

**Pointers**

SONET uses a concept called “pointers” to compensate for frequency and phase variations. Pointers allow the transparent transport of synchronous payload envelopes (either STS or VT) across plesiochronous boundaries (that is, between nodes with separate network clocks having almost the same timing). The use of pointers avoids the delays and loss of data associated with the use of large (125-microsecond frame) slip buffers for synchronization.

Pointers provide a simple means of dynamically and flexibly phase-aligning STS and VT payloads, thereby permitting ease of dropping, inserting, and cross-connecting these payloads in the network. Transmission signal wander and jitter can also be readily minimized with pointers.

Figure 10 shows an STS-1 pointer (H1 and H2 bytes) which allows the SPE to be separated from the transport overhead. The pointer is simply an offset value that points to the byte where the SPE begins. The diagram depicts the typical case of the SPE overlapping onto two STS-1 frames. If there are any frequency or phase variations between the STS-1 frame and its SPE, the pointer value will be increased or decreased accordingly to maintain synchronization.

**VT Mappings**


There are several options for how payloads are actually mapped into the VT. Locked-mode VTs bypass the pointers with a fixed byte-oriented mapping of limited flexibility. Floating mode mappings use the pointers to allow the payload to float within the VT payload. There are three different floating mode mappings – asynchronous, bit-synchronous, and byte-synchronous.

**Concatenated Payloads**

For future services, the STS-1 may not have enough capacity to carry some services. SONET offers the flexibility of concatenating STS-1s to provide the necessary bandwidth. Consult the Glossary for an explanation of concatenation. STS-1s can be concatenated up to STS-3c. Beyond STS-3, concatenation is done in multiples of STS-3c. Virtual tributaries can be concatenated up to VT-6 in increments of VT-1.5, VT-2, or VT-6.

**Payload Pointers**

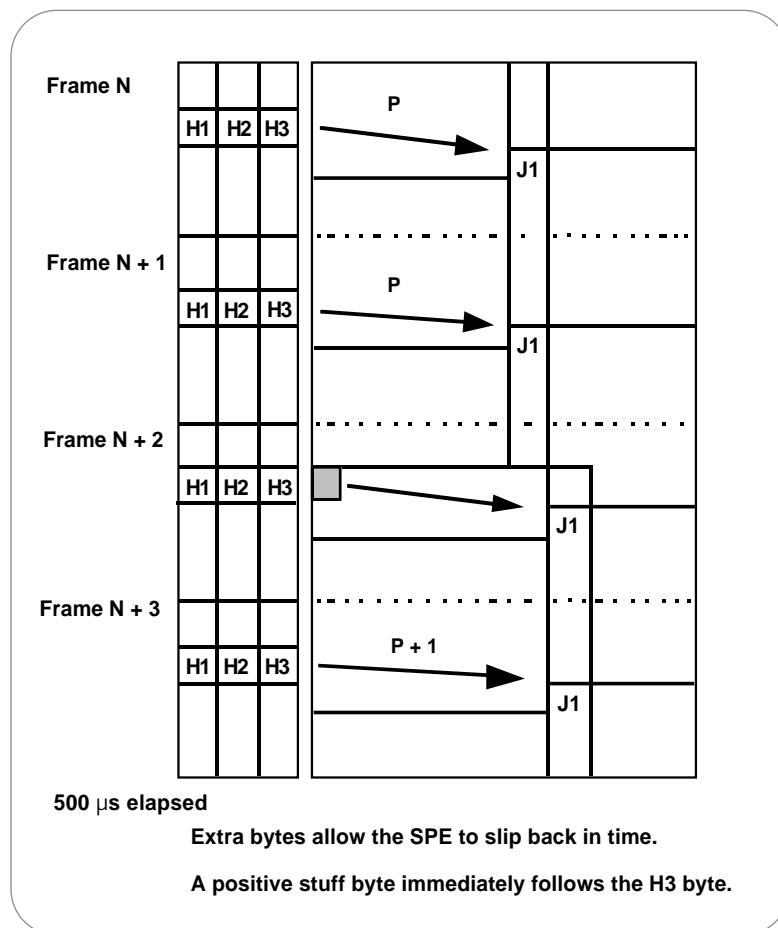
When there’s a difference in phase or frequency, the pointer value is adjusted. To accomplish this, a process known as byte stuffing is used. In other words, the SPE payload pointer indicates where in the container capacity a VT starts, and the byte stuffing process allows dynamic alignment of the SPE in case it slips in time.



