SS8 Networks
One Better

SS7 Tutorial

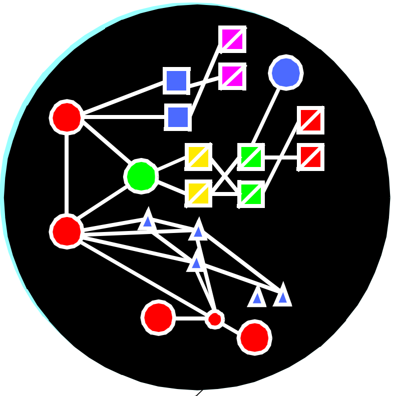

User Parts

OSI Layers

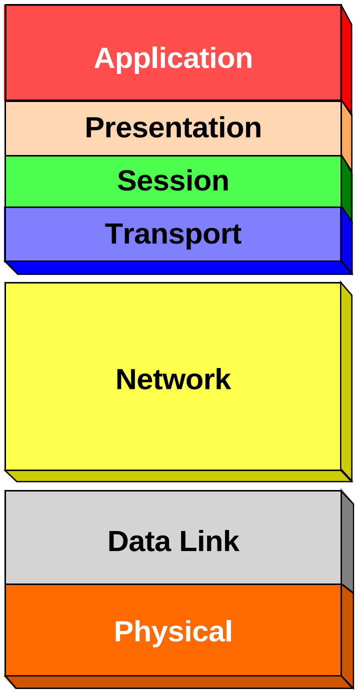

SS7 Functional Layers

The Layered Model

To understand the SS7 protocol, it is helpful to understand the concept of protocol layering. By 1983 some of the major Telecom companies had begun to realize the numerous problems that were developing because computers of numerous types were attempting to communicate with each other over network connections. They decided to sit down and create some specific interfaces that could be used by all.

In the process, they quickly became aware that by creating specific interfaces they might be slowing the development and implementation of future standards and future computer technologies. So, instead of creating specific interfaces, they decided to create a model of a layered architecture which could be used to develop future networks. The result became the Open Systems Interconnection (**OSI**) model later adopted by the International Standards Organization (**ISO**).

A layered protocol consists of modular programs, each designed to perform certain groups of functions, and each designed to be able to offer its functionalities to other modules. Thus, each module became a “part” of the whole protocol architecture, and each part which offers functional services to another part became a “user part”.

SS7 Layers

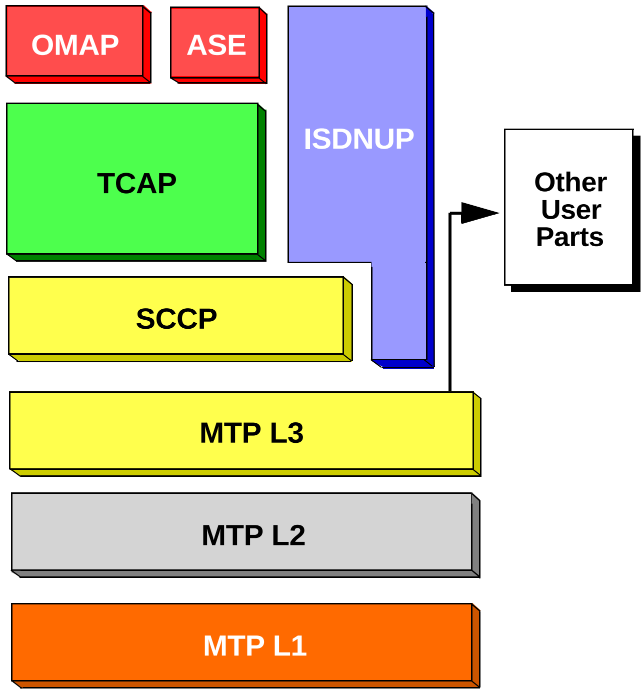

SS7 Functional Layers

The Layered Model

There are numerous benefits to this approach. One example lies in the fact that all networks have common needs. For example, every network needs to be able to transfer it messages between network locations over the connections used in that network. In the SS7 network these connections are called “links”. The message transfer must take place in some reliable fashion that guarantees message integrity and somehow prevents messages from becoming lost. The function used to accomplish this is generally known as a “transport mechanism”.

Because the transport mechanism is modular, new functionality in the network need not be concerned with how the messages will be transported. To add a new “user part” to an old network becomes a simpler matter because the new user part has only to develop its own functionality, and then to make use of the other functionalities it needs, including message transport.

Likewise, protocol conversions and combining of networks becomes easier because messages of one network type (for example X.25) can be transported in another network (for example SS7) using the network services (such as message transport) available there.

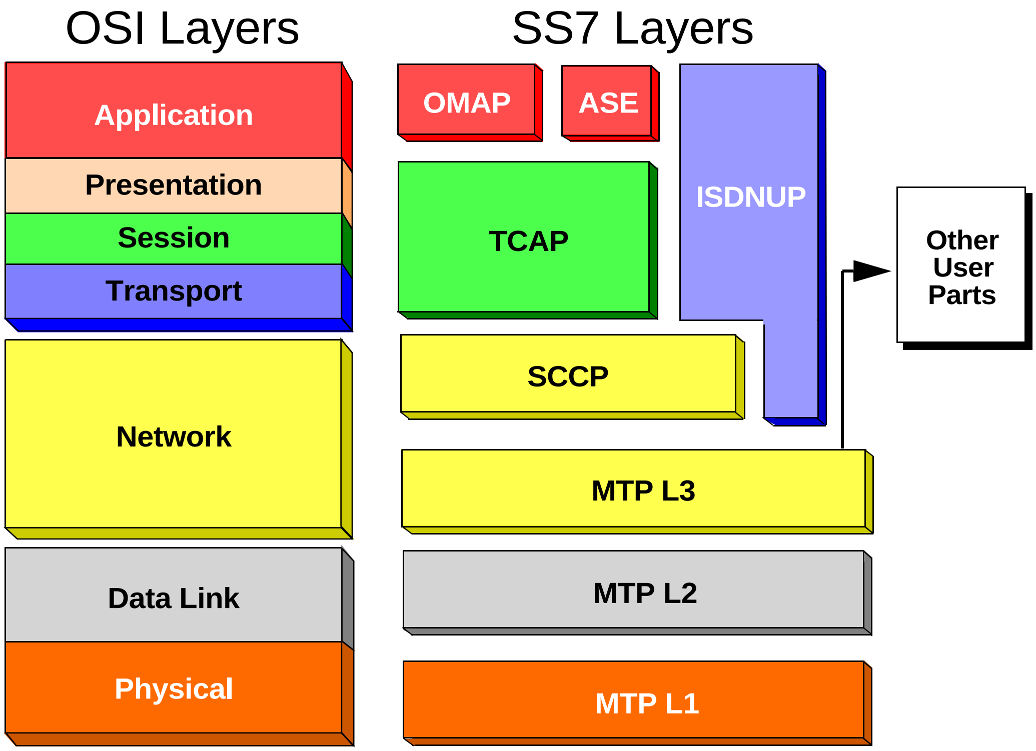

SS7 Functional Layers

The Layered Model

It would be difficult to find a protocol which incorporates all of the seven layers of functionalities grouped in exactly the same way as the seven layers of the OSI model. The SS7 protocol is no exception. However, for those who may be familiar with the OSI model it can be helpful to show the OSI alongside the unfamiliar protocol for reference. The comparison of OSI and SS7 is given here for that purpose.

Note that, in this drawing, while most of the lower layers incorporate the functionalities of the OSI directly, some of the functionalities of the upper layers become more mixed. For example the ISDNUP (Integrated Services Digital Network Users Part) extends all the way from the network layer to the application layer. Other User Parts not shown here include the Telephone Users Part (**TUP**) and the Data Users Part (**DUP**) which, while still currently in use, have generally fallen out of favor except in certain parts of the world (**e.g** China).

By the way, since most protocol architectures are shown with the layers "stacked" in order of the way they depend on each other reading from the top down, the protocol layers are often referred to as a "stack".

It is beyond the scope of this tutorial to fully investigate the OSI model. Those who need more complete information will find it at such locations as http://whatis.com/osi.htm .

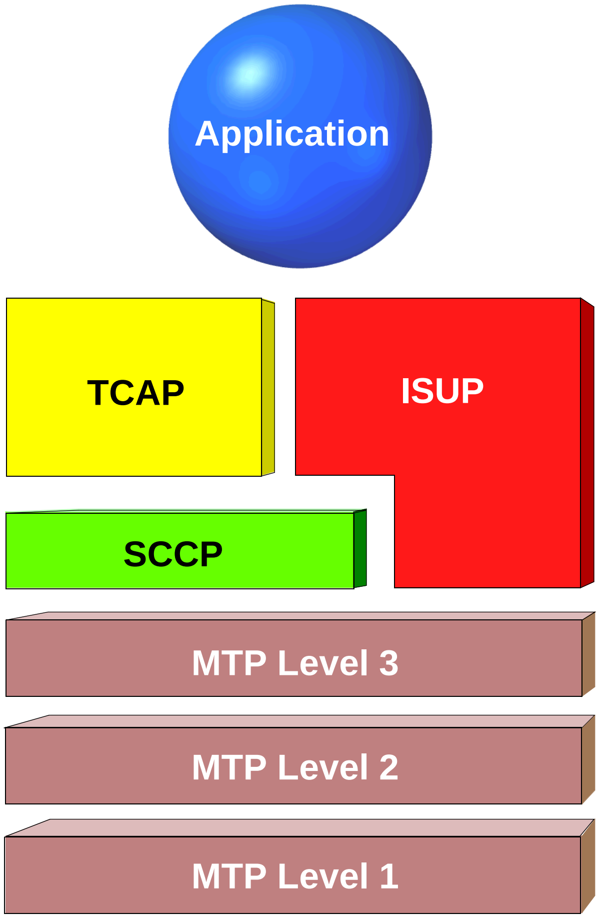

SS7 Functional Layers

The SS7 Stack

The User Parts (functional layers) illustrated here are not all of the parts of the SS7 protocol. The earlier drawing, for example, also made at least oblique reference to TUP and DUP.

That drawing also showed the OMAP (Operations, Maintenance and Administration Part) and the ASE (Application Service Element). These two are at the application level and as such perform duties which are quite different from the message handling portions of the stack. We'll finish our discussion by coming back to applications.

The OSI model recommends that communications between layers be kept as simple as possible. This is done to prevent loses in performance brought about by inter-layer connections. The SS7 handles this requirement through the use of "primitives" which are simple codes passed between layers to identify the services required of the receiving layer. These are characterized as requests, indications, responses and confirmations.

