IEEE P802.3
ETHERNET IN THE FIRST MILE
JANUARY 8,9 2001

David Closs

david_closs@adc.com
203-639-7617

Threshold moment for Ethernet

- Ethernet is emerging as an end-to-end standard
- FSAN/ITU PON standard is 10 yrs old, ATM based, only 622Mbps, needs to be renovated
- Copper infrastructure giving ground to fiber
- Companies will start providing fiber-based Ethernet to the home, market penetration will be hampered by fragmentary approaches

As far as standards go, Ethernet has probably had more effect on peoples lives than any other standard save for, perhaps, the standards that were wrought to bring telephone service and television into peoples homes. I believe that just through the sheer momentum of its installed base Ethernet will go on to eclipse those two familiar mediums as it supplants those two technologies, and then renders them obsolete as it surely will.

Our most pressing problem at the moment is getting that Ethernet data to and from the home in a fashion that won't require infrastructure reconstruction every few years as bandwidth needs steadily increase, as high-bandwidth Ethernet access to the home changes from a "premium service" to a service providing basic needs. (At this time, premium service can be used to refer to anything over a megabit-per-second of guaranteed usable, symmetric, full-duplex bandwidth.)

Ironic, that last bit: What would be a premium service today for data transport to the home is not fit, from the bandwidth perspective, to provide the basic telephone and television services of today. But that is what we must keep in mind during this journey down the standards process, the essential kernel of truth that will decide if this standard will become as ubiquitous as the standard it seeks to augment. What we are seeking to do is to turn the equation around, to stop shoving as many bits as possible down an existing copper infrastructure, and to start to plan and build an infrastructure that will deliver all our services down one pipe, over one medium, over one protocol.

That infrastructure may initially terminate in copper for the last mile, and later from the curb onwards, but it will surely become fiber all the way to the home as coaxial-cable based systems share their bandwidth across more homes, and as DSL service depletes the number of usable twisted pairs through actually providing high-speed services to the home, and polluting other pairs with crosstalk.

Goals: ***what are we looking for?***

- Provide standard for a fiber-to-the-home and fiber-to-the-curb based delivery mechanism
- Use the simplest protocol that meets flexibility/upgrade needs
- Allow for simplest CPE equipment that meets the need
- Allow for initial deployment with existing technology
- Right now need copper-terminated variants
- FTTH/FTTC will be much closer by the time a standard is ratified

A broad scope is required

- Many different industries involved (CATV, large carriers, CLECs, etc.)

It is important to remember that the companies that will be buying equipment based on this standard will have different preferred ways of architecting, building, and maintaining their networks. Flexibility is the key, and that needs to be reflected in the standard.

- Delivery mechanism, format of data on the medium

- Provide for non-VOIP (legacy) telecom service

This could have a beneficial effect on industry acceptance of the standard: the equipment operators can provide phone service over their new network before VOIP is widespread or popular, and not have to change anything except a plug card on the CPE to change over to full VOIP.

It also allows the overbuilder market to compete not only for cable and data services, but also for POTS, with the added benefit of not needing access to the incumbent's equipment.

- Codify the philosophy behind the decisions made on the protocol, the equipment, and the features as seen by the service provider and the end user.

- Standardize the methodology behind upgrading provided services via scalability requirements

This is one of my key concerns. I like to run through an initial installation, and then work it through several iterations of upgrades, both for increased service level and for technological upgrades. Working out how well the equipment specified by this standard meets this test will help us guage market acceptance.

- Provide for future delivery of all basic services on the fiber.

The big items here are of course telephone service, either VOIP or over the 64k channels, video, and the ethernet data service.

End-User Objectives

- Pay for guaranteed bandwidth with always-on service

This is what people like about DSL and DOCSIS based services, it's just the install/service headaches and the lack of bandwidth that they dislike.

- Symmetric bandwidth

Symmetric bandwidth is important in part because TCP assumes that the forward and reverse paths have same bandwidth and latency.

