RAD

data communications

www.rad.com

Everything about TDMoIP

Unique Access Solutions

PWE3 – 52nd IETF

12 December 2001

Classic Telephony

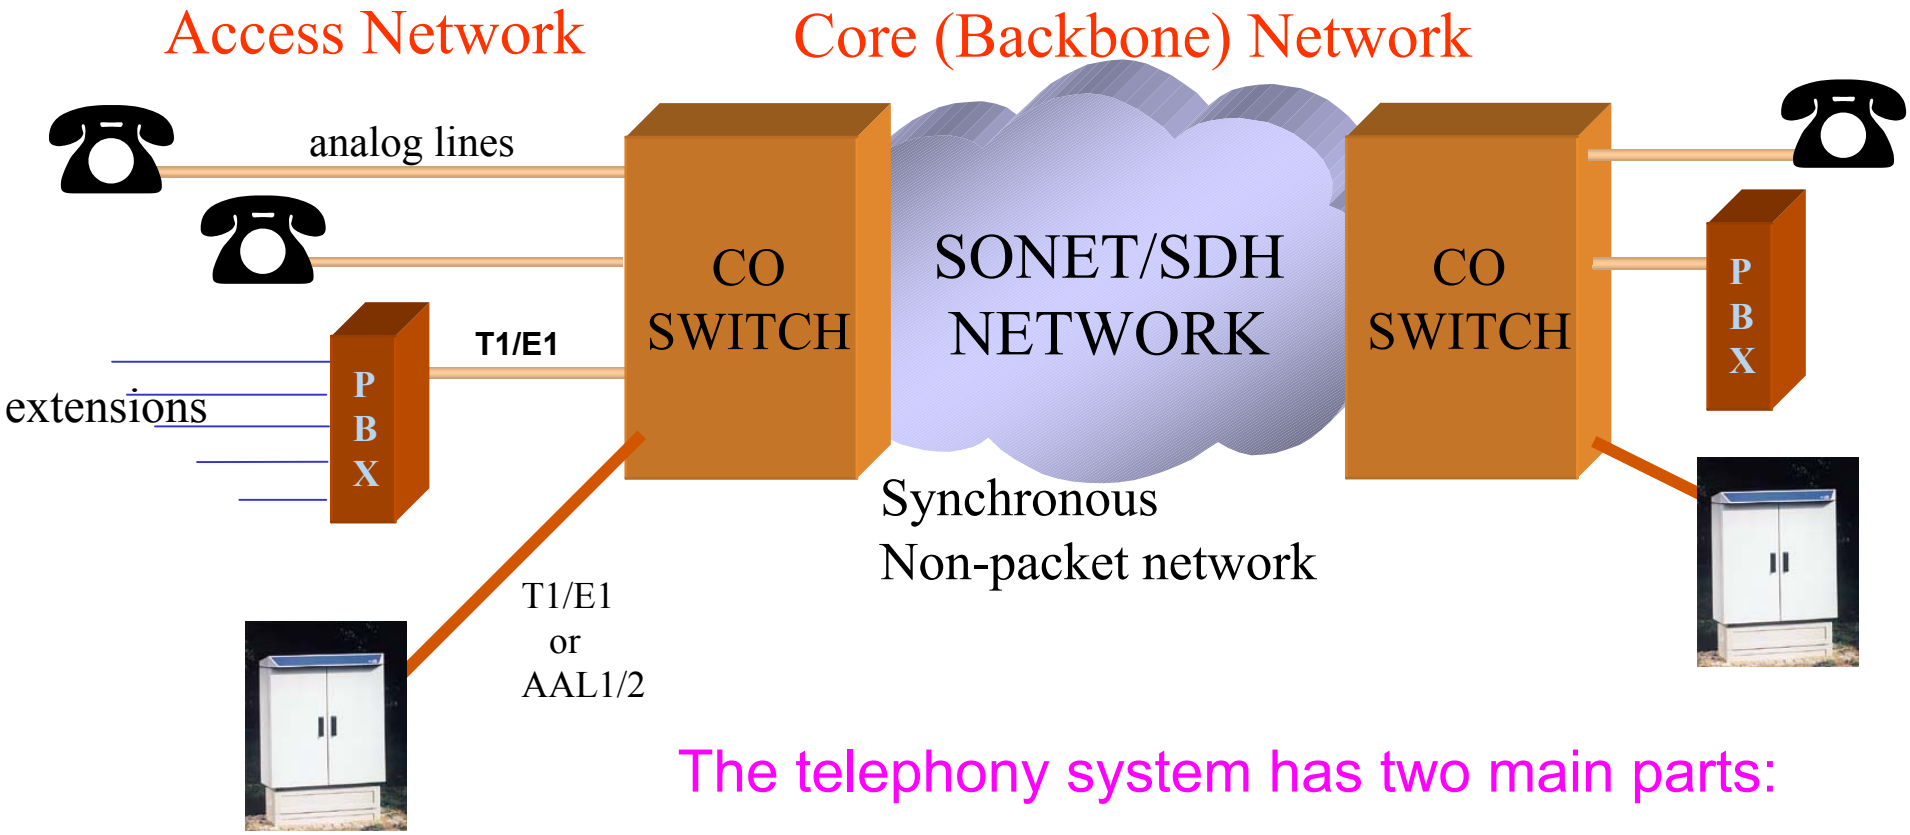

The telephony system has two main parts:

- Access network (analog, T1/E1, AAL1/2)
- Backbone network (SONET/SDH)

