

FPGAs Enable Pseudo Wire Emulation Edge-to-Edge

Network convergence is possible using PWE3 for transport and delivery.

by Anthony Dalleggio
VP, Marketing
Modelware
dalleggio@modelware.com

Pseudo Wire Emulation Edge-to-Edge (PWE3) provides methods to carry networking services such as ATM, Ethernet, and TDM over a packet switched network (PSN). It has been gaining in popularity because of compelling economic and operational factors.

Let's say that you are a network operator who needs to add capacity for ATM or E1/T1 services. Should you invest in growing those networks or carry those services over your IP network, where you may have spare capacity? In most cases, the choice is clear – and the answer is PWE3.

PWE3 can consolidate the edge-to-edge delivery of multiple service offerings and traffic types such as E1/T1, ATM, and PPP/HDLC onto a single PSN. This is, of course, not without technical challenges such as accurate clock recovery.

In this article, I'll outline the networking problems that PWE3 solves, how PWE3 works, and describe a PWE3 implementation on Xilinx® Virtex™ and Spartan™ parts.

Reference Model and Services

The Internet Engineering Task Force (IETF) has done extensive work on PWE3. The IETF PWE3 reference model

in Figure 1 shows two customers connected through an attachment circuit (AC) to their respective service providers. The native service exists between the customer edges (CE1, CE2) and their provider edges (PE1, PE2). The PW (pseudo wire) spans the PE-to-PE segment and carries the emulated service over the PSN.

PWE3, as any networking layer, must support data plane, control plane, and management functions.

Table 1 lists general and specific requirements for the PWE3 data plane. Some requirements are optional or not applicable depending on the native service being emulated.

Control Plane and Management Considerations

In addition to the data plane requirements, PWE3 must provide the following maintenance requirements:

- Setup and teardown of a PW
- Support and handling of CE-to-CE in-band and out-of-band maintenance messages (such as OAM)
- Support of PE-initiated maintenance messages

Management information bases (MIBs) are required to support provisioning, per-

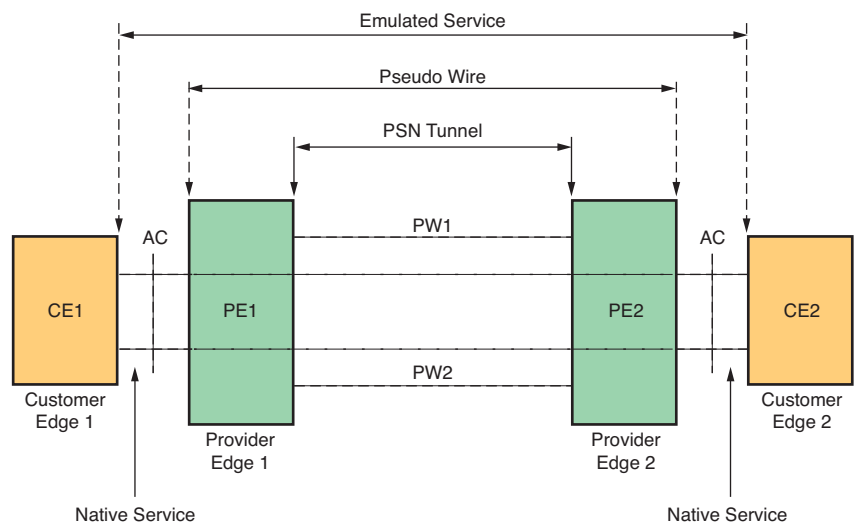


Figure 1 – PWE3 reference model (Source: RFC 3985)

General Requirement	Specific Requirement	Description
Encapsulation	L2 Header Information	To convey L2 header information to allow correct processing of the packet at the PW egress
	Variable-Length PDUs	To support variable-length PDUs if they are allowed by the native service
	Multiplexing/Demultiplexing	To support “trunking” mechanisms like multiple ATM VCCs in one VPC
	Validation of PW-PDU	To support payload validation, the L2 checksum is either passed through the PW or removed at the ingress and recalculated at the egress if the L2 header is modified at the PW entry and exit points
	Payload Type Information	To allow the differentiation between PWE3 and other traffic (such as IPv4) to reduce the probability of packet misordering by load-balancing mechanisms
Frame Ordering		To ensure in-order packet delivery for services that require it; this can be accomplished using a sequence number in the PW header
Frame Duplication		To ensure that duplicated frames are not delivered for services that do not allow it
Fragmentation		To fragment packets whose length, including PW and PSN headers, exceeds the maximum transmission unit (MTU) of the PSN; long packets may optionally be dropped
Concatenation		To concatenate short PDUs into a single PWE3 packet to increase efficiency over the PSN, considering the resulting delay and jitter
Quality of Service		To ensure low delay and jitter characteristics for services that require them
Timing		To generate a service clock at the egress of the PW with a frequency that is close or equal to the ingress service clock

Table 1 – PWE3 requirements

formance monitoring, fault management, and connection verification.

PWE3 Protocol Mapping

Depending on your target network, you will want to specify one of several forms of packet encapsulation. For example, you can carry PWE3 traffic over an IP or an MPLS network. A “light” version also exists for point-to-point solutions.

PWE3 over IP

Figure 2 shows how to map PWE3 traffic onto an IP-based PSN. The payload of the native service is normally passed as is without processing. You can rely on RTP for timing and sequencing; otherwise, a sequence number is added to the PW encapsulation. For PW demultiplexing, you can use the MPLS label, L2TP session ID, or the UDP port number.

PWE3 over MPLS

Figure 3 shows the PWE3 to MPLS mapping. The MPLS mapping is more efficient because you can compress some of the PWE3 information into one control word. For PW demultiplexing, you can use an inner MPLS label. The sequence number is carried in a control word.

