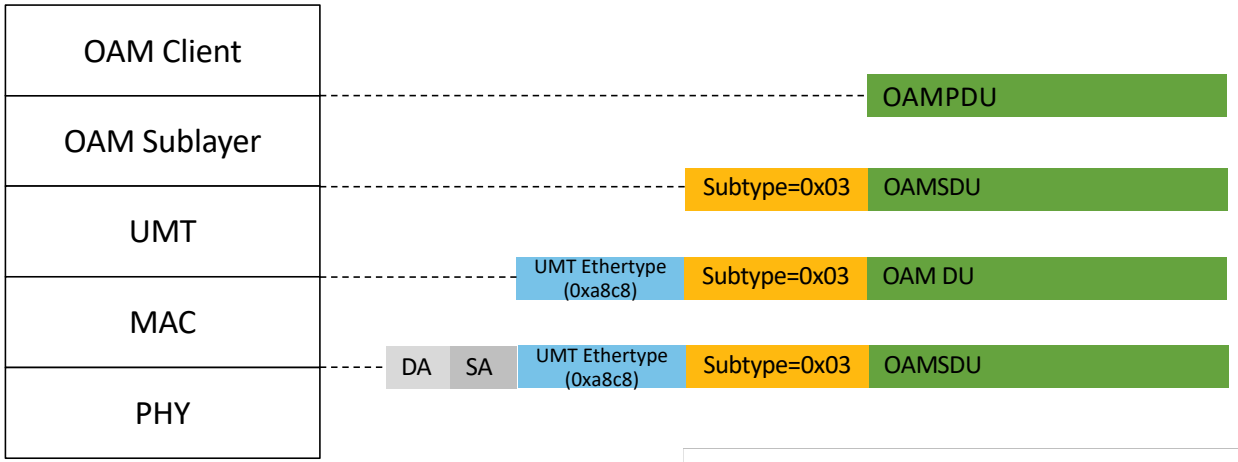

- OAMPDU: Is defined in IEEE 802.3 Clause 57
- OAMSDU: OAMPDU without the SA, DA, Slow Protos Subtype, Ethertype, and FCS

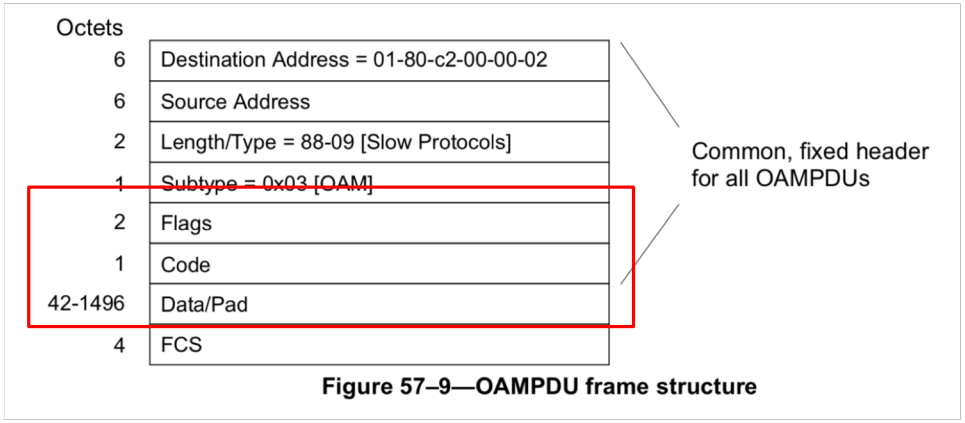

Figure 57–9—OAMPDU frame structure

TiBiT communications inc

What's Missing?

- ✓ What is the format of a UMTPDU?
- ✓ How is the OAM encapsulated in a UMTPDU?

- How does UMT Peer C know about UMT Peer D?
- How do UMT Peer C and UMT Peer D know that OAM is to be transported?
- Can there be multiple tunnels on a UMT Peer?
- Can there be multiple tunnels between the same two UMT Peers?
- What defines a tunnel? Is it simply the UMT peer-to-peer, or is each UMT client carried in its own tunnel?

