

14 Networking

Introduction

While the final top level deliverable out of a truck today is still in most cases a stream, or a group of audio and video streams, increasingly, the workflow to do that is orchestrated by IT. Information Technology, or IT, is the blanket term that covers how multiple computer network aware nodes communicate, and share information with each other. The truck industry tends to refer to this as I/P Networking, most likely a cross between IT and Production. Historically this I/P networking, which we will refer to as simply networking from here out, was numerical data that represented everything from characters, pieces of a picture, addresses, to number values for their own sake. The fact that the vast majority of computers from the 1940s on dealt in number values meant they were digital. The “digital revolution” started over three generations ago.

Networking on the Road

“Modern television equipment employs networking technology to both provide control and process audio and video signals. As an example, today’s production switchers use LAN architecture to connect the switchers control panel and auxiliary panels to the electronics frame. Almost all equipment today is built this way; a control surface connected via a LAN to a processing frame. Audio consoles work this way as well. Routing switcher control panels connect to the physical matrices through a controller via a LAN.

Graphics generators and video servers are really nothing more than specialized computers with video or graphics processor cards, these computers can be networked together, just like business systems. And of course, all trucks need to support traditional IT functionality of web access, e-mail and printing. Other networks can be used to transport multiple audio channels over single network cable point to point and point to multi-point.

List of common networks at a Venue

1. Public Internet—for general web surfing by the crew and e-mail
2. The Public Internet behind a firewall, spam filter, anti-virus, etc, etc- used to get to the web for files going to servers, graphics machines, troubleshooting, etc.
3. Production Switcher- panel to main fame, aux busses
4. Cameras—CCU to RCPs & Master Setup Panels
5. Router Panels
6. Audio Console
7. Graphics
8. One of more replay networks
9. Terminal Equipment control panels for control of Frame Syncs, Color Correctors and other Modular Gear
10. Tallies & Under Monitor Displays
11. One or more Audio Networks
12. Enhanced Graphics (PVI, Sport Vision, Sports Media, Score Board Interfaces, Bug Boxes)
13. VOIP Phones

The danger on a remote site is intermingling the various networks together. The consequences of connecting the public network to the switcher LAN, creating a scenario where e-mail checking and web surfing are occurring on the same LAN can be disastrous. The switcher becomes infected by a virus, or the switch execution times are slowed by network traffic. In large events, multiple trucks can be connected together,

including their LANs, creating collisions if IP addressing is not planned out properly, such as if two devices have the same IP address. The technical portion of networking on a truck generally needs to be known by a couple of people, the EIC and E2 (engineering assistant), the practical implications of networking, such as what to plug your lap top into need to be understood by everyone on site.

In most organizations, the IT department is in charge of all things LAN, TCP/IP networking, and usually all applications in use, especially from a strategic enterprise wide view. Even many local broadcasters have adopted this approach. While broadcasters will often refer to separate business and operational networks, today in many facilities, if it is IT related, from switching to air, or working with Excel spreadsheets, it is under the watchful supervision of IT. Even though the first computer and LAN use was in engineering and operations, this IT focus reflects broadcasting’s shrinking share of revenues, as more and more cable or satellite outlets compete for end viewers.

The name of the game today for broadcasters is to do more with less. Most decided long ago that meant one IT department, and the loss of email or the station’s traffic system, the two largest IT tasks, seem to trump operational concerns, short of the station being off the air.

Remote Television Production grew up differently. While computers have been found on trucks for 25 years, they were usually stand-alone boxes. The pressure to automate many processes using computers linked together was never as pressing out in the field. Any system that had the early markings of a computer LAN was often an island unto itself, and often even those involved with them didn’t think they were “networking”.

While email started showing up in the field when trucks started seeing dialup connectivity, none had email servers on board, and email clients simply “dialed” into a server back home. About 10 years ago, equipment started to reach out, instead of being comprised only of pieces specific to its own system. It allowed outside clients, usually run on generic PCs, to setup and operate a system, whether audio/video routers and intercoms, video switchers, cameras, or processing equipment along the audio and video sequences.

So in the truck world, it is the engineering effort that still far outweighs the office side. Even with all the advanced technology, the constant dis-assembly and re-assembly, the exposure to the elements, the sometimes continuous mix of different people and personalities doesn't lend itself to rigid adherence to firewall and LAN sub-netting policies. Then, add the fact that live remote production is much more than simply routing real time video to air; it is the creation of real time video that often requires the network layer to get out of the way while it is occurring.

The truck engineer needs to have a very good understanding of the IT realm, as that's how the engineer constructs the desired client workflow. But that the same engineer is expected to understand the real time urgent need is getting the program on air—missing that objective is disastrous to anyone who plans a career in the industry. So in the truck realm, since the networking effort is a part, but not the whole of the overall truck engineer's effort, it is often referred to as IP Networking on the truck.

The Computer Layer on Top

The digital wave has totally displaced analog video except at the very ends, with the conversion of light energy to electrical potential at the camera, and presenting the scene captured onto a display. Some might still argue that some displays are quantitative or discreet, and thus digital, in how they light up the display. Analog was continuous and infinitely variable, digital is sampled into discreet chunks and restricted to a finite number of values. It is often hard to determine when a computer is doing traditional data processing tasks and when a computer is masquerading as a video processing box.

Computers exist in two camps; servers and clients. Servers do just that, they serve. Every time you visit a website your browser is dealing with a web server that ‘serves up’ web pages and software snippets, HTML code, CGI scripts and Java Applets that make up the page displayed. A server has to be able to handle lots of external requests, and have lots of resources, including memory and storage. They must be fast, with enough redundancy, and backup power supplies, to ensure that they are up and running as close to 100% of the time as possible.