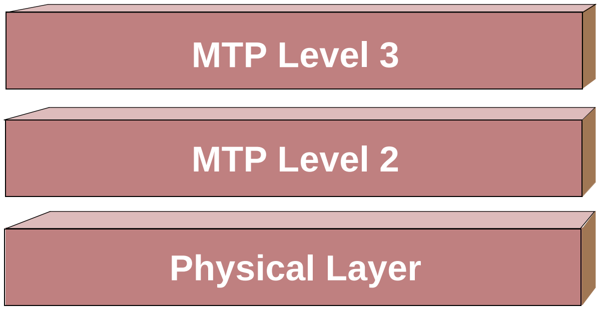

SS7 Functional Layers

The Message Transport Layers

MTP Level 1

We'll begin the examination of the layers at the lowest levels. These are the levels which are the first to handle an incoming message and the last to handle an outgoing one.

You'll note from the drawing above that MTP (Message Transfer Part) level 1 represents the physical layer. What this means is that this is the layer that deals with connecting the computer on which the programmed functionalities run, into the network with which it needs to communicate. MTP level 1 considers the links, the control of clocking, and all of the physical considerations of sending messages over wires. These are concerns of Electrical Design Engineers who need to consider what types of transmission lines may be available in the network infrastructure (T1/E1 etc.) This layer, then, consists of hardware. It is of little interest to those software designers who need to create stacks and applications to communicate with the network. Primarily software engineers will be concerned with level 1 only to the extent that they understand what is necessary to interface and to properly interact with the physical layer.

Interface cards are available for telecom infrastructures all over the world. The SS7 standard has variations and variants to allow the SS7 to work whether the infrastructure is T1, E1, DS0 or any one of a number of other possibilities.

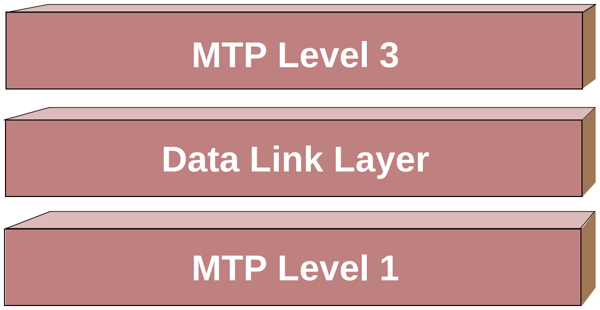

SS7 Functional Layers

The Message Transport Layers

MTP Level 2

At the next level, the SS7 becomes intelligent. That is, from this point on the layers become the concern of the software engineers. MTP layer level 2 performs the last intelligent handling of a message before it is sent out on the links. Likewise it performs the first intelligent handling of messages received from the network.

Because level 2 is so closely associated with the links, one of the main tasks assigned to it is that of link monitoring. Monitoring link congestion is a prime concern. Strictly speaking, it really isn't possible to congest a link. A link will carry only as much traffic as the transmission lines allow (for example, 64,000 bits per second). Congestion really refers to the buildup of messages in queues. Queues are limited in size, of course. Eventually messages being delivered to queues which are not emptying at an equal rate will be lost

The SS7 has numerous mechanisms for the handling of congestion. But MTP level 2 simply monitors and reports, it doesn't control. That job is left up to the network layer at MTP level 3.

Another job for level 2 is the assembly of outgoing messages into packets known as signalling units. There are three packets, all of which share common housekeeping characteristics. The packet which is used to include SS7 messages is called the Message Signal Unit. Another packet is used only to carry information about the links. This packet is called the Link Status Signal Unit. A final packet is used to make sure that there are no gaps in transmission and that there is recognizable data on the link as long as the link is operating. This is called the Fill In Signal Unit.

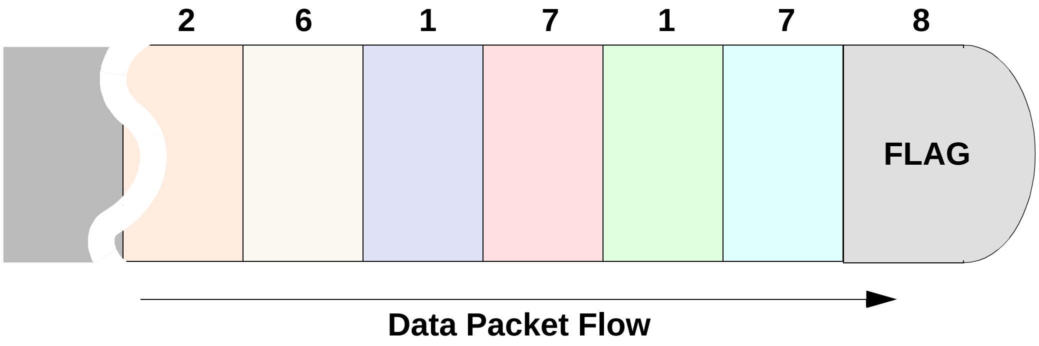

SS7 Functional Layers

The Message Transport Layers

MTP Level 2 (Continued)

In any stream of data it is important for the receiving end to recognize where to start reading a packet. Because of this the MTP level 2 has the task of applying a “start reading here” code to each packet. ANSI networks use the code at the beginning of very packet, while ITU networks may use it at both the beginning and end of each packet. The code or “flag” is defined as a byte with zeros at both ends and 1s in the middle.

Signal units can be quite large (279 octets). The odds of the same 01111110 code being repeated somewhere are, therefore, quite high. If that were to occur, then the reading sequence would be terminated (**ITU**) or reinitiated (**ANSI**). To prevent such occurrences in outgoing messages, the code stream is examined to find anyplace where there are five sequential 1s. Then, in a procedure known as “bit stuffing”, a zero is placed after the fifth 1. Finally, the flag is attached and the message is transmitted. On the receiving side, the MTP checks the incoming messages. After discarding the flag the MTP reads the stream looking for five 1s in sequence. After such a sequence, a zero is removed from the next bit position and the code has been restored to the original.

MTP level 2 has additional responsibilities in regard to message integrity. It needs to validate the integrity of received messages and to correct any messages that are received in unreadable condition. This cannot be done by the receiving MTP without the assistance of the transmitting side. To understand this it is necessary to understand the signalling packets and the interactions of the sending and receiving MTPs in assembling and disassembling the packets.

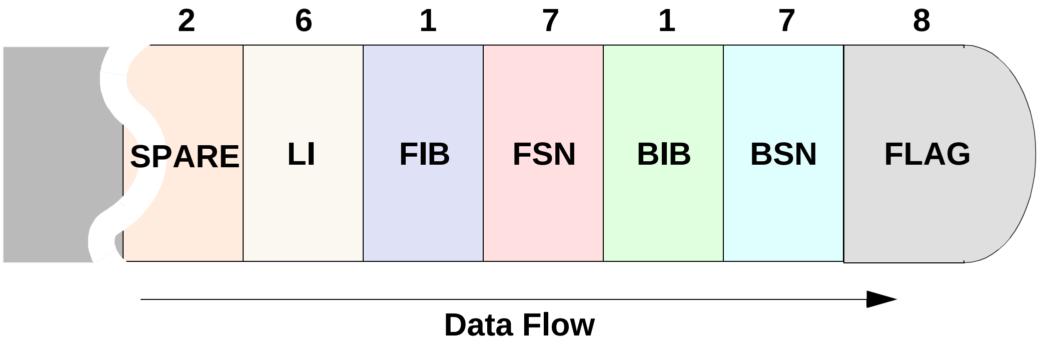

SS7 Functional Layers

The Message Transport Layers

MTP Level 2 (Continued)

The drawing above illustrates the common fields of all three packets, sometimes called the “housekeeping” fields. Their purpose is to provide the MTP with the information it needs to correctly read the packet, to ensure that the data has not been corrupted, to request copies of any received message whose data has been compromised, and to allow the receiving MTP to acknowledge the receipt of packets which are intact.

The arrow in the drawing shows the direction of the flow of data so you will understand that it is the flag that is seen first by the receiving MTP. The next field is called the Backward Sequence Number (**BSN**). As we will see, it is changed by the receiving MTP when it sends a packet back to the transmitting MTP. The Backward Indicator Bit (**BIB**) is also changed by the receiving MTP in the packet it returns to the transmitting MTP when a copy of the message is requested. The Forward Sequence Number (**FSN**) is the field used by the transmitting MTP to place the cyclic value (0-127) which identifies each packet. The Forward Indicator Bit (**FIB**) is set by the transmitting MTP and examined in returned acknowledgment packets to determine whether the packet is a request for a copy. The numbers over the fields indicate the number of bits in each field.

Messages are transmitted from queues that are link specific. Each queue is really two queues. From one the messages are sent directly, so it is often called the “transmit buffer”. At the same time a message is sent, it is also copied into a second buffer commonly called the “retransmit buffer”. With its number applied for reference (known as the Forward Sequence Number) the packet waits in the retransmit buffer until informed by the receiving side MTP (Acknowledgment) that the message has been received intact. When this Acknowledgment comes, the copy is destroyed along with any in the buffer that had been sent previously.

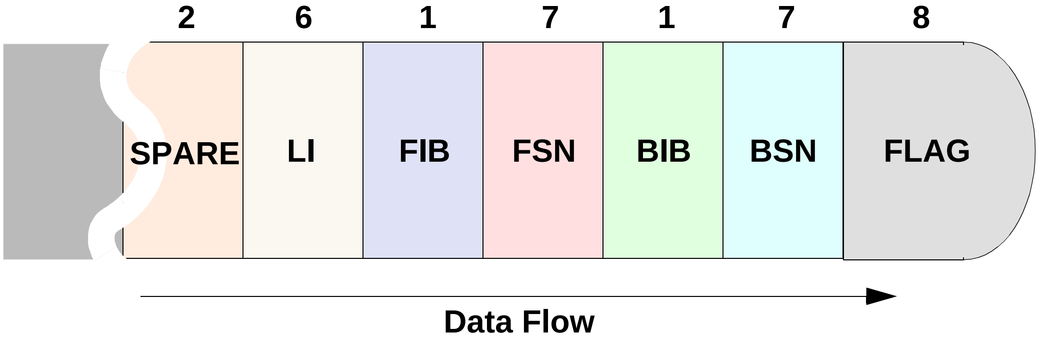

SS7 Functional Layers

The Message Transport Layers

MTP Level 2 (Continued)

Even though you are not yet familiar with the methods used by the MTP to determine message integrity, we'll trace the interaction between the sending and receiving sides to see how the integrity of the packet is guaranteed. First, the sending side places a value (0-127) in the Forward Sequence Number field of the message it transmits. It also retains the state of the last Backward Indicator Bit received from the receiving side. When the message goes out, this value (1 or 0) is in the BIB and the Forward Indicator Bit is set to be the same. The message as transmitted then contains a numbered value and the BIB and FIB are set either both to zero or both to one. The sending side transmits the message and sends a copy to its retransmit buffer.