- What, in the UMT peers, relays the OAMPDU to the UMT layer?
- What if there are multiple tunnels on host C? How does it know which tunnel to send the OAM through?

UMT Peer Discovery and Tunnel Configuration

- UMT Peers need to
  - Know the existence of one another
  - Know that a tunnel* is expected between UMT peers in a given set
  - Know which protocols are expected in a given tunnel

* These questions raise additional questions about the fundamentals of UMT…

Fundamental Questions:

- What is a UMT Peer?
  - A single instance of the UMT Layer on a given station
- What defines a tunnel?
  - Two UMT peers agreeing to exchange UMTPDUs
  - What constitutes agreement?
    - A Matching configuration on the peers: Client protocol(s), SA, DA
- Is more than one tunnel allowed or needed between a given set of peers?
  - Yes – The desired SA/DA could be different. This is especially important for multicasting UMTPDUs
  - Alternatively, a UMT peer could be limited to a single tunnel

TiBiT communications inc

Multiple Tunnels: Single Stack or Multiple Stacks?

Multiple Tunnels per UMT Layer?

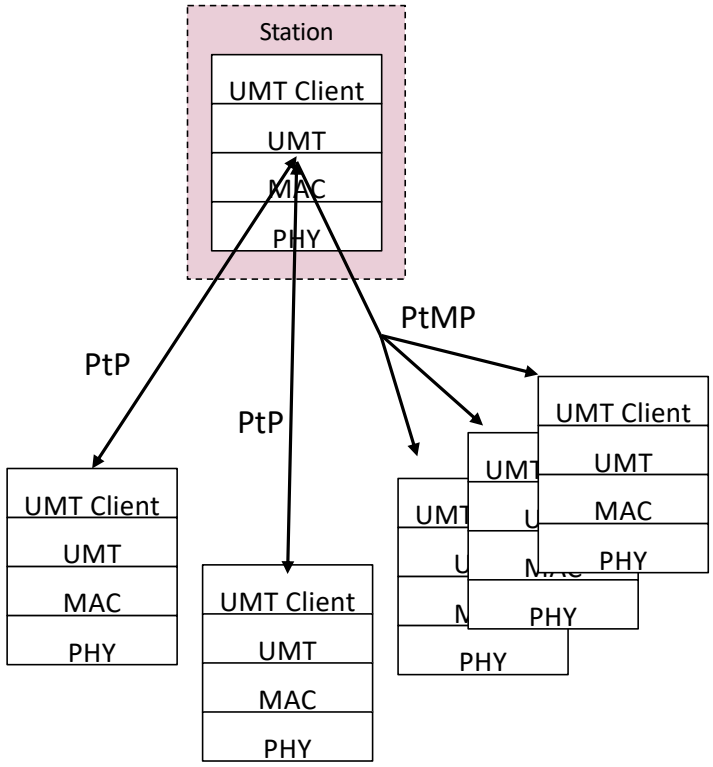

Single Tunnel Per UMT Layer?