TDMoIP

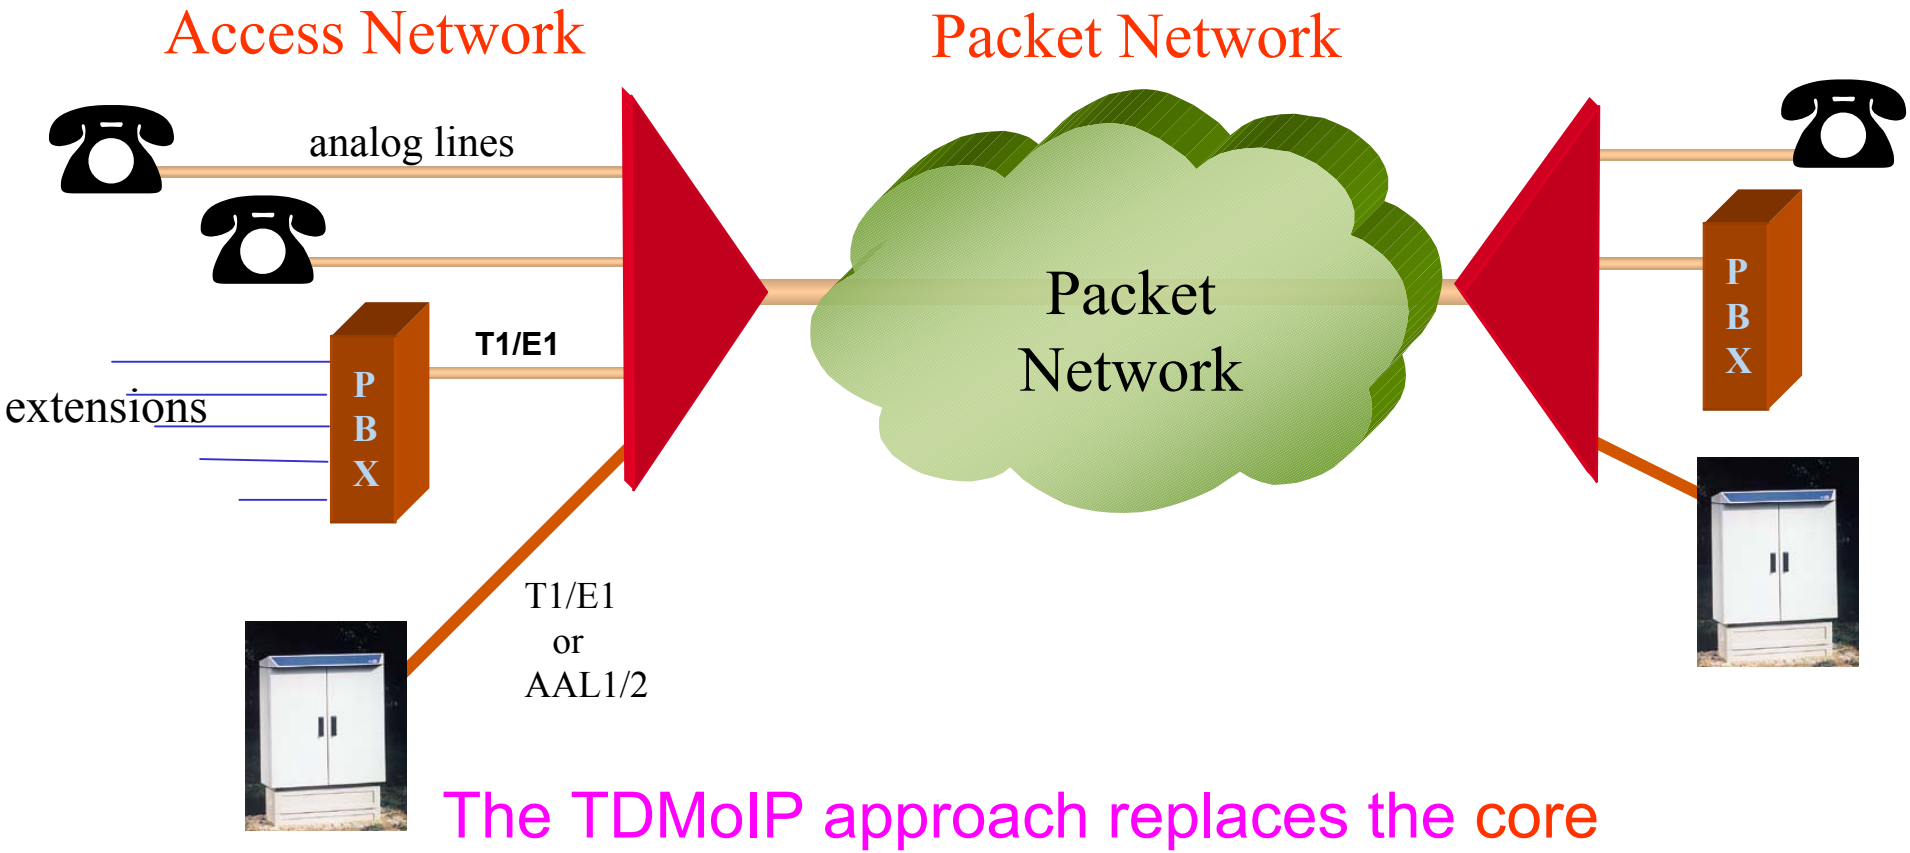

The TDMoIP approach replaces the core
with a packet (IP or MPLS) network
The access networks and their protocols remain !