Every time the receiving side sees a good message, it retains the Forward Sequence Number and the state of the Forward Indicator Bit. It does this for good messages received because the same information may not be readable in a garbled message. When a garbled message is received, the receiving side uses the information saved from the last good message to request a copy. It does this using a different packet. Typically the MTP is sending link status information periodically using a packet called the Link Status Signal Unit. The MTP will use the next packet to be sent back as the "envelope" to return the negative acknowledgment and this will often be the Link Status Signal Unit

In the return packet, the MTP uses the saved FSN of the last good message as the BSN of this packet. It restores the FIB and then places a BIB (Backward Indicator Bit) in the packet which is in the toggled state to the FIB. That is, if the FIB is a zero, the MTP places a one in the BIB. If the FIB is a one, the MTP places a zero in the BIB. When the transmitting side receives the packet, it reads the value of the last good message received (which is now in the BSN). It also examines the FIB and BIB to see if they are alike or toggled to be opposite each other (0/1 or 1/0).

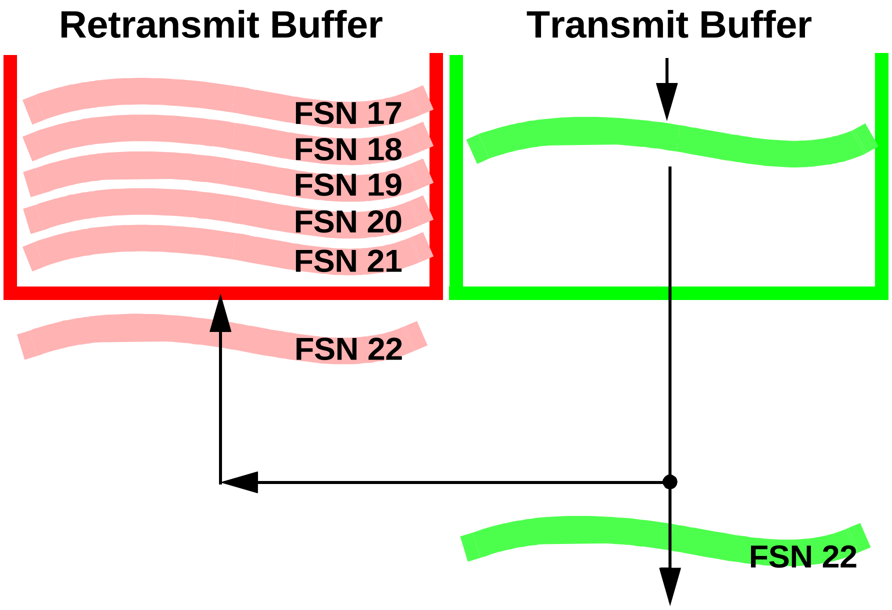

Pre-error Message Transmission

SS7 Functional Layers

The Message Transport Layers

MTP Level 2 (Continued)

The drawing illustrates the way in which the transmitting side fills in the Forward Sequence Number and copies the message to the Retransmit buffer before sending it. Let us assume that the message with Forward Sequence Number 19 was not received in good order. The receiving MTP will take the number of the last good message it saw (in this case 18) and place it in the Backward Sequence Number field of the packet it will be returning to the transmitting side. It also will ensure that the Backward Indicator Bit is set to the opposite state of the Forward Indicator Bit. This serves to request retransmission.

When the packet is received at the transmitting side the difference in the FIB and the BIB trigger a series of actions:

1. Transmission is stopped.

2. Messages beginning with FSN #19 are retransmitted in sequence.

3. Messages which were in the buffer before FSN #19 are deleted.

4. Transmission begins again

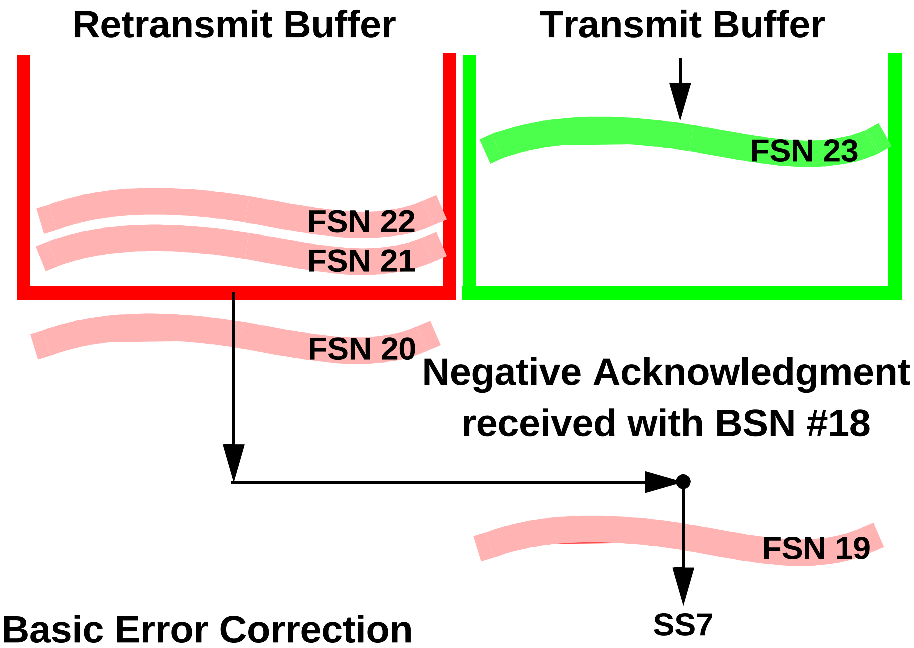

Basic Error Correction

SS7 Functional Layers

The Message Transport Layers

MTP Level 2 (Continued)

Using this approach, the receiving side is guaranteed of the integrity of the messages it receives, and the transmitting side is allowed to clean out the Retransmit Buffer, thereby avoiding congestion.

The receiving side must return responses even when all messages have been received accurately. It does so in a fashion nearly identical to the way in which it requests message copies. On a periodic basis (which it must do to convey Link Status information anyhow) the packet it sends contains a BSN which is the FSN of the last good message received. This time, the Backward Indicator Bit (**BIB**) is set to the same value as the Forward Indicator Bit (**FIB**). This is seen by the transmitting side as a positive acknowledgment, allowing the transmitting MTP to destroy all message copies up to and including the FSN indicated in the packet.

All of this packet traffic occurs over links, and messages relating to specified buffers are transferred only over the related links. This could mean that if a link failed utterly (perhaps unplugged from its port) all this positive and negative acknowledgment wouldn’t occur. In such cases other buffer overflow mechanisms may be needed. One of these is a timer (**T7**) which is set awaiting acknowledgment. If the timer times out, transmission can be stopped on that link. Another is that the local MTP will report a failure whenever there is no data at all on a link.

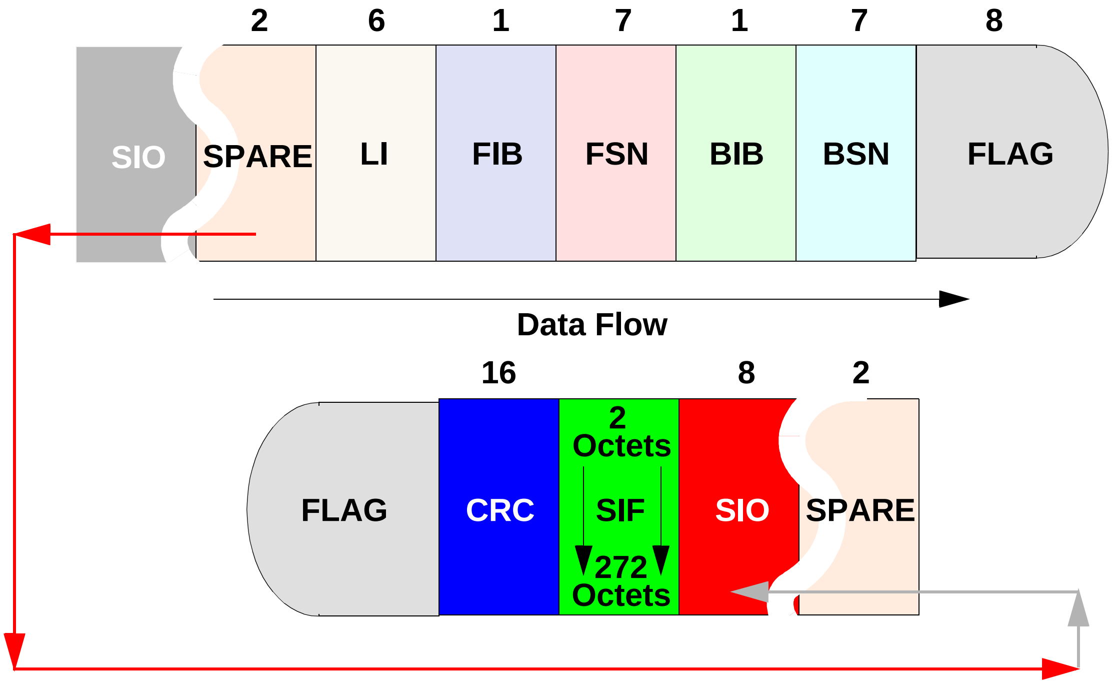

SS7 Functional Layers

The Message Transport Layers

MTP Level 2 (Continued)

All of the preceding is in reference to a set of procedures called **Basic Error Correction**. What we have shown here is how the process applies to handling non-intact Message Signal Units, which are the packets used to transmit all SS7 messages. Using further rules such as what happens to the FSN, BSN, BIB and FIB when transmitting other packets (Fill In Signal Units and Link Status Signal Unites) the standards also provide the means of detecting errors in the non-message packets. We have not covered them here because they fall outside of the scope of this Tutorial. There is another error correction approach called Preventive Cyclic Retransmission, used largely in Satellite communications, which also falls outside that scope. We'll only say that this method requires that all data be continuously retransmitted until acknowledgments are received.

Detecting Errors

Errors are detected using a variety of tests. To begin with, note that the fields of the Signal Units (a Message Signal Unit is shown above with all fields) are always an even number of octets. That is, if the total number of bits in the packet is divided by 8, the remainder is always 0. If the result of division by 8 results in anything but a 0 remainder, the packet is in error and can be discarded.