PWE3 “Light”

Certain point-to-point or private network applications do not require all of the PSN overhead. For these applications, you can carry the native service directly over Ethernet with minimal PW header information to identify any subchannels and carry timing information if required.

Pseudo Wire Timing Recovery

Not all native services require timing recovery. In general, non real-time services (such as data) do not need timing recovery at the destination. However, real-time services such as TDM do require timing recovery. There are three basic approaches to timing recovery: absolute, differential, and adaptive. Regardless of the approach, you can generate the clock using analog or digital techniques.

Absolute Mode

This is the normal method used for real-time protocol (RTP). The sending end generates a time stamp that corresponds to the sampling time of the first word in the packet payload. The receiving end uses this information to sequence the messages correctly but without knowledge of the sending end’s clock frequency. This method is used when absolute frequency accuracy is not required.

Differential Mode

In the differential mode, both sending and receiving ends have access to the same high-quality reference clock. The sending end generates time stamps relative to the reference clock. The receiving end uses the time stamps to generate a service clock that matches the frequency relationship of the sending end’s service clock to the reference clock. This method produces the highest quality clock and is affected least by network quality of service issues.

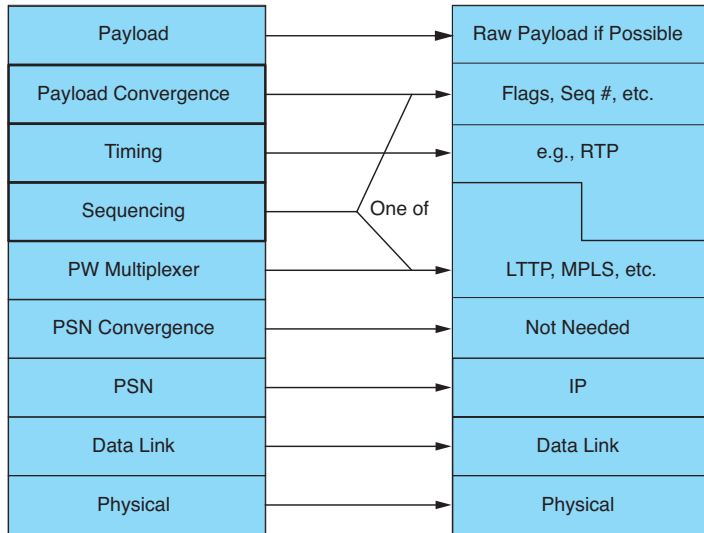


Figure 2 – PWE3 over IP (Source: RFC 3985)




Figure 3 – PWE3 over MPLS (Source: RFC 3985)

Adaptive Mode

The adaptive clock recovery mode relies on packet inter-arrival time to generate the service clock frequency. This method does not require time stamps or a reference clock to be present at the receiving end. However, it is affected by packet inter-arrival jitter.

Clock Generation

Modelware's PWE3 clock generator implements the adaptive and differential mode timing recovery for E1/T1, using digital clock synthesis, in Xilinx Spartan-3 and Virtex-4 FPGAs. A four-channel differential clock generator uses 679 slices, 8 block RAMs, and 1 global clock buffer.

Low frequencies such as E1 (2.048 MHz) and T1 (1.544 MHz) can be generated completely digitally using a high-speed clock while meeting jitter and wander requirements. A typical clock generation circuit is shown in Figure 4.

Note that the circuit behaves like a phase lock loop (PLL), where the frequency error is based on the buffer fill level for the adaptive mode and on the reference clock for the differential mode.

At Modelware, we have implemented adaptive- and differential-mode timing recovery for E1/T1 using digital clock synthesis. The advantages of our solution are the linearity of the digitally controlled oscillator (DCO) and the ability to configure the frequency resolution to very small values.

Conclusion

In this article, I have discussed the motivation, architecture, and services provided by PWE3. Because of its attractive cost structure and the promise of network consolidation, PWE3 is gaining momentum. This is true not only in public and private networks, but also as a low-cost approach for point-to-point intra-system connectivity.

Modelware offers PWE3-related IP cores and design services. For more information, please contact Anthony Dalleggio at (732) 936-1808 x222, or e-mail dalleggio@modelware.com.

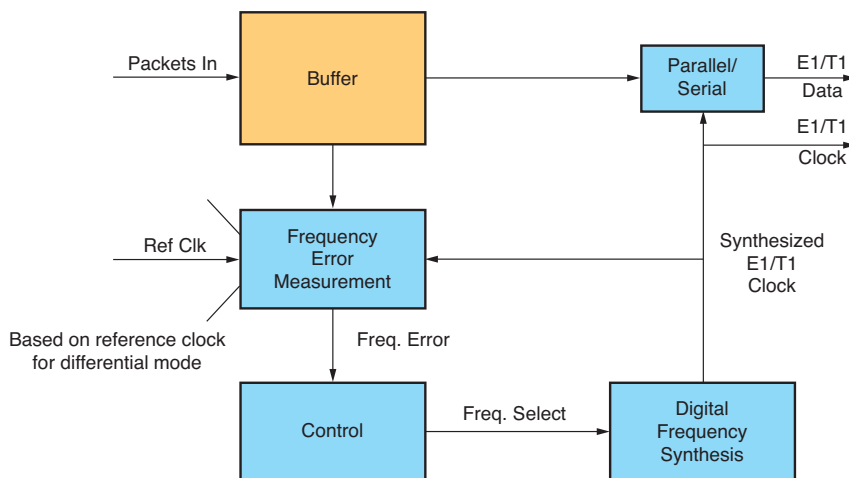


Figure 4 – E1/T1 clock generation