The compelling reason for this, though, is that there are going to be businesses operating on this network, and they are going to be operating servers. Also, there's going to be an inreased demand from people to operate servers and Napster-type services from their homes. This is something the operators currently forbid or at least discourage, but it could be an extra revenue-generating feature for them.

- Bandwidth to users scaleable in small increments

This is important because the operators will want to offer a basic service, then different grades of more expensive services, with the thought that people will buy as much bandwidth as they can afford. If the service levels are 1Mbit, 10Mbit, amd 100Mbit, there's a lot of people that could not afford a 10Mbit service, but could afford and buy a 2 or 3 Mbit service.

- Quick turn-around on bandwidth adjustment requests.

This could be important for businesses that need to turn up their bandwidth for short periods of time for special purposes, such as collaborating on a project, or a filmmaker sending a digital film out for some special editing.

Service Provider Objectives

- Ease of management
  - Ability to bill by usage (by-the-bit)
  - Flexible subnet granularity
  - Control and feedback about latency, statistics, security
  - Add new users without service interruption
- Fault isolation capabilities
  - Limit how many homes can be affected by a CPE fault

If 300 homes are on the same fiber and a transmitter gets stuck on, or fails to talk at the right time, 300 homes could be affected with little information about where the problem is. A piece of CPE should be restricted to being able to affect only, say, 25 subscribers or so, limited to a relatively small geographic area.

  - Allow problem source detection while faulted
- Flexibility of configuration
  - Multi-Dwelling-Unit uses same components
  - Same system for homes, businesses, “power users”

Also important is the ability for the service provider to have a solution which allows him progressive deployment capabilities.
There’s no need for the service provider to install home equipmet or even fiber until that subscriber is ready to be turned on, and a system that requires him to do so is very unattractive.

local-switching an option but not a requirement of the standard

Protocol Objectives (1)

- Stop using the same hammer (SONET/ATM) for every problem that comes up

The industry has its share of pundits and detractors for these technologies, and I feel that SONET has its useful places, ATM somewhat fewer useful places, but that neither belongs at the edge of the network, in the last mile. We have an apllication with unique but rather simple needs, so we should find a solution that has a better level of equipment and protocol simplicity than either SONET or ATM. To that end, I feel that...

- Simple TDM and TDMA structures provide simple and adequate method for transporting Ethernet and legacy voice
- No transport of control information via the encoded Ethernet packets
  - Fewer SNMP agents than subscribers: Headend equipment is SNMP agent for connected field equipment
  - CPE requires no MAC, IP address
  - Enhances security
  - More robust

No home or business user will ever be able to access the control information, or deny access to it.

In ADC scheme the control protocol rides in the overhead.

Any desired level of security is obtainable via encryption of the data in the pipe.

Protocol Objectives (2)

- Simple scheme with more security than one based on VLAN tags

Leaving the packets untouched will simplify the processing requirements of the equipment, and there is often a relatively small limit to the number of vlans an individual piece of equipment can process. (shared address table model)

- QoS via bandwidth: total of subscriber guaranteed bandwidth does not exceed WAN capabilities.

The ultimate goal here is full 10 or 100Mbit service to all the homes, which will allow Quality of Service to be handled by the software on the PCs in the home, and the servers they are connecting to.

- Protocol independent transport of the Ethernet data so future protocols can be transported without upgrade

For example, if the packet is not looked at deeper than the MAC level to allow low-cost equipment, the upgrade to IPv6 and beyond would be transparent to the network. (almost think of it as just an extension of the physical layer)

- Forward path data destination fully determined at the headend equipment

HE equip. required to place data into standard format (sorting) but can do it in different ways (L2, L3, VLAN Tag). right mode not always the same

- Minimum bandwidth guaranteed, full 10 or 100Mb service given on the fly if available

- 64kbps telephony provisions

Legacy telephony services should be a part of the protocol (in the overhead) so service can be provided immediately.

Equipment Objectives (1)

- Initial install supports many low data-rate subscribers, then headend equipment is added to increase data rate to each subscriber
- Concentration of hi-end services in the headend
- Customer-premises equipment as simple as possible
- Specify the required subset of switching rules

Some obvious must haves are filtering errored packets, port mirroring, statistics, MAC/IP address management and security features.