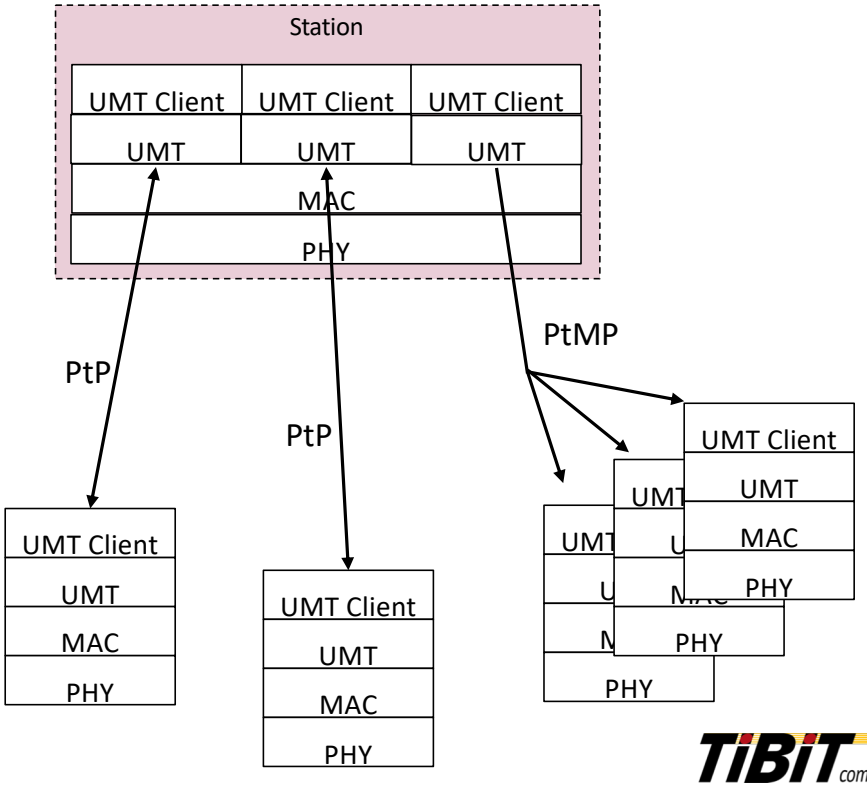

TiBiT communications inc

UMT Peer Discovery and Tunnel Configuration

How does UMT Peer C know about UMT Peer D?

How do UMT Peer C and UMT Peer D know that OAM is to be transported?

- In the proposed base standard, UMT peers are aware of one another and agree on transported protocols via an unspecified mechanism.
  - Manual configuration by an administrator

- The base standard must leave room to specify an automatic method at a later time or in an annex/appendix/amendment

TiBiT communications inc

What's Missing?

- ✓ What is the format of a UMTPDU?
- ✓ How is the OAM encapsulated in a UMTPDU?

- ✓ How does UMT Peer C know about UMT Peer D?
- ✓ How do UMT Peer C and UMT Peer D know that OAM is to be transported?
- ✓ Can there be multiple tunnels on a UMT Peer?
- ✓ Can there be multiple tunnels between the same two UMT Peers?
- ✓ What defines a tunnel? Is it simply the UMT peer-to-peer, or is each UMT client carried in its own tunnel?

- What, in the UMT peers, relays the OAMPDU to the UMT layer?
- What if there are multiple tunnels on host C? How does it know which tunnel to send the OAM through?

TiBiT communications inc

Relaying to the UMT Layer

- What is this relay function?
  - A function that is defined for each client protocol – An adaptation layer
  - Performs necessary transformations on the client protocol to allow it to use the UMT Layer
- How should it be defined/specified?
  - An annex to the base standard
  - The base standard describes an adaptation layer in abstract form, leaving the specifics to the annex