▶ **Figure 10.** Pointer – SPE position in the STS-1 frame.

## Positive Stuffing

When the frame rate of the SPE is too slow in relation to the rate of the STS-1, bits 7, 9, 11, 13, and 15 of the pointer word are inverted in one frame, thus allowing 5-bit majority voting at the receiver. (These bits are known as the I-bits or Increment bits.) Periodically, when the SPE is about one byte off, these bits are inverted, indicating that positive stuffing must occur. An additional byte is stuffed in, allowing the alignment of the container to slip back in time. This is known as positive stuffing, and the stuff byte is made up of non-information bits. The actual positive stuff byte immediately follows the H3 byte (that is, the stuff byte is within the SPE portion). The pointer is incremented by one in the next frame, and the subsequent pointers contain the new value. Simply put, if the SPE frame is traveling more slowly than the STS-1 frame, every now and then “stuffing” an extra byte in the flow gives the SPE a one-byte delay. See Figure 11.



► **Figure 11.** Payload pointer – positive justification.

## Negative Stuffing

Conversely, when the frame rate of the SPE frame is too fast in relation to the rate of the STS-1 frame, bits 8, 10, 12, 14, and 16 of the pointer word are inverted, thus allowing 5-bit majority voting at the receiver. (These bits are known as the D-bits, or Decrement bits.) Periodically, when the SPE frame is about one byte off, these bits are inverted, indicating that negative stuffing must occur. Because the alignment of the container advances in time, the envelope capacity must be moved forward. Thus, actual data is written in the H3 byte, the negative stuff opportunity (within the Overhead); this is known as negative stuffing.

The pointer is decremented by one in the next frame, and the subsequent pointers contain the new value. Simply put, if the SPE frame is traveling more quickly than the STS-1 frame, every now and then pulling an extra byte from the flow and stuffing it into the Overhead capacity (the H3 byte) gives the SPE a one-byte advance. In either case, there must be at least three frames in which the pointer remains constant

before another stuffing operation (and therefore a pointer value change) can occur. See Figure 12.

## Virtual Tributaries

In addition to the STS-1 base format, SONET also defines synchronous formats at sub-STS-1 levels. The STS-1 payload may be subdivided into virtual tributaries, which are synchronous signals used to transport lower-speed transmissions. The sizes of VTs are shown in Table 8.

In order to accommodate mixes of different VT types within an STS-1 SPE, the VTs are grouped together. An STS-1 SPE that is carrying Virtual Tributaries is divided into seven VT Groups, with each VT Group using 12 columns of the STS-1 SPE; note that the number of columns in each of the different VT types – 3, 4, 6, and 12 – are all factors of 12. Each VT Group can contain only one size (type) of Virtual Tributary, but within an STS-1 SPE, there can be a mix of the different VT Groups.

For example, an STS-1 SPE may contain four VT1.5 groups and three VT6 groups, for a total of seven VT Groups. Thus, an SPE can carry a mix of any of the seven groups. The groups have no overhead or pointers; they're just a way of organizing the different VTs within an STS-1 SPE.