- DWDM not required at outset
  - DWDM can be transparently added to network later
  - Use and non-use of 1310/1550 WDM part of standard

DWDM on the pole will be more reasonable when laser frequencies are managed by some means other than temperature control.

Equipment Objectives (2)

- Components essentially the same for different targets
  - Equipment suitable for use in Multi-Tenant-Unit applications

What I envision here is that while the packaging would probably be different, the same circuit boards would be used, and laser and receiver card slots in some of the equipment would also accept electrical interfaces

  - Same equipment used for home and business subscribers

- CPE has generic plug-in components

The equipment in the home would be able to accept an…

  - Ethernet plug-in as CAT-5 copper, fiber, Bluetooth, HPNA

, a plug-in for

  - Basic telecom

services that would be used for the 64kbps channels or VOIP, and a future plugin for

  - Video

over IP services.

- Multiple standard interconnects at the headend (Gigabit Ethernet, 10Gig Ethernet , SONET)

Equipment Objectives (3)

- Ethernet based connections to WAN support 802.3 LACP
- WAN connection equipment required to place data into standard format (sorting) but can do it in different ways (L2, L3, VLAN Tag)

I think this is pretty important. The aggregating equipment that connects to the WAN is going to be used by different service providers with different philosophies, and for different uses. Having a variety of equipment available will allow service providers to meet unique and emerging needs.

- Standardize pre-aggregated return path data rate at 0.622GHz, 1.2GHz, 2.5GHz, to allow inexpensive initial deployments

The specific data rates here are not what's important, but the importance of having the customer premises equipment be as low-cost as possible really can not be overstressed. The units in the homes can be made to accept a laser plug-module of any speed, and indeed there's little reason that lasers of varying speeds can't be on the same fiber.

Compatibility

- EFM transport is transparent to 802.3 features
  - IEEE standard Virtual Lans (VLANs)
  - Spanning-Tree
  - Link Aggregation Control Protocol

Once again, this is to ensure the equipment in the field and in the headend can go untouched for years, evan as the 802 standards are upgraded and expanded.

A Sample Architecture

Some Preliminary Data for Your Consideration

A hybrid fiber/co-ax cable headend can reach 50,000 people: a passive optical network architecture we've been looking at can reach far enough to serve those homes with 2.5Gbit/s symmetric bandwidth (with 2Gb/s ethernet payload), one fiber for upstream and downstream, 1550nM downstream, 1310 upstream

We've come up with a simple line-oriented framing scheme with efficient packing of packets for mutiple dwellings per line, service-level data, addressing data for terminating units, ranging data, telecom service, control, monitoring, and maintenence communications.

For the service provider installing this system there needs to be a pretty big differentiator from DSL/DOCSIS based services, and we think this architecture has the potential to do that.

24 homes, 10/100 service mix

The equipment that connects to the WAN and serves as an aggregator/distribution element is referred to as the OLT, for Optical Line Terminator. In a typical confifuration, it has 2 1Gb/s trunked ethernet links to a router or layer-3 switch.

The next active element in the field is referred to as the FMUX, for Field Multiplexer.

The next and last active element residing in the field is the Broadband Termination Unit, or BTU. It is the equipment ion the home or business.The OLT sends out framed lines of data to the FMUX, which except for minor alterations sends the data out to the BTUs unchanged. The framed data to the BTUs does not have specific timeslots assigned to homes, but is dynamically addressed.

Every home connected to the FMUX receives the same data, but only pulls off the data specifically addressed to it. This gives a notional amount of security, which is enhanced by having the payload data of each home encrypted using a separate key for each home.

Each BTU is given a window of opportunity to transmit data back to the FMUX in a burst-mode fashion, and a BTU can dynamically request more opportunities to transmit. Automatic ranging functions obviate the need to pre-measure the optical path delay.

All data in the forward path, towards the home, is sent with a 1550nm laser, and all return path data is sent with 1310nm lasers. There’s more of the return path lasers, so it makes sense to use the cheaper 1310nm lasers for the BTUs.