draft-anavi-tdmoip

SONET/SDH CEP

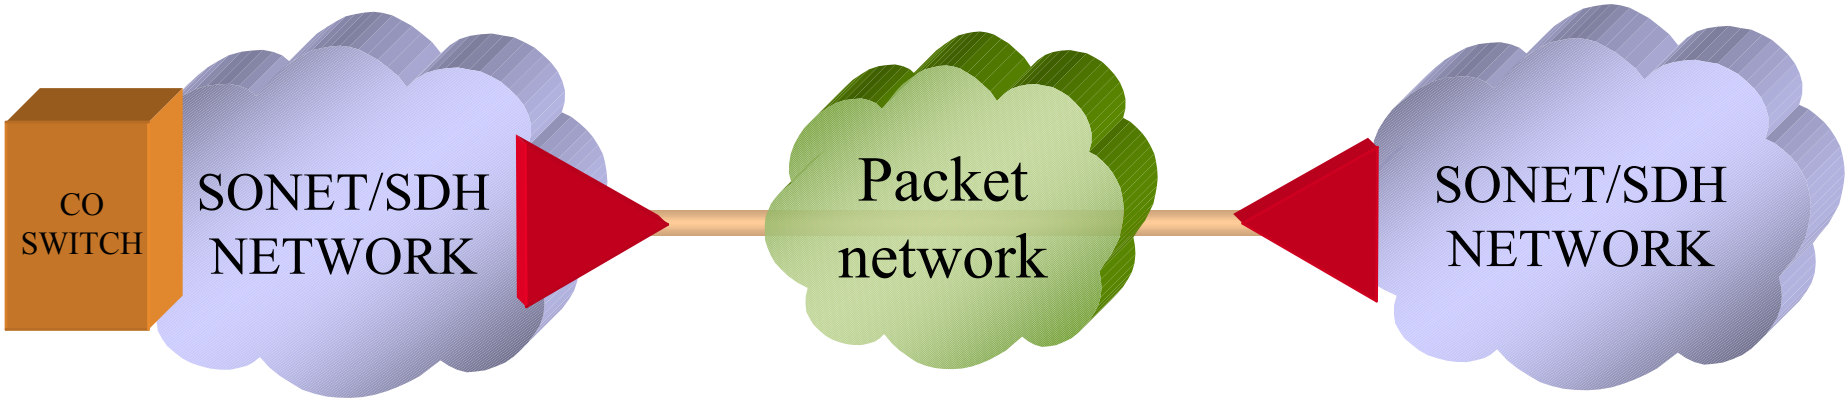

Circuit Emulation over Packet interconnects

different SONET/SDH networks

The packet network becomes a carrier's carrier

draft-ietf-pwe3-sonet

Related (but different) Applications

VoIP connects individual users over IP networks
replacing all signaling with new protocols

VoDSL connects users over DSL connections
using VoATM technologies

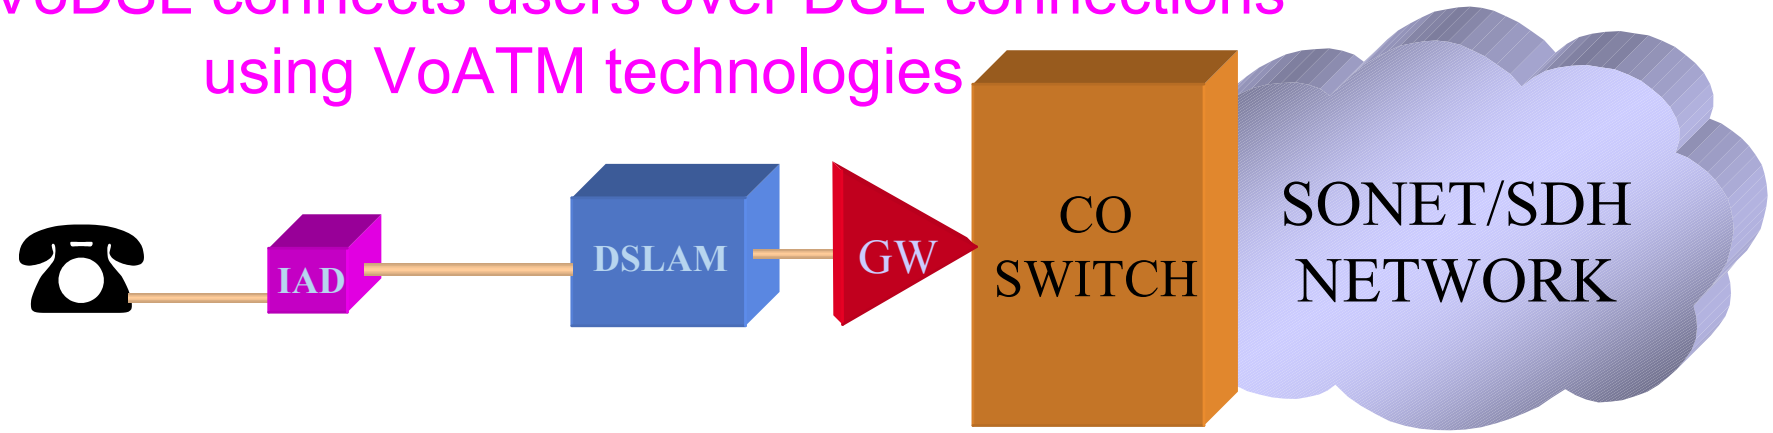

Functionality

What needs to be transported from end to end?