Since each of the VT Groups is allocated 12 columns of the Synchronous Payload Envelope, a VT Group would contain one of the following combinations:

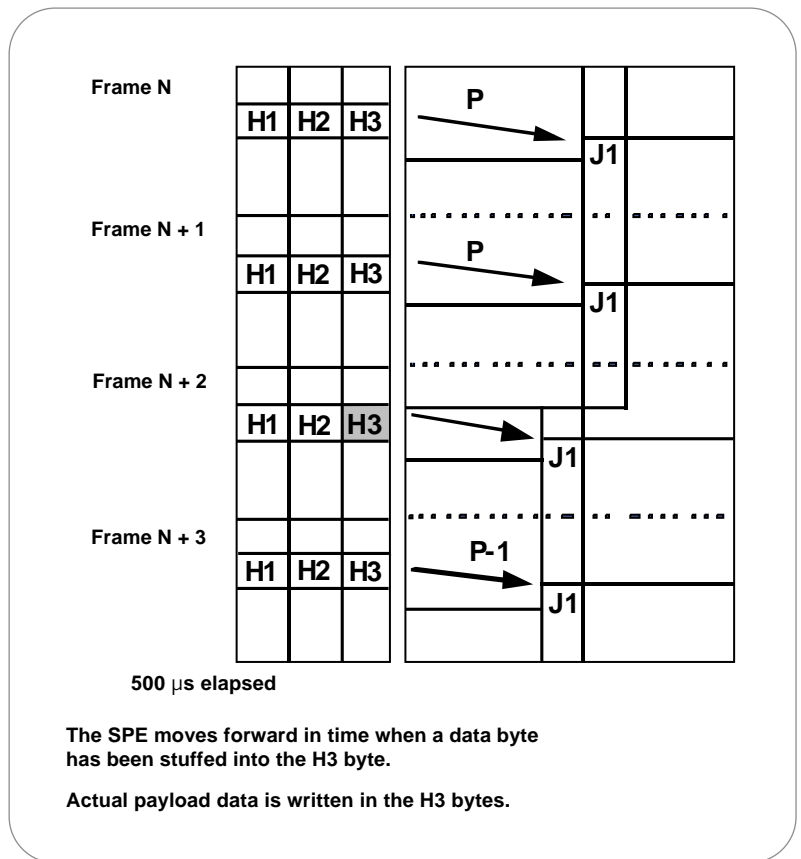
- ▶ Four VT1.5s (with 3 columns per VT1.5)
- ▶ Three VT2s (with 4 columns per VT2)
- ▶ Two VT3s (with 6 columns per VT3)
- ▶ One VT6 (with 12 columns per VT6)

The 12 columns in a VT Group are not consecutive within the SPE; they're interleaved column by column with respect to the other VT groups. As well, column 1 is used for the Path Overhead; the two columns of "fixed stuff" are assigned to columns 30 and 59.

The first VT Group, called Group 1, is found in every seventh column, starting with column 2, and skipping columns 30 and 59. That is, the 12 columns for VT Group 1 are columns 2, 9, 16, 23, 31, 38, 45, 52, 60, 67, 74, and 81.

**Table 8. Virtual Tributaries (VT)**

VT Type	Bit Rate	Size of VT
VT1.5	1.728 Mb/s	9 rows, 3 column
VT2	2.304 Mb/s	9 rows, 4 columns
VT3	3.456 Mb/s	9 rows, 6 columns
VT6	6.912 Mb/s	9 rows, 12 columns



▶ **Figure 12.** Payload pointer – negative justification.