The only optically sophisticated devices in this network are the quad-taps, which split the downstream data 5 ways, that is to 4 homes and to the next tap, and in the reverse path they directionally combine the light from the 4 BTUs onto the fiber towards the FMUX only.

To summarize, this is a bus architecture that implements a logical star

Utilizing a splitter, not shown here, and 4 receivers, an FMUX/OLT can service 96 people with an 80%/20% split on 10bT/100bT service. 521 OLTs with 2 LACP GbE trunked interfaces each, for a data rate of 1.1Tbit/s, would be required to service 50,000 homes.

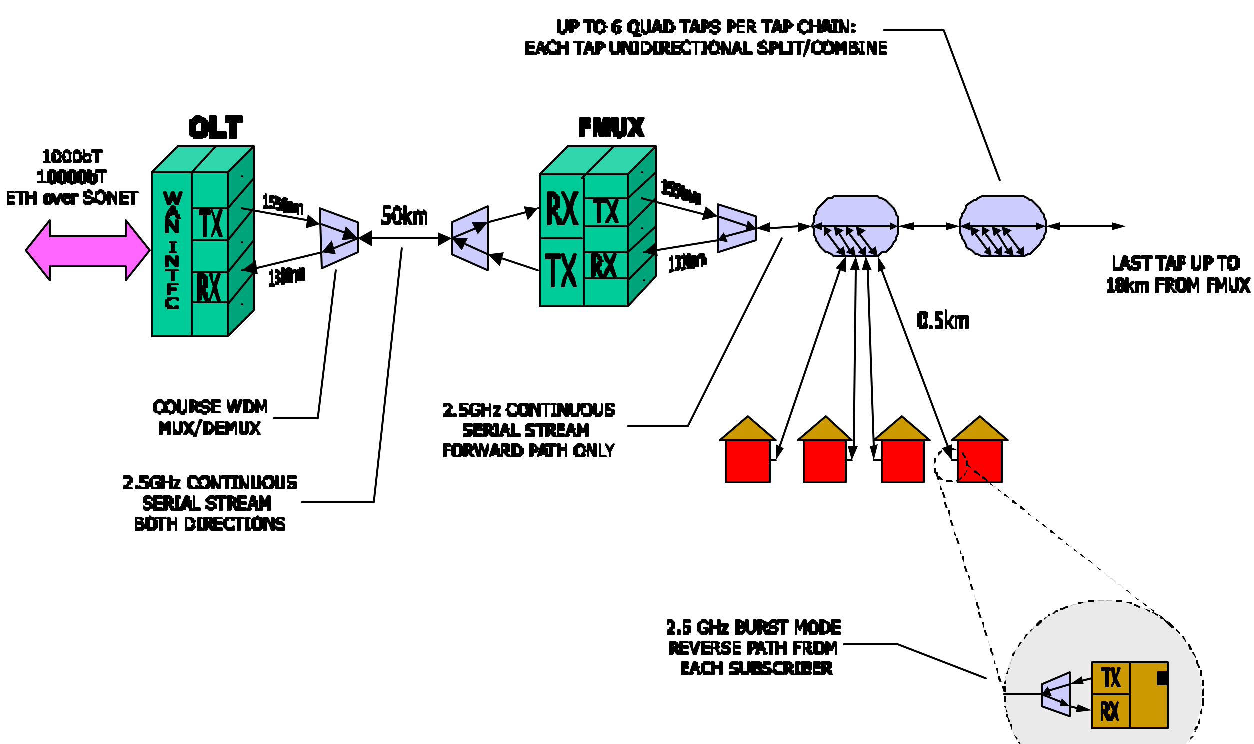

FTTC using curb / pole electronics to serve 4 homes

If the service provider is using a fiber-to-the-curb based network, the same equipment is used to connect to a pole- or pedestal-based BTU that serves four homes. The optical tap only splits the data two ways, and the access to the home is over a copper-based medium. The possibilities for that copper medium are

- a dedicated set of twisted pair wires,
- a twisted pair shared with phone service via an extension of HPNA or using a simplified DSL arrangement,
- or a coax cable with a significantly simplified “individual cable modem” at each end,

allowing video and data to share an existing coax connection to the home.
A fiber-to-the-curb multi-subscriber BTU can reside on the same fiber as a fiber-to-the-home single sub. BTU.