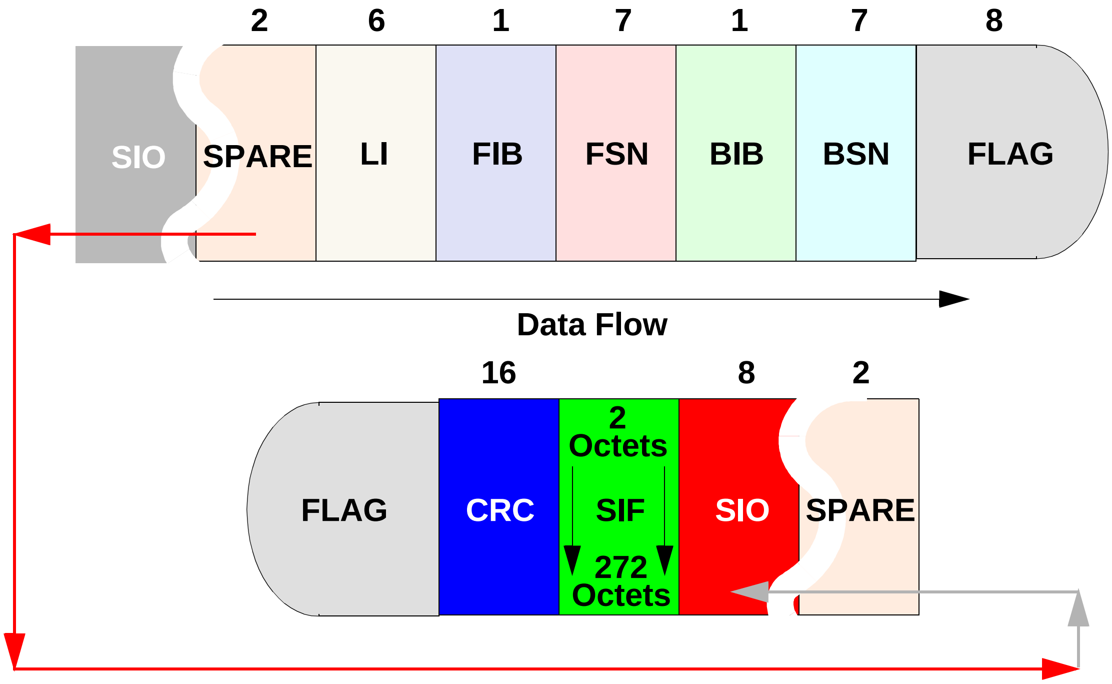

SS7 Functional Layers

The Message Transport Layers

MTP Level 2 (Continued)

In addition, note the field labeled **CRC.** This field contains a code (the **C**yclic **R**edundancy **C**ode) which is a representation of the total number of bits placed in the Message Signal Unit by the transmitting side. It is derived by algorithm so as not to lend itself to misinterpretation. At the receiving end this value is checked to determine if it is a match to the number of bits actually counted by the receiving MTP. If it matches, the message is deemed to be correct; and if it doesn't match it is certainly a tainted message.

Throughout all of this testing and validating the MTP is tolerant of transmission problems. They are accepted as a part of the process. MTPl2 watches not only the messages, but also the links themselves. Two procedures are employed to determine the health of the links. One of these is used during the process which places a link in service. This process is known as the *alignment* procedure and is employed when the link is first started up, and every time the link has been taken out of service and is about to be reactivated.

During the procedure the MTP uses a counter called the **A**lignment **E**rror **R**ate **M**onitor. Since the standards provide for two different proving periods the AERM is generally configurable (by the user) to either of these two choices.

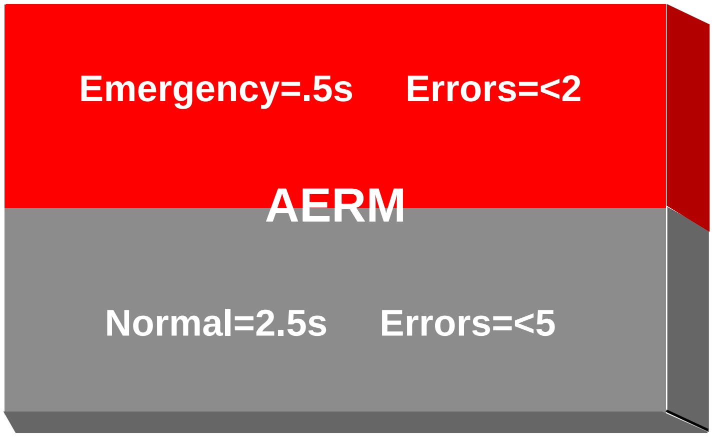

Alignment Error Rate Monitor

SS7 Functional Layers

The Message Transport Layers

MTP Level 2 (Continued)

One of the choices is referred to as *Normal* alignment and the other is referred to as *Emergency* alignment The standards specify the amount of lost data per time that is permitted. Since this is rate related, it might be easier to use the actual figures for a specified rate. In the drawing the rate chosen was 64Kbs.

After some early establishment of the link connection, the MTP enters a proving period. During this period, Fill In Signal Units are sent to be monitored. If the configuration calls for Emergency alignment, the FISUs **(F**ill **I**n **S**ignal **U**nits) are monitored for .5 seconds. The Fill In Signal Unit carries no data. That is, remove the SIO and SIF fields from an MSU and you have a FISU. During the proving period a single error will be ignored. However, if 2 errors occur, the link is not brought to service and the alignment procedure starts all over again.

During the alignment procedure MTP Level 2 is also reporting what it is doing with the link. It begins by sending "O" (Out of Alignment) in a LSSU **(L**ink **S**tatus **S**ignal **U**nit). This packet contains all the same data as an MSU except for where the MSU carries messages (**SIF**) and service information (**SIO**). Instead of these two fields, the LSSU has a one or two byte field used to convey link status information

When it enters the proving period it sends "N" or "E" (Normal or Emergency). It also sends "OS" (Out of Service) to indicate the link should not be used.

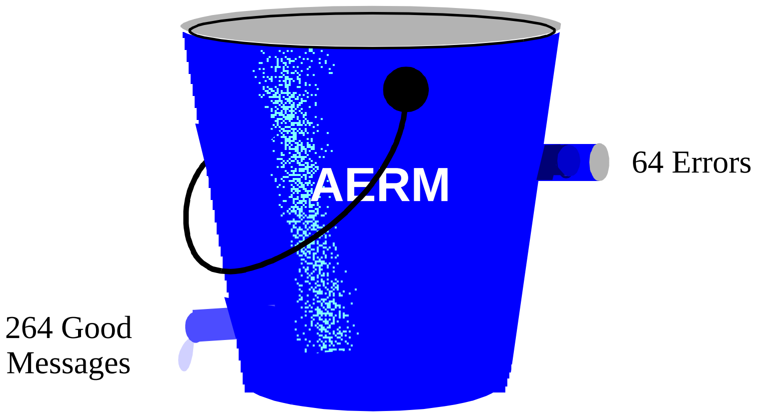

Signal Unit Error Rate Monitor

SS7 Functional Layers

The Message Transport Layers

MTP Level 2 (Continued)

If the configuration calls for Normal alignment, the FISUs (**F**ill **I**n **S**ignal **U**nits) are monitored for 2.5 seconds. During that time 4 errors will be ignored. However, if 5 errors occur, the link is not brought to service and the alignment procedure starts all over again

Establishing a proving period works fine for alignment, but the MTP must also monitor the links during normal transmission. Since this is continuous, monitoring must also be continuous.

Monitoring during normal transmission makes use of another monitor known as the Signal Unit Error Rate Monitor (**SUERM**). Since time is not the issue, this monitor is a simple counter. It counts the number of units in error by incrementing the counter. If the counter reaches a count of 64, the link is removed from service and realigned. But, since rate is normally a matter of events per time, how can a rate be established using only an incremental count?

The answer lies in what is normally referred to as the leaky bucket technique. Every error increments the counter, but 256 error-free signal units causes the counter to decrement. This “leak in the bucket” hopefully prevents the bucket from becoming full enough to remove the link from service.

*Trail flag not used in U.S.

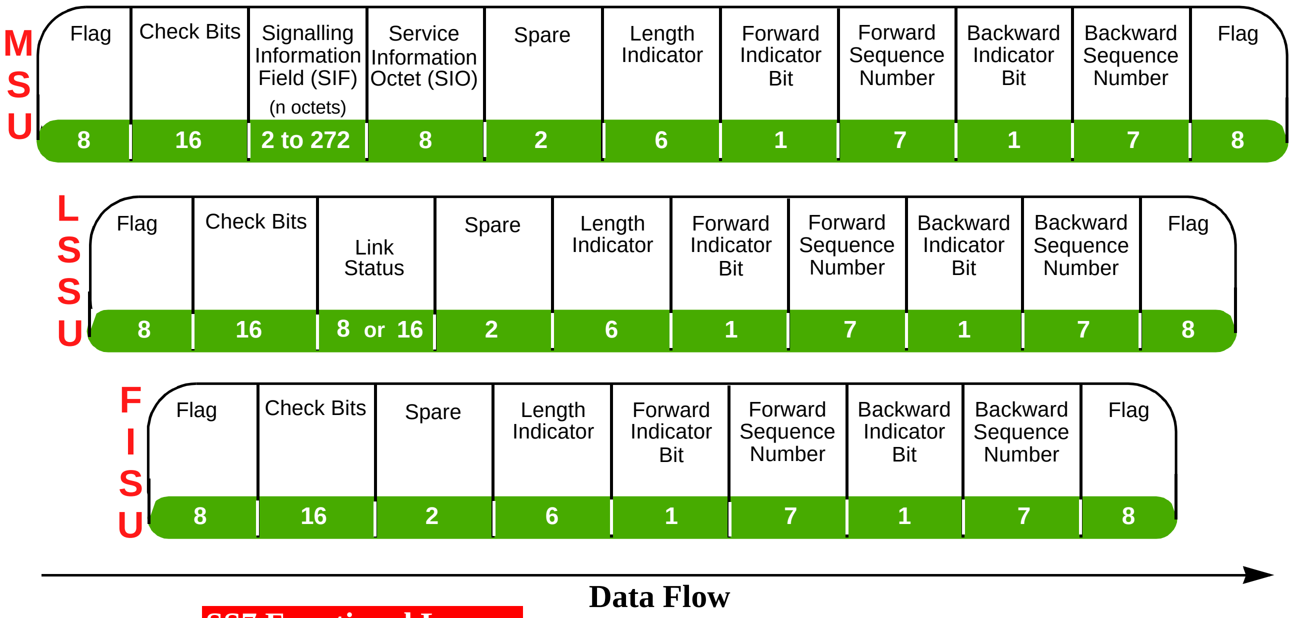

MSU

Flag	Check Bits	Signalling Information Field (SIF) (n octets)	Service Information Octet (SIO)	Spare	Length Indicator	Forward Indicator Bit	Forward Sequence Number	Backward Indicator Bit	Backward Sequence Number	Flag
8	16	2 to 272	8	2	6	1	7	1	7	8

LSSU

Flag	Check Bits	Link Status	Spare	Length Indicator	Forward Indicator Bit	Forward Sequence Number	Backward Indicator Bit	Backward Sequence Number	Flag
8	16	8 or 16	2	6	1	7	1	7	8

FISU

Flag	Check Bits	Spare	Length Indicator	Forward Indicator Bit	Forward Sequence Number	Backward Indicator Bit	Backward Sequence Number	Flag
8	16	2	6	1	7	1	7	8

Data Flow

SS7 Functional Layers

The Message Transport Layers

MTP Level 2 (Continued)

Before we move on to MTP level 3, let’s take one last look at the three types of signal units used to transport the information.