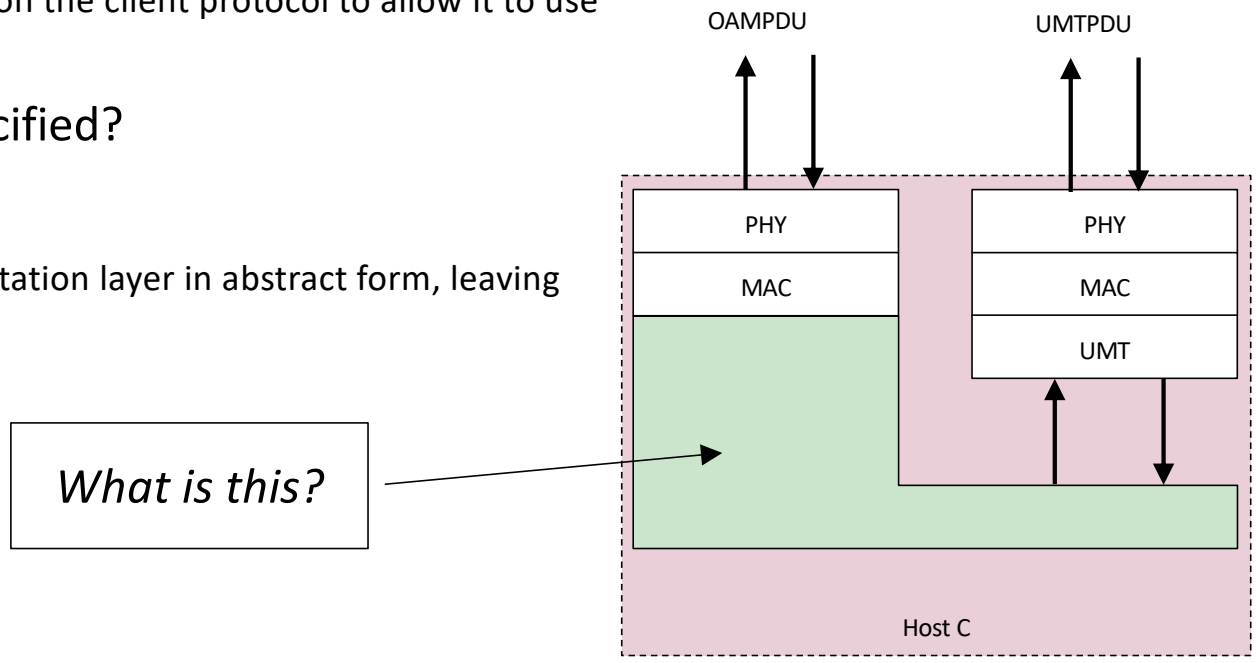

TiBiT communications inc

Which Tunnel to Choose?

1. IF: Multiple Tunnels per UMT Layer
   - The UMT Layer needs to present an interface that allows the client to choose a tunnel

2. IF: Single Tunnel per UMT Layer
   - The client needs to choose the tunnel by connecting to the correct UMT Layer

Option 2 is simpler to describe

TiBiT communications inc

What's Missing?

- ✓ What is the format of a UMTPDU?
- ✓ How is the OAM encapsulated in a UMTPDU?

- ✓ How does UMT Peer C know about UMT Peer D?
- ✓ How do UMT Peer C and UMT Peer D know that OAM is to be transported?
- ✓ Can there be multiple tunnels on a UMT Peer?
- ✓ Can there be multiple tunnels between the same two UMT Peers?
- ✓ What defines a tunnel? Is it simply the UMT peer-to-peer, or is each UMT client carried in its own tunnel?

- ✓ What, in the UMT peers, relays the OAMPDU to the UMT layer?
- ✓ What if there are multiple tunnels on host C? How does it know which tunnel to send the OAM through?

Key Principles to Take Away (1/2)

Based on the OAM Use Case…

- We know the UMTPDU format
- We know how an example client would be encapsulated
- We know that there is a need for an Adaptation Layer that will perform the necessary transformations on the client protocol
- We know that UMT Peers must be aware of one another
  - Base Standard assumes an external method (manual or other TBD)
- We know that UMT Peers must agree on tunnel configuration (which protocols to carry)
  - Base Standard assumes an external method (manual or other TBD)
- We assume that each instance of a UMT Layer represents a single tunnel
- We define a tunnel by {Source UMT Peer MAC Address, Destination UMT Peer MAC Address} where the DA could be Unicast, Multicast, or Broadcast

TiBiT communications inc

Key Principles to Take Away (2/2)

Based on the OAM Use Case…

- We assume that it is up to the UMT Client to connect to the desired UMT Layer (tunnel)
  - Method for “choosing” and “connecting” is unspecified in the base standard
  - Could be manually configured or some TBD method
- We assume that interoperability is important, therefore the standard must define how each UMT Client protocol
- We need to determine the structure of the standard to allow extensions to be added easily
  - Amendment, Annex, Appendix?

TiBiT communications inc

Thank You!
Additional Q&A

TiBiT communications inc