- Voice (telephony quality, low delay, echo-less)
- Tones (for dialing, PIN, etc.)
- Fax and modem transmissions
- Signaling (there are 1000s of PSTN features!)
- Timing

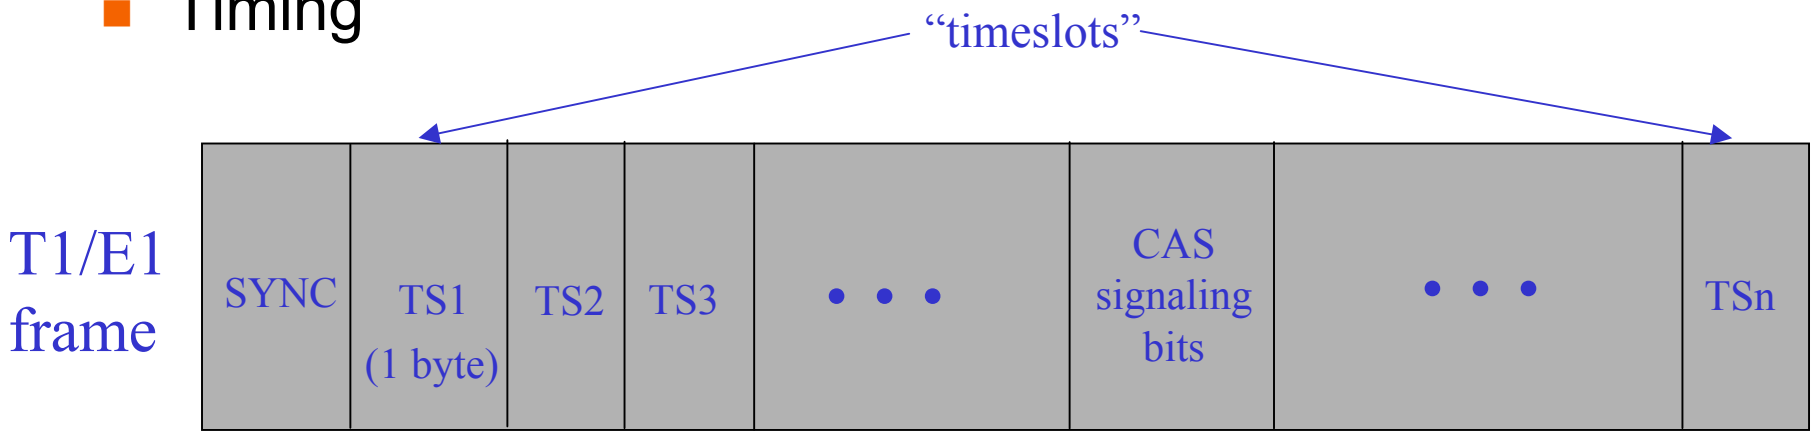

Note:
Various proposed extensions to RTP that multiplexed voice sessions are not applicable since they only handled the voice!

Why isn't it easy?

Why don't we simply encapsulate the T1/E1 frame?