The **M**essage **S**ignal **U**nit (**MSU**) contains a Service Information Field to carry the data identifying the type of data contained in the packet. This may be data related to an ISUP or TCAP application. It may also be network management or testing information. The same field can be used to provide a priority value to the packet. The signaling information itself, such as IAMs, ANMs etc. (ISUP message types we’ll cover in the messages section), is contained within the Signaling Information Field (**SIF**), preceded by a routing label.

The **L**ink **S**tatus **S**ignal **U**nit (**LSSU**) has a single field in place of the SIF and SIO (it is called the Link Status field). In that field, one or two bytes of data identify congestion, current alignment states, etc.

The Fill In Signal Unit has no field in place of the SIF or SIO. It is used as a filler between MSU or LSSU transmissions so that the receiving MTP can always see activity on the links. It is also used by MTP to perform error checking on links that are being aligned

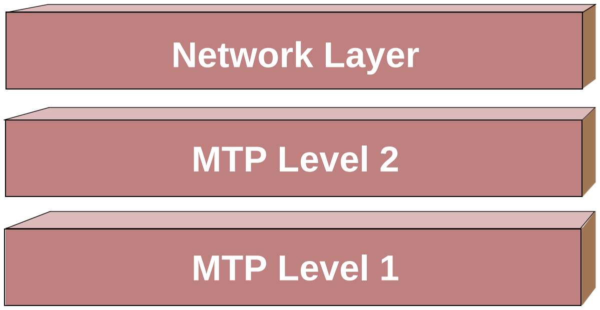

SS7 Functional Layers

The Message Transport Layers

MTP Level 3 - **S**ignalling **N**etwork **M**anagement

The functionality of MTP Level 3 is broken down into two groupings. One of those groupings has to do with where the MTP send the messages it receives and is referred to as **S**ignalling **M**essage **H**andling (**SMH**). The second grouping deals with how MTP Level 3 handles the traffic, the links, and the routes available to it. This group is referred to as **S**ignalling **N**etwork **M**anagement (**SNM**).

Since we have just finished a discussion of the activities of MTP Level 2 which provides much of the information required for network management, we'll discuss Signalling Network Management first. MTP level 3 continually receives information (from the local Level 2 as well as from remote nodes in the network) about the status of links, linksets, routes and routesets. A major part of the job of Level 3 is to make use of this information to control traffic on the traffic resources over which it has control. This means, for example, that when one its links becomes congested, it will direct that traffic normally intended for the congested link will now be redirected to a non-congested link.

MTP Level 3 sees it resources as being available or unavailable. There are many reasons a link may become unavailable. A link which has failed, become congested, been labeled as inhibited or blocked, or is not currently activated is seen by Level 3 as unavailable. Under normal circumstances, MTP Level 3 is generating a code (the Signalling Link Selection or **SLS**) which it will use in a round robin fashion to distribute the load across available links. When it receives an indication that a link is unavailable, it changes the destination of new messages to another link in the same linkset. This is a procedure known as **Changeover**. When the condition which resulted in the Changeover clears, the MTP is no longer using the available links efficiently. To change that, it must now direct a proportionate amount of traffic back to the restored link. This is a procedure known as **Changeback**

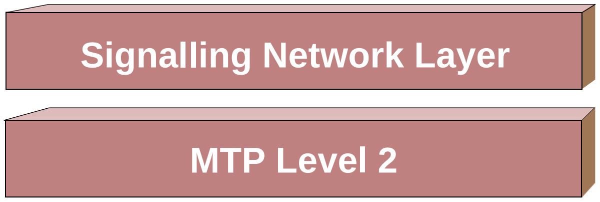

SS7 Functional Layers

The Message Transport Layers

MTP Level 3 - Signalling **N**etwork **M**anagement

In addition to handling link traffic, MTP Level 3 must handle route traffic. It may be that traffic directed to a remote destination cannot get there because a route (one of the MTP's own linksets) becomes unavailable. When this occurs the route table is consulted to see if an alternate route to the same location is available. If such is the case the outgoing route is changed in a process called **Forced Rerouting**. When the original route again becomes available, the MTP Level 3 at the originating node buffers the outgoing data, changes back to the original route, transmits the buffered data, and resumes normal transmission. This procedure is referred to as **Controlled Rerouting**.

The MTP may receive information from the network that some destination has become unavailable. This allows the MTP, once again , to consult a routing table to determine if alternate locations are possible; and, if such is the case, to reroute traffic to these alternates

It should be noted that the MTP Level 3 always returns to the original link/linkset traffic configuration whenever the condition which required a change in the traffic configuration has been cleared.

All of this activity of redirecting traffic in response to link and network conditions is referred to as **S**ignalling **N**etwork **M**anagement (**SNM**).

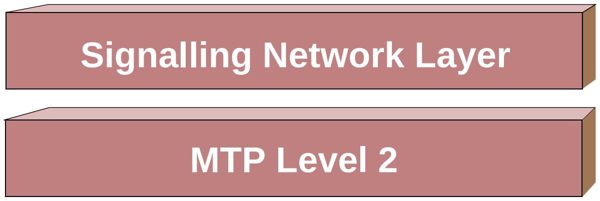

Signalling Network Layer

MTP Level 2

SS7 Functional Layers

The Message Transport Layers

MTP Level 3 - Signalling **N**etwork **M**anagement

Still another Signalling Network Management task for which MTP Level 3 is responsible is called **MTP Restart**. If a node should lose all of its access to the network, special considerations need to be made when the linksets have been restored. During the absence of the node from the network, conditions may have changed. For example, destinations which were available at the instant of the shutdown of the node may no longer be available The node which is reentering the network needs time to update its network condition information. However, nodes sending messages to the reentering node will see it as available and will begin to send messages immediately. The restarting node needs to prevent this. In addition, the restarting node wants to be sure that an adequate amount of resources (**e.g.** a majority of all of its links) are available before inviting new traffic.

To do this, the first outbound messaging will usually tell the transmitting nodes not to send messages. This can be done using such things as a Traffic Restart Waiting message followed by a Traffic Restart Allowed. Numerous other considerations are made in the restart process. For example, outgoing traffic may have been buffered before the shutdown and needs to be transmitted at restart.

The considerations also vary with the type of node which has been isolated. An STP, for example, needs to take steps to complete the routing of unrouted messages. Restarting may involve a fairly complex set of tasks designed to allow the node a graceful reentry to the network. On the other hand a full restart may not always be necessary for some nodes.

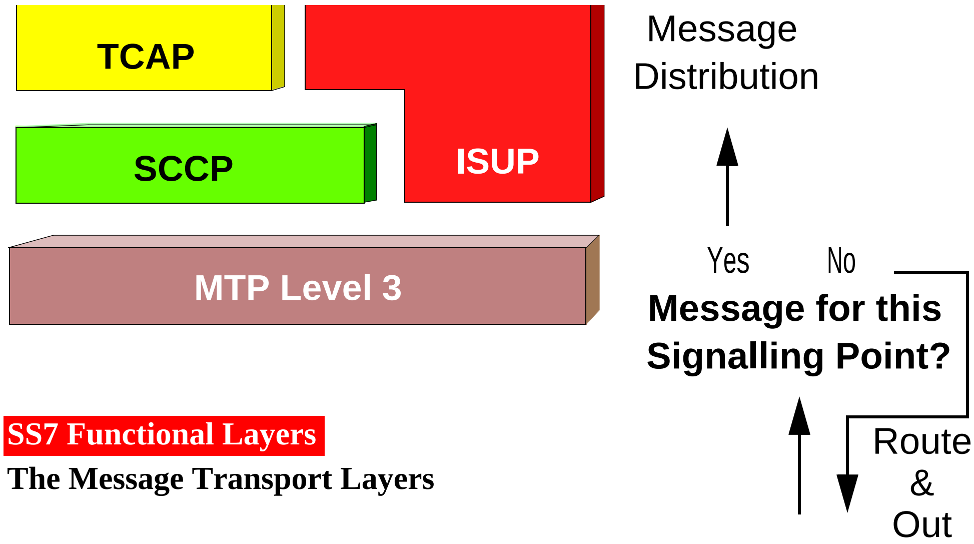

SS7 Functional Layers

The Message Transport Layers

MTP Level 3 - Signalling **Message Handling**

The other side of MTP Level 3 functionality is the **S**ignalling Message **H**andling (**SMH**). This can be further broken down into Message Discrimination and Message Distribution.

The drawing illustrates **Message Discrimination**. The message is examined to determine whether it is addressed to this node (Destination Point Code). If it isn't, the rule is "route and out". Since an STP bears major network routing responsibilities, this is a most invoked function at an STP. For other nodes, the function may be limited or even non-existent. Some nodes may be required to do some limited routing, particularly intra-network routing using F links. In other cases the node will have no routing requirement and will simply "dump" a message which is not addressed to it

If the message has arrived at the correct node then the Message Discrimination process delivers the message to the Message Distribution functionality. Most of the layers have no option with an incoming message other than to pass it up to the next User Part. For MTP Level 3, however, a selection must be made. This is because Level 4 can consist of several User Parts. SCCP and ISUP are the most commonly used, but there may be a TUP (**T**elephone **U**sers **P**art) or even **D**UP (Data **U**sers **P**art). Data found in the Service Information Octet (**SIO**) of the Message Signal Unit (**MSU**) will help Level 3 make the determination.

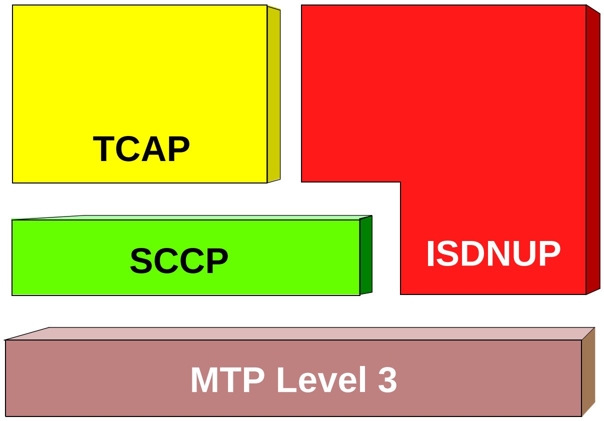

SS7 Functional Layers

The Message Transport Layers

MTP Level 3 - Other Services

Before moving on to other levels, it might help to get a clearer understanding of the term “User Parts”. Each layer, through its set of functionalities, offers services to other layers. Layer 3, for example, uses the services of Layer 2 to keep it informed of the status of links, linksets and routes. Layer 3, in turn, offers its services to Level 4.