# SONET Telecommunications Standard

► Primer

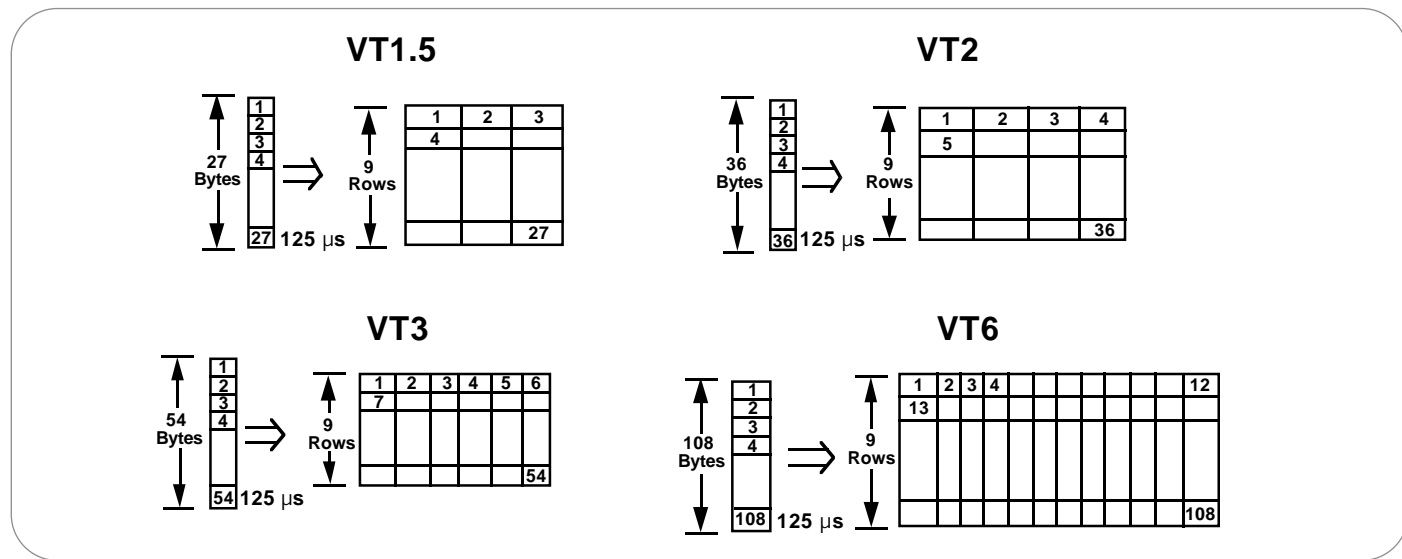
Just as the VT Group columns are not placed in consecutive columns in an STS-1 SPE, the Virtual Tributary columns within a group are not placed in consecutive columns within that group. The columns of the individual VTs within the VT Group are interleaved as well. See Figure 13.

The VT structure is designed for transport and switching of sub-STS-1 rate payloads. There are four sizes of VTs: VT1.5 (1.728 Mb/s), VT2 (2.304 Mb/s), VT3 (3.456 Mb/s), and VT6 (6.912 Mb/s). In the 87 column by 9 row structure of the STS-1 SPE, these VTs occupy columns 3, 4, 6, and 12, respectively.

To accommodate a mix of VT sizes efficiently, the VT structured STS-1 SPE is divided into seven VT groups. Each VT group occupies 12 columns of the 87 column STS-1 SPE, and may contain 4 VT1.5s, 3 VT2s, 2 VT3s, or 1 VT6. A VT group can contain only one size of VTs; however, a different VT size is allowed for each VT group in an STS-1 SPE. See Figure 14.



► **Figure 13.** SONET tributaries – VT structured STS-1 SPE.



► **Figure 14.** VT structure, VT sizes.

**STS-1 VT1.5 SPE Columns**

One of the benefits of SONET is that it can carry large payloads (above 50 Mb/s). However, the existing digital hierarchy can be accommodated as well, thus protecting investments in current equipment. To achieve this capacity, the STS Synchronous Payload Envelope (SPE) can be sub-divided into smaller components or structures, known as Virtual Tributaries (VT) for the purpose of transporting and switching payloads smaller than the STS-1 rate. All services below the DS3 rate are transported in the VT structure. Figure 15 shows the VT1.5 structured STS-1 SPE. Table 9 matches up the VT1.5 locations and the STS-1 SPE column numbers, per the Bellcore GR-253-CORE standard.