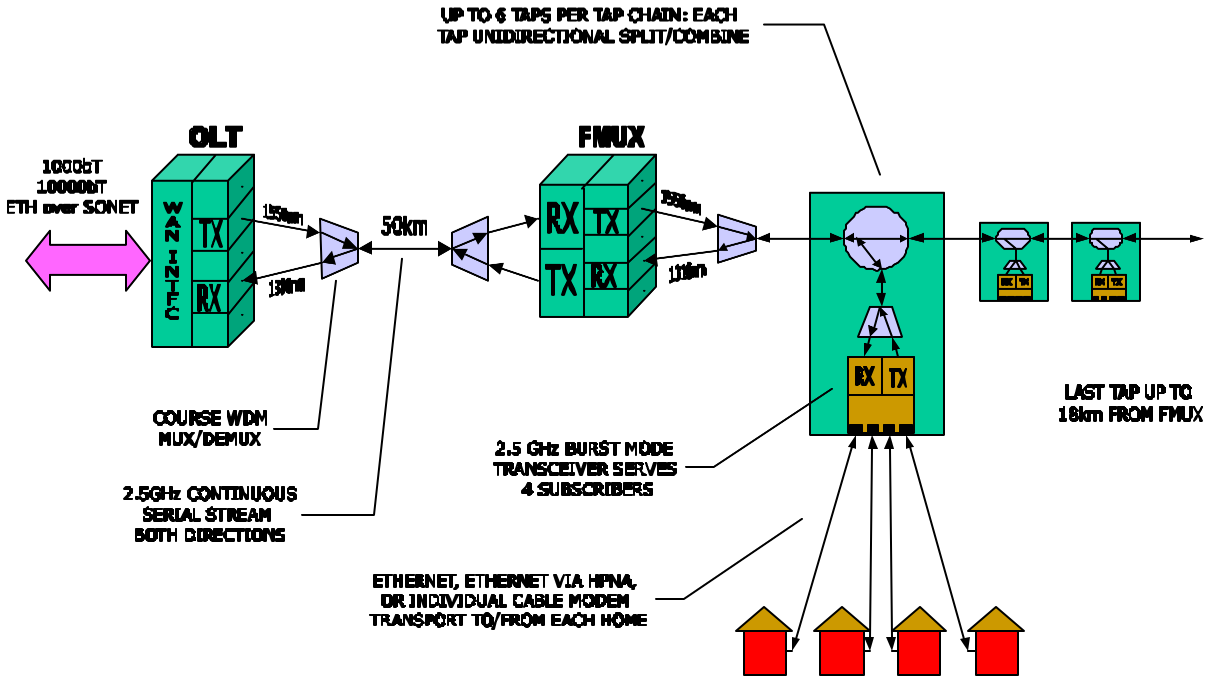

Single OLT, Quad FMUX, Quad TAP chains, 384 homes served

It is anticipated that in an initial deployment, the goal would be to serve many more homes than discussed above, but with a lower data rate for each home. The infrastructure would later be improved through deploying more OLTs at the headend only, reusing and not even disturbing the FMUXs and BTUs.

In this figure there are 384 people served by each OLT, and each OLT has a single Gbit/s ethernet connection to the WAN. Each subscriber has 1 Mbit/s guaranteed service, and 130 OLTs are required for 50,000 subscribers. The Gbit ethernet link is less than half utilized, and would be fully utilized if each home had 2.6Mbit/s guaranteed service.

Note that in the 1Mbit service case the total data rate of the headend is a comparatively manageable 50Gbit/s.

384 people per OLT, 5 Mbit/s guaranteed service each, 130 OLTs with 2 LACP GbE trunked interfaces each (260Gbit/s data rate at headend)

The fiber count for a 50,000 subscriber headend with this scheme would be 50,000 subs x 1fiber/96 subs, or 521 fibers.