One prime example of this lies in the use of the Signalling Link Selection Code provided by Level 3. This code is normally used by Level 3 to indicate the link on which each packet should be sent. Usually Level 3 keeps changing the code so that each packet gets sent on a different link and the traffic is evenly distributed across the available links.An additional bit in the ANSI standard is “rotated” (swapped front to back) to change linkset selections.

SCCP provides services of various classes. Two of these require that the data be delivered in sequence. There is only one way this can be guaranteed. If data is delivered over separate links, the Application at the adjacent (directly connected) node may not see the messages in the same sequence in which they were sent. The reason for this is that each link has its own message queues. The quantity of messages in these queues tends to fluctuate in operation with the result that, at any given time, it is unlikely that any 2 queues hold the same number of messages.

When requested to perform in-sequence delivery, the SCCP makes use of the SLS services of the MTP. When requested to do so, MTP Level 3 will “freeze” the SLS for the duration of a transaction. By delivering each message to the same link repeatedly, Level 3 assures that all messages end up in sequence in the message queues and cannot be delivered to the application out of order.

The Integrated Services Digital Network User Part Layer

ISUP Services

ISUP offers two types of services, known as **Basic** and **Supplementary**. Basic Services consist of those services employed in the process of setting up and tearing down a call. Supplementary Services consist of those services employed in passing all messages that may be necessary to maintain and/or modify the call.

ISUP functionality can be further broken down into three procedural categories. The first of these is **Signalling Procedure Control** (SPRC) which directly interfaces with the services of the MTP. The SPRC, in turn, provides support for **Circuit Supervision Control** (CSC) and for **Call Processing Control** (CPC). The application which deals with the circuit connection requirements of the switch, and simultaneously with SS7 signalling, is usually referred to as a Call Control application.

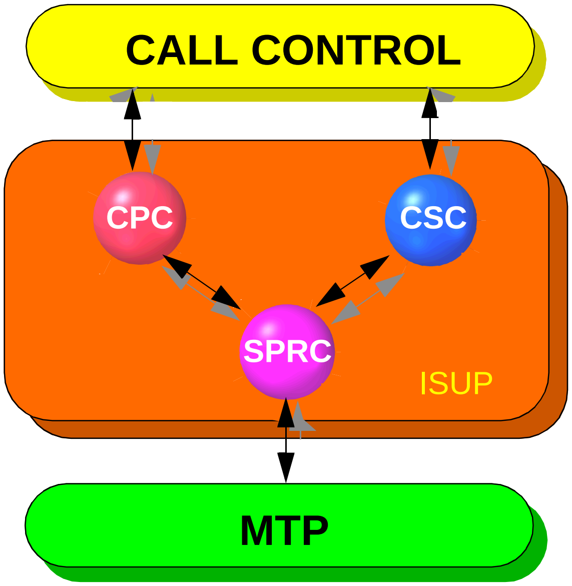

CPC = Call Processing Control CSC = Circuit Supervision Control

SPRC = Signalling Procedure Control

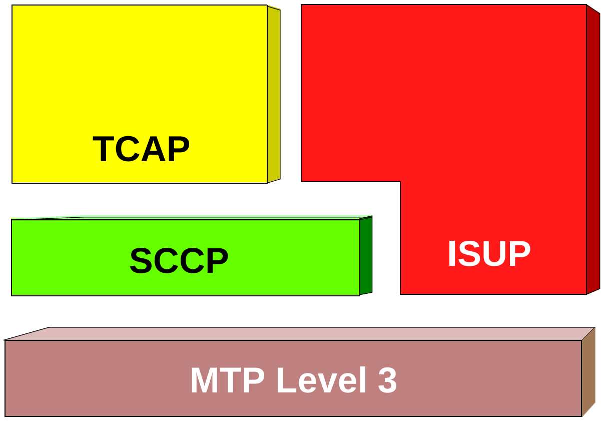

SS7 Functional Layers

The Integrated Services Digital Network User Part Layer

ISUP Services

While this User Part is really the ISDNUP, you will more often see only ISUP used. This is because the standard refers to all portions of PSTN (Public Switched Telephone Network) use (for example the line side) and not just to setting up and tearing down circuit connections. The ISUP term is generally used in reference to circuit related usage. The applications which have the job of controlling these circuit connections are referred to as “Call Control” applications.

ISUP conveys all the signalling information necessary to establish and maintain call connections. Each switch gets this information from the previous switch in the circuit as the connection is being established. Thus, ISUP messages move through the SS7 network from switching node to switching node parallel to the voice path being established.

The process begins with the originating switch program analyzing the dialed digits and consulting a routing table to determine which of the switches with which it shares trunks is the correct way to route the call. Then it selects a circuit in the trunk connecting to that location to which it will connect the line of the caller. It then constructs a message to the next switch to indicate the circuit it will use and to contain all the information which that switch will require to select the next connection.

The first message sent is an ISUP message called the **IAM** (Initial Address Message). This message contains all the information necessary for each switch to be able to consult its routing table and to select circuits which will result in connecting the circuit from end to end. The sender of the IAM receives a message confirming that the switch to which the IAM was sent is now in possession of all the required address information. This message is called the **ACM** (Address Complete Message).

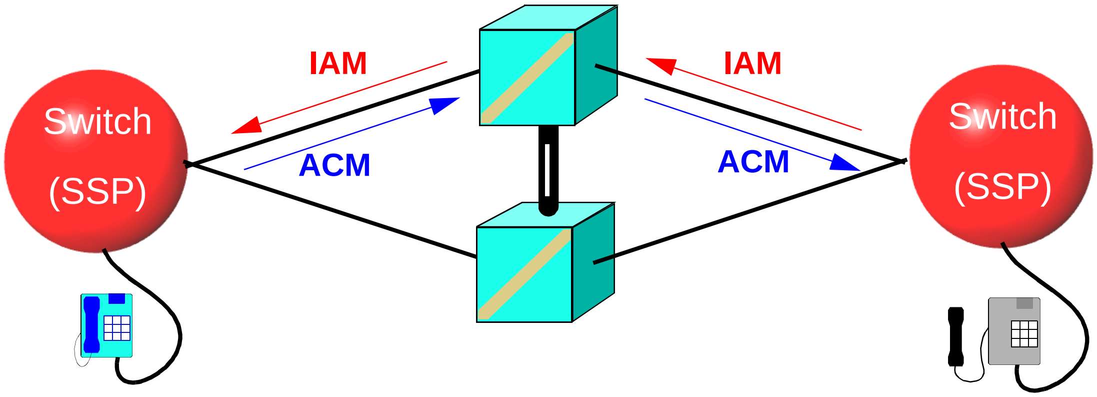

SS7 Functional Layers

The Integrated Services Digital Network User Part Layer

ISUP Services

At the destination End Office of the circuit, the Subscriber Line Interface is checked to determine the status of the line being called. If it is “on hook” the end office then sends a signal to start the phone ringing. At the same time it uses the SS7 to send an ACM signal backward into the network that results in “ringback” being sent to the caller. If the called party’s line is “off hook” the return to the caller is the “busy” signal.

Before the SS7 none of this signalling was done in digital packets. Instead, each switch connected its circuit and then sent the information about the call forward using the voice circuit to do so. Since switches were set up to deal with sound this signal was sent using a slightly smaller range of frequencies than that established for voice (a range of roughly 3200 cycles). This meant that this signalling was “In-(the voice)band. While several variations of analog signalling were tried, eventually the most prevalent was MF (multi-frequency). While much In-band signalling has disappeared, one form still exists largely because it exists on the “line side” to homes and businesses where digital signal generating equipment is usually not available. We all know it as touch tone; but to the telephone industry it is known as **DTMF** (Dual Tone Multifrequency).

While today’s switches can make connections very rapidly, the time to make a connection is not zero. When circuit connections had to be made to carry the signalling information, these connection times were added to by each switch in the circuit. Today, when a switch establishes a circuit it need not connect it. Instead, it can simply mark the circuit as “reserved”. If the response from the end of the circuit indicates that the phone is busy these reserved circuits can immediately be released. Then, no time has been lost either connecting or disconnecting.

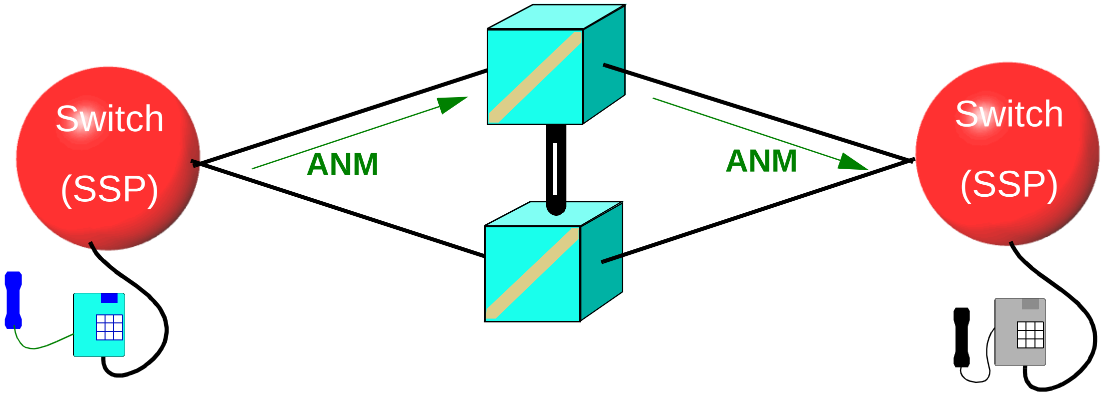

The Integrated Services Digital Network User Part Layer

ISUP Services

If the called phone rings, each switch can choose to make the connection. In such a case, connection times will not be additive because all switches in the circuit can be connecting at the same time. It is even possible for the switch to "reserve" until the phone is lifted off the hook. The connection can be made as soon as the receiver is lifted.

But, back to our basic call scenario. Once the phone is ringing, there is no further signalling being exchanged for a time. When the phone comes "off hook" the End Office sends an **ANM** (ANswer Message) backward into the SS7 network. Each switch is thus notified that the full circuit must now exist. Any switch which has not yet done so must now complete the connection. If this is a normal phone call (as opposed to such special case types as conference calls), a conversation now takes place and all switches do little more than maintain the connections.

Eventually, of course, someone hangs up. The phone line once again goes "on hook" and that is sensed at the subscriber interface of the End Office serving the customer who hung up. That office now sends a **REL** (release) message on to the previous switch in the circuit. Upon receiving the REL, each switch releases the circuit connection. At the same time, it returns an **RLC** (release complete) back to the switch which sent the REL. In this way, the switch confirms that it got the message about disconnecting. Switch by switch, the scenario of REL forward and RLC back continues until each switch in the circuit has released its circuitry and confirmed that action to the previous switch.