**DS1 Visibility**

Because the multiplexing is synchronous, the low-speed tributaries (input signals) can be multiplexed together but are still visible at higher rates. An individual VT containing a DS1 can be extracted without demultiplexing the entire STS-1. This improved accessibility improves switching and grooming at VT or STS levels.

In an asynchronous DS3 frame, the DS1s have gone through two levels of multiplexing (DS1 to DS2; DS2 to DS3) which include the addition of stuffing and framing bits. The DS1 signals are mixed somewhere in the information-bit fields and cannot be easily identified without completely demultiplexing the entire frame.

Different synchronizing techniques are used for multiplexing. In existing asynchronous systems, the timing for each fiber-optic transmission system terminal is not locked onto a common clock. Therefore, large frequency variations can occur. "Bit stuffing" is a technique used to synchronize the various low-speed signals to a common rate before multiplexing.

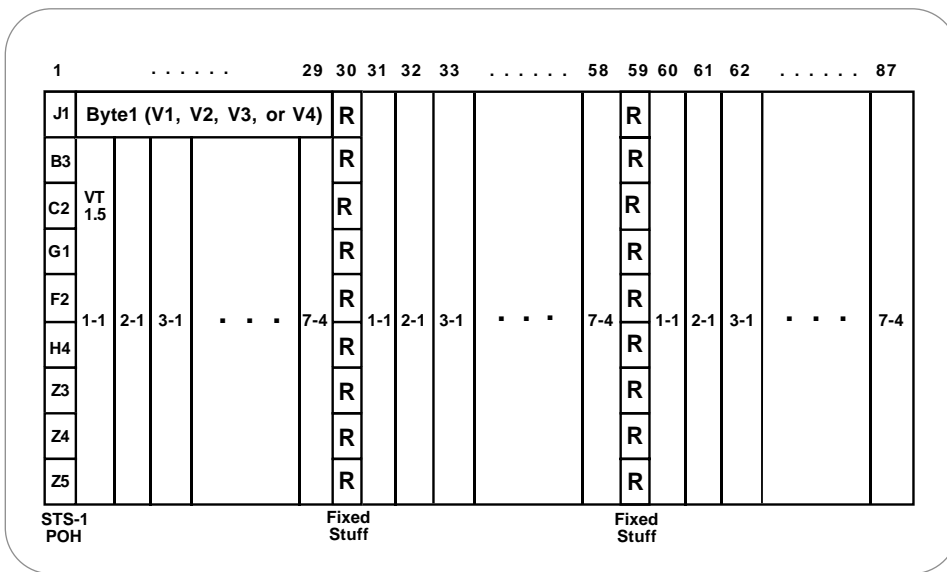
**Table 9. VT1.5 Locations matched to the STS-1 SPE Column Numbers**

VT Group #, VT #	Column #s
1,1	2,31,60
2,1	3,32,61
3,1	4,33,62
4,1	5,34,63
5,1	6,35,64
6,1	7,36,65
7,1	8,37,66
1,2	9,38,67
2,2	10,39,68
3,2	11,40,69
4,2	12,41,70
5,2	13,42,71
6,2	14,43,72
7,2	15,44,73
1,3	16,45,74
2,3	17,46,75
3,3	18,47,76
4,3	19,48,77
5,3	20,49,78
6,3	21,50,79
7,3	22,51,80
1,4	23,52,81
2,4	24,53,82
3,4	25,54,83
4,4	26,55,84
5,4	27,56,85
6,4	28,57,86
7,4	29,58,87