A critical thing to note here is that there are 130 OLTs serving 50,000 homes with up to 2.6Mbit/s service each, and by increasing the number of OLTs in the headend and not touching a single piece of field equipment, the data service to the subscribers can be increased to 10Mbit/s for about 80% of the subscribers, and 100bT for the remainder. The end result is that every fiber at the headend now terminates in its own OLT.

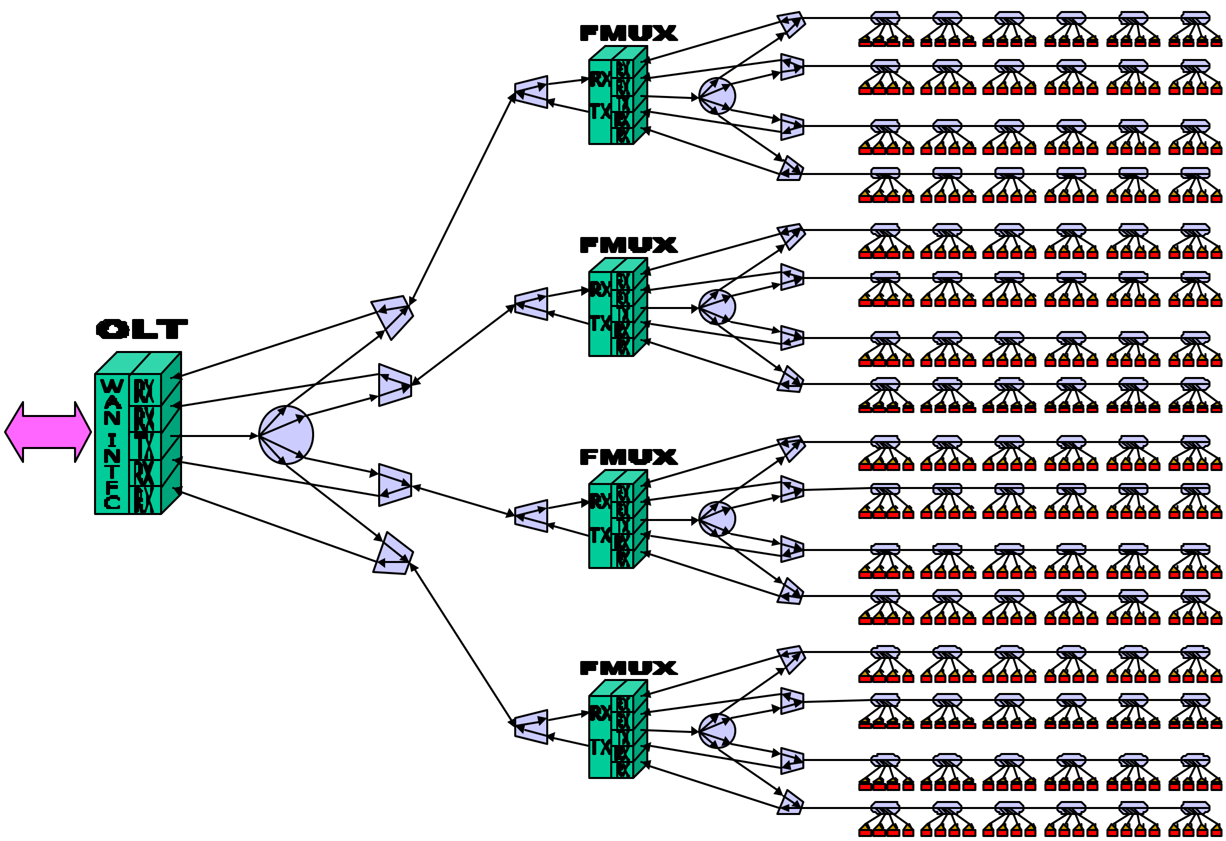

FRAMING SCHEME

Here's a quick view of the downstream framing scheme we've developed to date. It addresses the payload transport, the control information for the ranging required by the burst-mode return path, 64kbps POTs channels, subscriber service level, and addressable packetized information for control, statistics, maintenance, and provisioning.
As mentioned, all control and ranging, etc., data is segregated in the overhead.