For all the User Parts, timers play an important role. ISUP employs numerous timers, many of which play a crucial role in making sure the network works smoothly. For example, when the REL is sent, the switch expects to receive an RLC in reply.

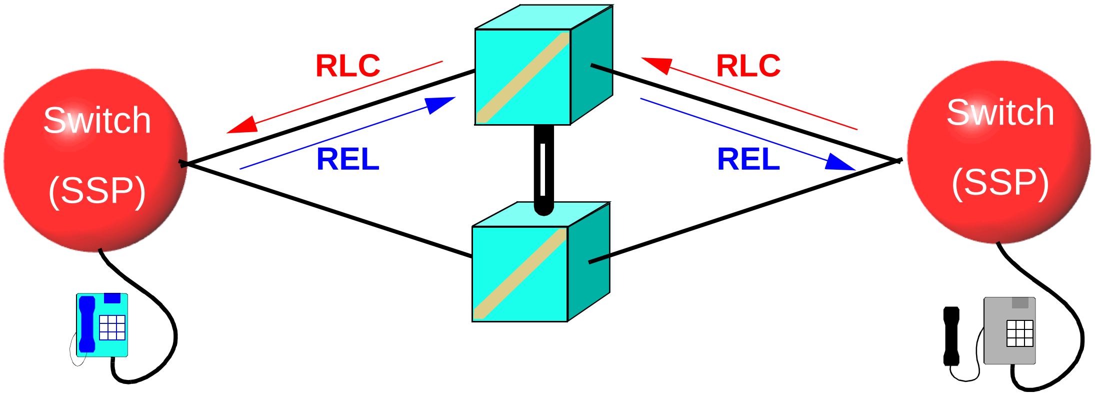

SS7 Functional Layers

The Integrated Services Digital Network User Part Layer

ISUP Services

If there is no return of an RLC, it might mean that the switch to whom the message was sent didn't get the message and it still maintaining the circuit. So the switch sending the REL also sets a timer. The timer will be disabled by a returning RLC. If not, the timer will time out, thereby alerting the switch application to a potential problem. ISUP uses many specific timers, most of which are set when a message is sent and terminated by the receipt of an anticipated response.

Switches along the circuit also use other timers which differ according to which office is setting the timer and to the responsibility that office has relative to the phone call. For example, one office needs to set a timer on receipt of the ANM and let that timer run until the REL is sent. The office which does this is the one which has "charge office" responsibilities for the call. It is this office which determines the length of the phone call.

ISUP is used not only to make simple phone call connections, but also to extend those services to **CLASS** (**C**ustom **L**ocal **A**rea **S**ignaling **S**ervices). For these services messages are sent to the switch directing it to take some "non-standard" (but rapidly becoming "standard") actions. For example, the switch can be directed that when a connection to one phone line is requested, the connection should be made to a different phone line. This can be predicated upon a timed period of allowing the original phone to ring before making the new connection or upon diverting the call to another phone without ever trying to connect to the original phone. This, of course, is call forwarding.

Other services falling generally under the heading of CLASS are Calling Number Delivery (Caller ID), Call Waiting, Automatic Callback, Selective Call Rejection, etc.

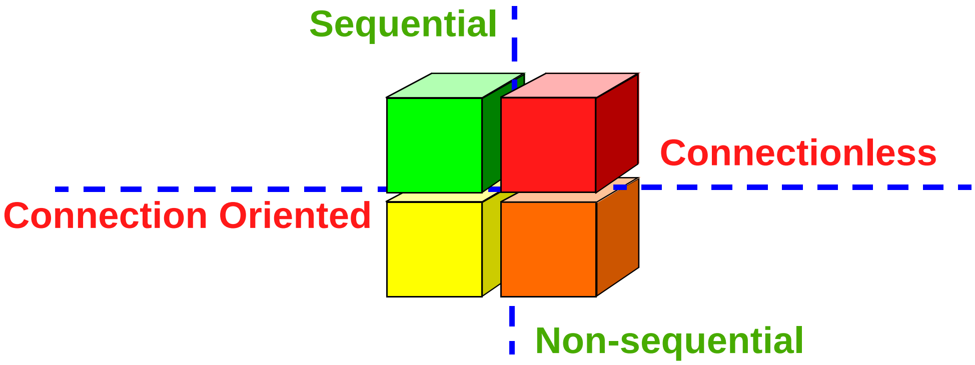

SS7 Functional Layers

The Signalling Connection Control Part Layer

SCCP Services - Connection Management

As the name implies, the SSCP deals with signalling connections. To be clear, SS7 connections are continually present in the physical sense. That is, they are not like voice circuits which need to be connected. So the “Connection” part of the name refers only to whether one set of rules regarding the transfer of data needs to be applied or not. The connection is a virtual connection. In a virtual connection, the locations at both ends of a linkset first communicate about establishing a connection before starting the transfer of the data which prompted the connection.

The rules of connection start with one location making a request of the other to establish the connection. The side receiving the request responds. Then there is a transaction in which the two sides provide and respond to information about the connection. Requests fall into certain category types and those types determine what information is exchanged next. At the very least, a connection identifier is established along with information about the type of data to be delivered while the connection is in place. All this occurs before any data is transferred. After the data transfer, the connection is released. This occurs in a pattern of messages concerning the release which resembles the release of a voice connection.

The SCCP also offers data transactions which occur without establishing a connection. In truth, the vast number of SCCP transactions occurring in real networks is connectionless.

All told the SCCP offers four classes of data transaction services. Connection oriented services fall into two sub classes. In one of these sub classes, data for the entire transaction is delivered over a single link to guarantee delivery in the same sequence as transmitted. The same two subclasses of Connectionless services exist as illustrated in the drawing.

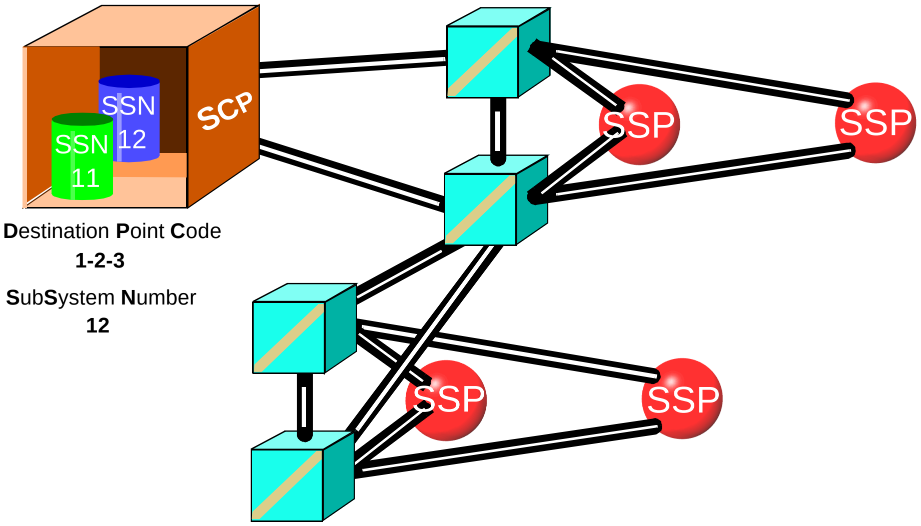

SS7 Functional Layers

The Signalling Connection Control Part Layer

SCCP Services - Extended Addressing

Largely the intent of the SCCP is to provide resources for data transfers that are not resources available at the MTP layer. Among these are special addressing capabilities. The MTP at any node is concerned only with getting the message to the next node. Because of this it uses only the Signalling Point Codes of the adjacent (link connected) nodes. Even when routing to a distant node MTP only considers which adjacent node should be sent the message.

SCCP, on the other hand, is concerned with end-to-end routing. To support this it uses multiple addressing mechanisms. It makes use of Signalling PointCode for all locations in the network, whether linked to its local node or not.

SCCP also confronts the problem of sub-addressing. For example, any given location may be home to more than one database. Addressing to the Signalling Point Code will only get a message to a specified node. With no further addressing there would be no way for that node to know which of its databases is being queried. SCCP resolves that by making use of identifiers of data services (including databases) known as **S**ub-**S**ystem numbers (SSN).

Subsystem numbers must be represented in the packet (**M**essage **S**ignal **U**nit) using a single byte of data to do so. For this reason, they range in value from 0 to 255. Some of the lower values are reserved. The drawing illustrates how any of the SSPs might address the SCP.

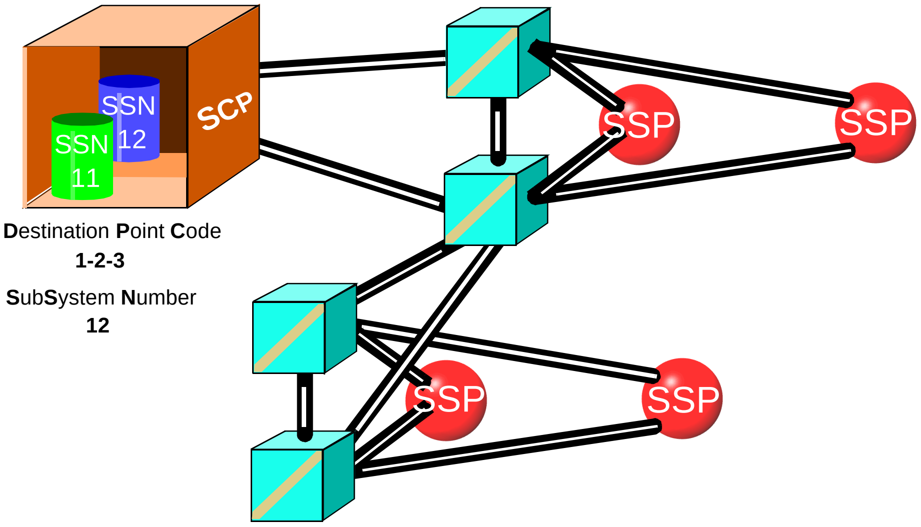

SS7 Functional Layers

The Signalling Connection Control Part Layer

SCCP Services - Global Title

In addition to **D**estination **P**oint **C**ode and **S**ub**S**ystem **N**umber addressing, the SCCP provides another means of directing queries. This means is called Global Title (**GT**).

Global Title addressing is essentially alias addressing. To send a message (**e.g**. an 800 number database query), the switch does not necessarily need to know where the database is or what its subsystem number may be. Instead, the switch simply needs to know a location (very often an STP) which does know where the information is to be found.

The switch sends a message coded as Global Title. Essentially the coding simply identifies the type of information that is required, and provides the data to work with. For example, when someone calls an 800 number, the switch will encode a message as global title, indicate that what is needed is a translation of the digits that were dialed into a normal (North American Numbering Plan) telephone number, and enclose the actual digits dialed. The switch sends this message to a location (probably STP) which is listed in its SS7 routing table as the place to send Global Title dialed digit translation requests.