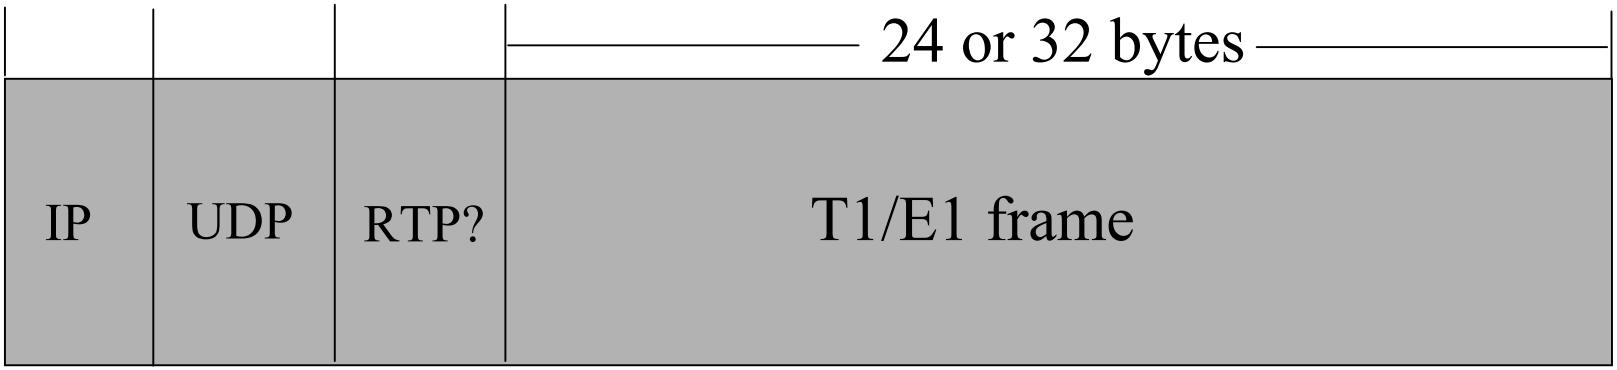

Because a single lost packet would cause service interruption

- CAS signaling uses a superframe (16/24 frames)
- superframe integrity must be respected

Because we want to efficiently handle fractional T1/E1

Because we want a latency vs. efficiency trade-off

I have an idea!

Those problems can be solved by:

- adding a packet sequence number
- adding a pointer to the next superframe boundary
- only sending timeslots in use
- allowing multiple frames per packet

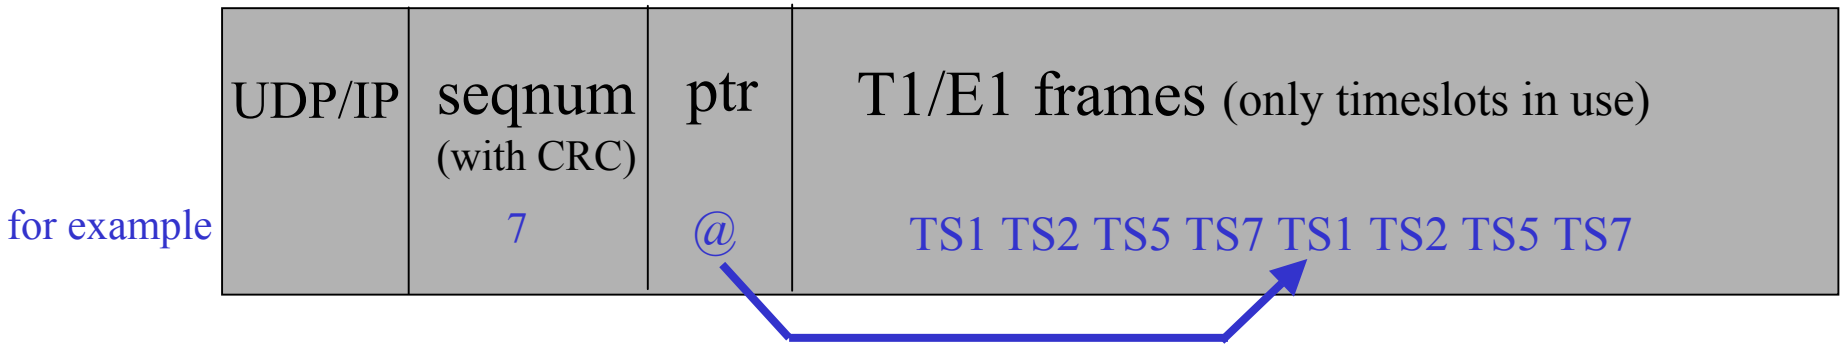

Good idea! That is precisely AAL1 !