Column 1 = STS-1 POH

30 = Fixed Stuff

59 = Fixed Stuff



▶ **Figure 15. STS-1 VT1.5 SPE Columns.**

# SONET Telecommunications Standard

► Primer

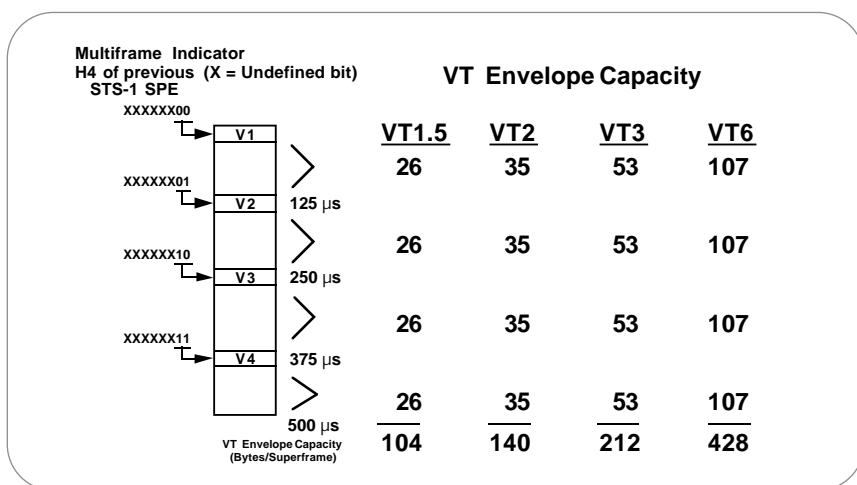
## VT Superframe and Envelope Capacity

In addition to the division of VTs into VT groups, a 500-microsecond structure called a VT Superframe is defined for each VT. The VT Superframe contains the V1 and V2 bytes (the VT Payload Pointer), and the VT Envelope Capacity, which in turn contains the VT SPE. The VT Envelope Capacity, and therefore the size of the VT SPE, is different for each VT size. V1 is the first byte in the VT Superframe, while V2 through V4 appear as the first bytes in the following frames of the VT Superframe, regardless of the VT size. See Figure 16.

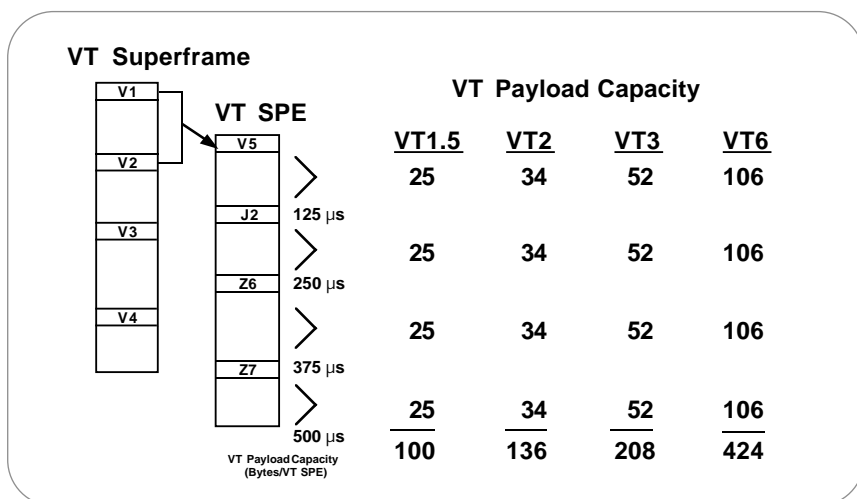
## VT SPE and Payload Capacity

Four consecutive 125-microsecond frames of the VT-structured STS-1 SPE are organized into a 500-microsecond superframe, the phase of which is indicated by the H4 (Indicator) byte in the STS POH.

The VT Payload Pointer provides flexible and dynamic alignment of the VT SPE within the VT Envelope Capacity, independent of other VT SPEs. Figure 17 illustrates the VT SPEs corresponding to the four VT sizes. Each VT SPE contains four bytes of VT POH (V5, J2, Z6, and Z7), and the remaining bytes constitute the VT Payload Capacity, which is different for each VT.



► **Figure 16.** VT superframe and envelope capacity.



► **Figure 17.** VT SPE and payload capacity.