Upon receiving the request, the Global Title location determines which of potentially numerous tables to consult, finds the Destination Point Code and/or subsystem number of the 800 database (SCP) and sends the query there. The database returns the answer to the Global Title location which returns it to the switch.

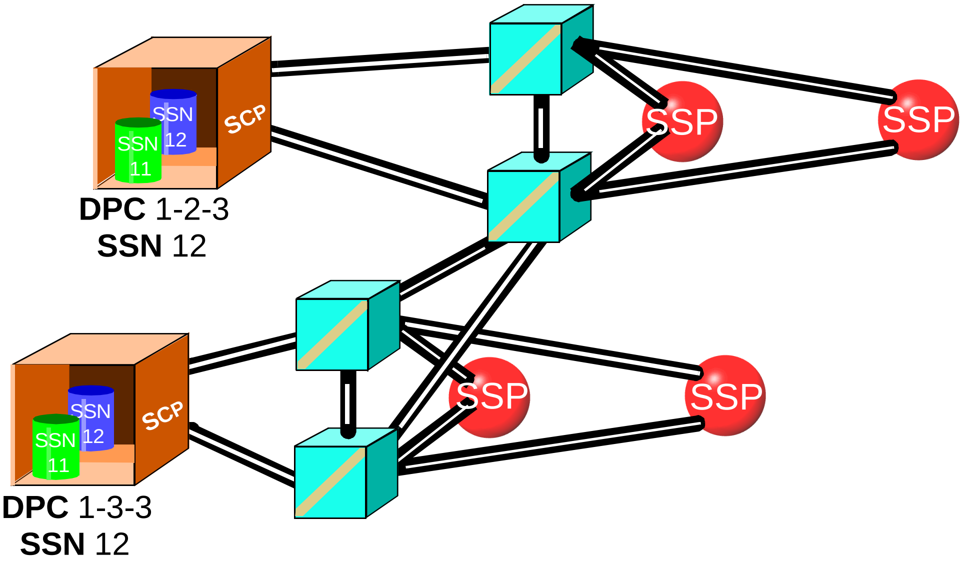

SS7 Functional Layers

The Signalling Connection Control Part Layer

SCCP Services - Global Title

The use of Global Title offers many advantages. For one thing, without Global Title every switch in the world would have to update its routing information to have access to any new database. Using Global Title, only the Global Title locations need to make table entries and then every switch which uses that Global Title location has immediate access to the new service.

Another advantage is that a network can hide services from the rest of the network. This provides the opportunity to maintain network proprietary services, or to offer the services for a fee.

Another way to use Global Title is to change the Global Title table lookup relative to network conditions. For example, a Global Title location may redirect queries away from a database which has been withdrawn from service and towards a backup database which contains the same information. The switch which sent the query is never aware of the change, and still receives the correct response.

The drawing illustrates how redundant databases might appear in a network. In the drawing the databases are at different locations and the subsystems are given the same number. Another way to handle redundancy would be to place both the database and its backup at the same location and use different subsystem numbers. Which brings us to the final SCCP functional grouping. SCCP is designed to handle redundant database management.

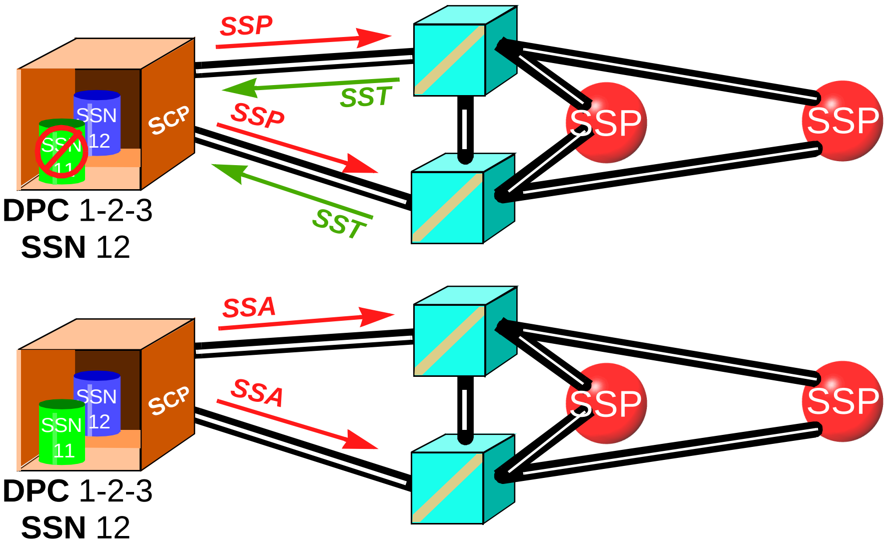

SS7 Functional Layers

The Signalling Connection Control Part Layer

SCCP Services - Database Management

The handling of redundant databases is another task performed by the SCCP. The process begins when one application (database) is scheduled to be removed from service. That application will use SCCP messaging to request permission of its peer for withdrawal. If the backup database application has no reason not to grant the request, it returns a grant for withdrawal and the withdrawing database begins the procedure.

Some locations in the network need to be kept informed of such withdrawals because they need to know how the queries should be routed while the prime database is down. To the database application, these locations are known as **C**oncerned **P**oint **C**odes (**CPC**). The SCCP at the database performing the shutdown first sends a signal to these locations warning them of the impending shutdown. This signal is called an **SSP** (**S**ub**S**ystem **P**rohibited).

The SCCP at the Concerned Point Code returns a request for status called an **SST** (**S**ubsystem **S**tatus **T**est). Upon receipt of this message, the application which is shutting down has confirmation that its message has been received, and it completes the shutdown process. The **SST** is sent by the concerned location and receives a response periodically. In the meantime, the **CPC** begins sending queries to the backup.

When the shutdown database is ready again, the SCCP at that location sends an **SSA** (**S**ub**S**ystem **A**llowed) to the Concerned Point code. If an SST is received in return, it means the SSA was not received; so it is sent until no SST is returned.

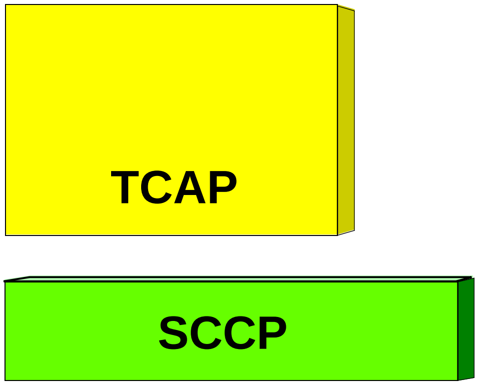

SS7 Functional Layers

The Transaction Capabilities Application Part

TCAP Services

From the very name of this part you will note a fundamental change. This is no longer referred to as a User Part. This is now an application part. The primary service offered to an application is the packaging of data.

Once you have absorbed that idea it becomes easier to deal with TCAP functionality because it doesn't resemble the functionality of any other user parts. The whole idea behind TCAP is to format and present data in various standardized formats which allows it to be used in multi-vendor environments.

TCAP and SCCP go together like blank and blank (you fill in the blanks). In the MSU, TCAP messages are found in the SCCP portion of the SIF. The SCCP provides the necessary Subsystem Number (SSN) addressing, handles deliveries in sequence when necessary (thanks to MTP level 3), provides segmentation services when the messages get overly long, and handles Global Title considerations. It does all of this while providing connectionless transport for TCAP.

Because TCAP messaging can invoke various actions, it is used with a broad variety of databases. But it can be used for other purposes such as invoking switch features. In short, TCAP is all about data and how to package it. Other parts, such as the Mobile Application Parts (GSM, IS41-C) sort of "piggy back" on TCAP messages. In general, they might be considered an extension of TCAP.

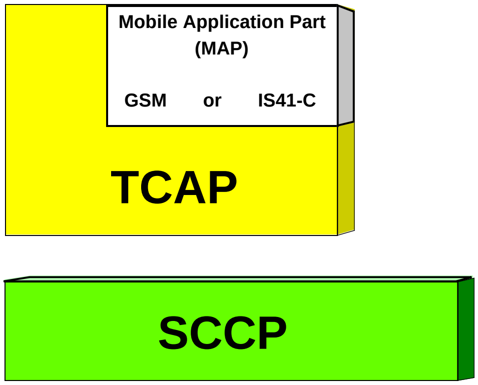

SS7 Functional Layers

The Transaction Capabilities Application Part

TCAP Services

When a TCAP query is sent, a number is attached to identify the transaction. This is because the switch will be sending queries and receiving responses in no particular order. The transaction identifier is copied into responses, allowing the switch to correlate the query and response.

TCAP coding makes use of discrete portions of the message known as components. The **Invoke Component** carries the specifics of the request and identifies the actions to be taken (**e.g.** dialed digits translation). The **Response Component** returns the requested data. A third component called the **Error Component** may be used when the answering database is having difficulty responding (**e.g** *You didn't send me enough digits*). A fourth component called the **Reject Component** may be sent when there is no way to comply with the request (**e.g**. *I don't have that information. What kind of a database do you think I am?*)

As you will see from the messaging section, TCAP, by its very structuring, presents a very "tight" protocol. That is, it would be difficult to miss an error in a TCAP message, The reason is that each specific code in the message is preceded by a flag to indicate that the code referred to has been included. Then there is a length indicator to identify the length of the upcoming specific. Then the specific code itself appears. For example, in the Error component, a flag first is set to say "Yes, an error code is coming". The next field says "This is how long the Error code is". And, finally, the next field says "And this is the cause of the Error" Similar sequences are found throughout TCAP messages.

For more information on the products on this page, visit our home page at
http://www.ss8.com or come to http://www.adc-adapts.com

SS8
Networks
One Better

SMserver™

Short Message Service Center (SMSC) for GSM and IS41 wireless networks.

OTAserver™

Over-the-Air Service Provisioning (OTASP) for CDMA and TDMA wireless networks.

CALEAserver™

SS8 offers the solution that allows carriers to meet Communications Assistance for Law Enforcement Act (CALEA) requirements today.

Signaling Gateway

There are a lot of gateways. But SS8's distributed environment and high performance make for a stand-out in the field.

Distributed7™

SS8's latest generation of SS7 development platforms is designed for high availability carrier applications. The clustered multi-host architecture enables SS7 to run on multiple computers simultaneously under a single SS7 point code. Distributed7 takes SS7 reliability beyond fault tolerant platforms.

Connect7™

Host independent controller board embedded with full Signaling System No.7 (SS7) functionality.

Easy7™

The ideal link concentration solution for smaller networks. This mini-STP can be in the same room with multiple SS7 signaling points with the result that fewer SS7 links can handle your network traffic. For some, the pay-back can be astonishingly short.