Why isn't that enough?

AAL1 is inefficient if the timeslots

- are "hard-wired", and
- not always in use

Although we can configure which timeslots are used

we can not change this configuration in real-time!

To allow dynamic allocation of timeslots

we can use AAL2

AAL2 buffers each timeslot and encapsulates it in a "minicell"

Isn't this just ATM?

AAL1 and AAL2 are adaptation protocols
originally designed to massage data into a format that can readily use

As we have shown, they are natural candidates for any application which needs to multiplex timeslots

For TDMoIP we do not put the AAL1/2 into ATM cells (no 5 byte header)

Rather we put the AAL1/2 directly into a UDP/IP packet

So, NO, this is NOT ATM

But it can easily interwork with ATM access networks!

What about RTP?

RTP is not a channel multiplexing protocol,
so this issue is orthogonal to that of the previous slides

RTP can be used to transport timing across IP networks

It does this by providing:

- a 16 bit sequence number
- 1 32 bit timestamp

at the expense of 12 additional overhead bytes per packet

Accurate timing is important in telephony
and IP networks add jitter

Don't we *need* RTP?

When RTP is *not* needed

RTP adds significant overhead – can we get away without it?

In many TDMoIP applications

all end-user equipment have access to accurate (stratum 3?) “station clocks”

So timing info need not be distributed over the IP network!

Even when adaptive (FLL/PLL) timing recovery is needed

the RTP timestamp does not improve accuracy as compared with a sequence number

since E1/T1 frames are sent at a precisely periodic rate as determined by the transmitting station clock!