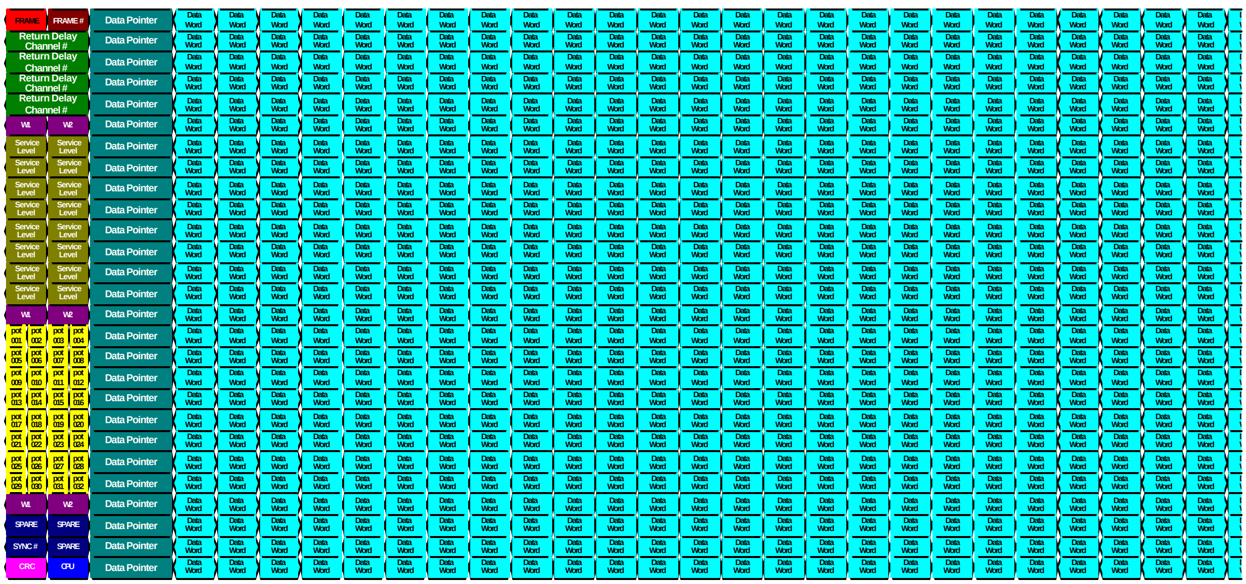

- FRAME — Framing "WORD" repeated at 192 kHz '0100 0101 1010 1101'
- FRAME # — Framing "COUNT" repeated and advanced at 192 kHz; when "24" reached (8 kHz) rolls over and advances Super Frame Counter
- Return Delay Channel # — Callculated Value for return path from TAP
- Data Pointer — First 10 Bits Indicate 'Owner' via MAC address pointer, next 3 bits Indicate 'Type', next 5 bits are 'Start Pointer' on word boundry, last 12 bits are 'End Pointer' on byte boundry.
- Data Word — Actual Data to be 'Drop Shipped' to User/Costomer
- W1 — Data Words sampled at 576 khz (2)
- Service Level — Word for Controlling Service's at a Customer Premisses
- pot 001 — Byte of POTS designated to user via Service Level Word.
- SYNC # — Time Stamp for reference to regenerate 10 Mhz or 100 Mhz as needed.
- CRC — Cycle Redundancy Check Word for entire frame for measurment of quality of conn
- CPU — Word to be used for CPUs to link up with each other.

Closing Remarks

- Fiber to the home will be here soon

It's surprising, but companies are finding out that in looking at the total cost of ownership, fiber is turning out to be cheaper than copper, mainly because of maintenance issues. There's been a little shakeup in the market recently, but this will come.

- Aim to provide as many services as possible

We need to work at preventing a chicken-or-the-egg syndrome

- Remember the broad mix of infrastructure providers that will be installing equipment based on this standard

Wide acceptance is going to depend on the standard being applicable and usable