TDMoIP frame structure

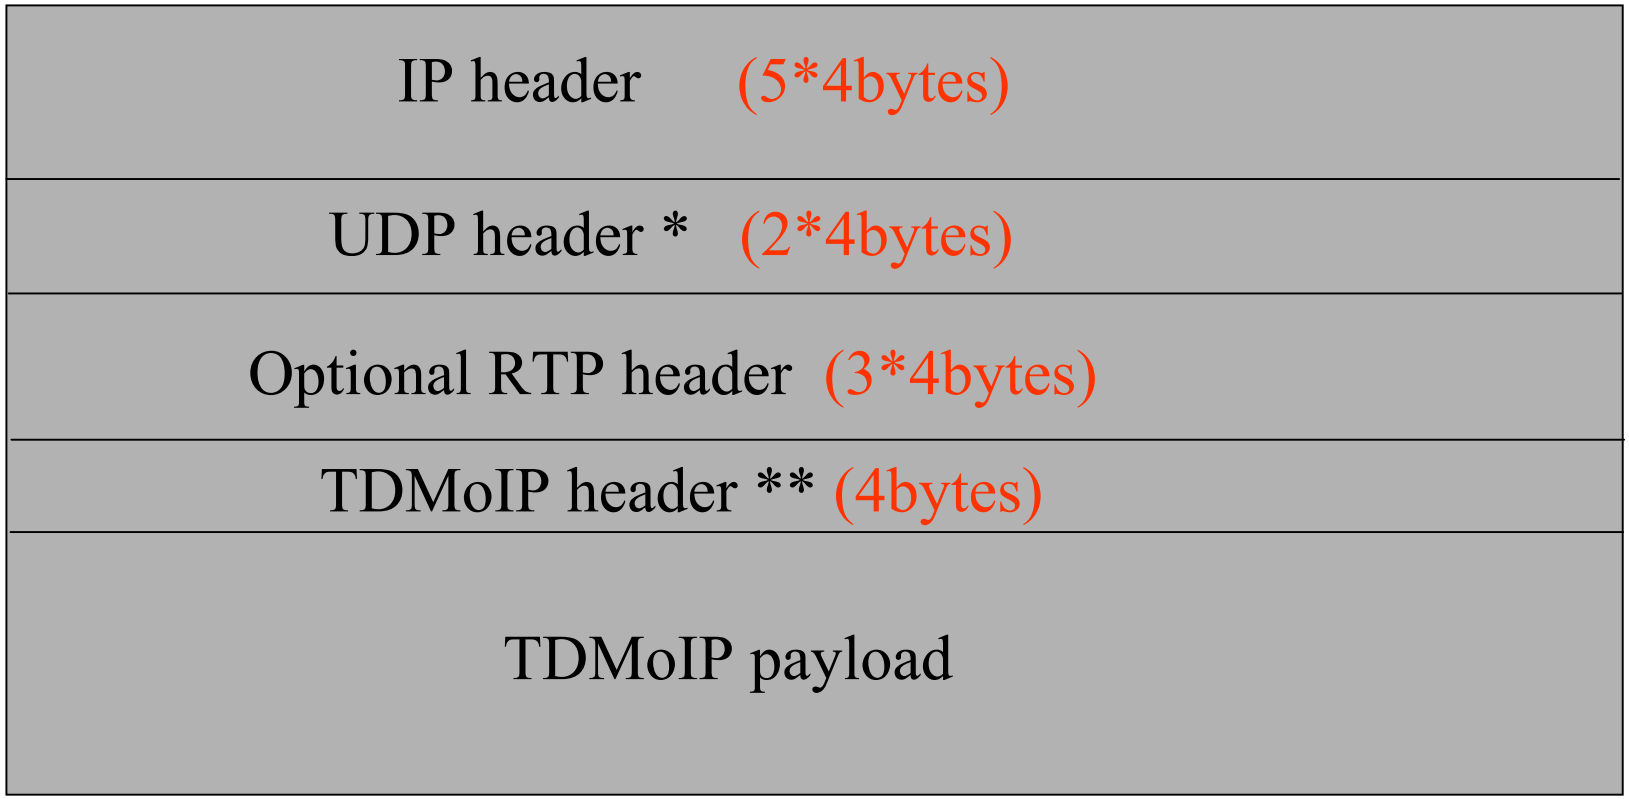

Notes

* The UDP source port number is used as a bundle identifier

** The TDMoIP is essentially the header defined in Martini et al

Further Advantages

HDLC support

- CCS signaling can be delivered

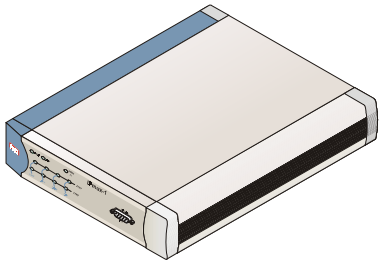

Simple implementation

- Processing for single T1/E1 performed by embedded CPU
- Large system price-per-channel is extremely low
- No “fork-lift” upgrade needed

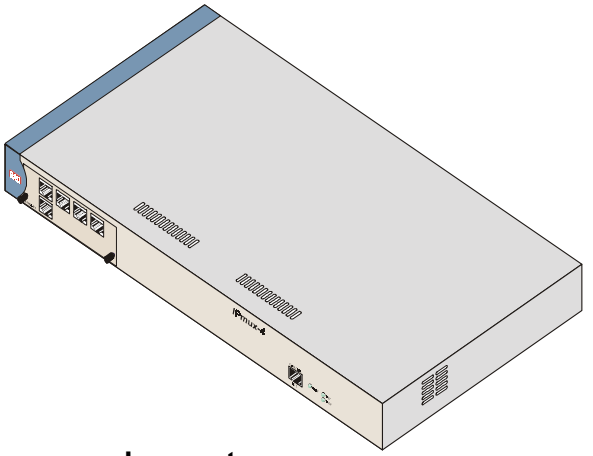

Field Proven Technology

- 1500 units in the field
- Over 5000 T1/E1 trunks

Municipal networks, school districts, business parks, etc.