Your browser is the client to that web server. It makes requests that the web server responds to. Many systems today operate like web servers. You find this synergy on trucks today, with web pages that access other resources, like video files, television router control, or truck environment control. So web servers can front other internal services in the computer, hosting the web server or services on other computers. Services like Databases that track all the media assets available, or what every router's crosspoint has selected as an input, are examples. Other services can be video servers, which are specialized servers that playback video on request, usually more than one clip at a time. Another would be a SNMP server that would provide information on the status of equipment throughout the truck, or even the whole compound.

A lot of equipment, even that which is not technically considered to be a computer, still has embedded web servers. Often, these web servers are the primary way to control and check the status of that equipment. Clients don't have to be near the server. The truck vendor's headquarters could have a client that looks at services on the truck. We'll see more on that later.

What are some common servers and client partnerships on a truck? There is the video server, a centralized server talking to multiple clients throughout the truck. There is also equipment that has a web server built in, where a browser on any computer in the truck could reach a SNMP server gathering info from SNMP aware equipment on the truck. There are also graphics servers, where a powerful machine containing a graphics engine has software running to deal with requests and commands from users throughout the truck. Editing systems are similar, with a centralized video server that has high resolution media, and serves low resolution copies, or proxies, to the editing stations. Here, edit decisions are made, such as clip in/out points, and then

the EDL (Edit Decision List) is shipped back to the central server. The EDL is then used as a log to play out the clips in order and length when requested.

IP Networking is pervasive not only across the media environment, but in a vertical sense as well. It's found at the top as part of the control system, and at the very bottom where the cables that connect one computer to another reside. This layering or stratification of various parts of IP Networking operation was codified into what is known as the OSI stack. OSI is short for Open Systems Interconnection model, a standard officially known as ISO/IEC 7498-1. This model has seven layers; we will start at the bottom and examine each layer.

Layer 1—Physical Layer

This is the bottom layer and it is generally the cable, connectors and other hardware that connects the components, nodes, and other items that comprise the computer network.

Cable

Network cable has historically been known as Cat#. Cat is short for Category.

- Cat 1 Voice Grade cable—not for data communications. Used for POTS (Plain Old Telephone Service). This uses 2 pairs of unshielded twisted pair.
- Cat 2 Uses 4 pairs of unshielded twisted pair. Good for data rates up to 4 Mbps (10 MHz). This rating is considered obsolete.


- Cat 3 Uses 4 pairs of unshielded twisted pair. Good for data rates up to 16 Mbps. Three twists per foot. Limited now to telecommunications and considered obsolete for data.
- Cat 4 Uses 4 pairs of unshielded twisted pair. Good for data rates up to 20 Mbps. Considered obsolete.
- Cat 5 Uses 4 pairs of unshielded twisted pair. Good for data rates up to 100 Mbps. This should no longer be used. Cat 5e should be used instead.
- Cat5e Uses 4 pairs of unshielded twisted pair. Good for data rates up to 100 Mbps. The e in Cat 5e stands for enhanced. This cable is capable of handling disturbances on all pairs caused by transmitting on all 4 pairs simultaneously. Cat 5e is generally the same price as Cat 5.
- Cat6 Uses 4 pairs of unshielded twisted pair. Good for data rates up to 1 Gbps (250 MHz). Length limit = 228 ft. It should at least be used for "riser" cables (between floors). Obviously not of much use within a truck but might be necessary for a compound.

It should be kept in mind that "yanking" on cable can change the cable's bandwidth or throughput because of changes to the twist and spacing of the various pairs in the cable.


Cable Termination

Often, for maximum flexibility, you don't run cables straight from one node to the next, but to what is known as punch down blocks; the solid copper wires are "punched down" into short open-ended slots.

The most common type of block is known as a 110 block. One side of this block has the wires punched down, the other side has RJ-45 connectors. These blocks are sized to handle from 25 to over 500 pairs. It can pass 1 Gbps traffic using Cat 6 cable, but Cat 6 cable is difficult to work with due to its size.



Typical punch down block and the tool required to insert wires into the block.



Most connections today involve RJ-45 connectors. There are patch panels for LANs, and they use RJ-45 connectors.

Speed/Bandwidth

The original speed used for Ethernet traffic was called 10baseT, which was half duplex. This meant only one node at either end of the line could send at once, the other had to listen. This made available bandwidth use only 30 to 40% efficient, which meant that only 3 to 4 Mbps could be passed back and forth per second. Half duplex is not common today. The 10 in 10baseT is the speed in Mbps. T means twisted pair cable was used, and it had to be less than 100 meters. Originally a single coax cable was used to connect all nodes on a segment together, and the cable was actually tapped, outer shield and inner conductor "bit into", so an AUI (Attachment Unit Interface) could be added that actually connected to the node.

Faster 10/100BaseT devices followed (10 and 100 Mbps speed) that no longer used coax but rather Cat 3, 4, or 5 cable. Devices or nodes today sense what bandwidths are available from other nodes and what the cable path will support; they can automatically sense the speed of the node it is connected to and will connect accordingly. Today many paths, especially if handling video data, are 1000 Mbps or 1 Gbps.

Fiber

Fiber is now a common physical path. Fiber is divided into single mode and multimode. Contrary to what you would think, single mode is better, and can carry coherent information longer than multimode fiber. Single mode has a glass optic path that is narrow enough that light is directed down that path without reflecting off the outside surface of the glass strand. Multimode means that light is repeatedly reflecting off the sides at various rates. The more reflections a beam of light undergoes mean a longer total path than light that sticks to the center of the glass. Some wavelengths and light rays encounter more reflections than others. This means that over long distances, packets of light can become compressed together, making it impossible to discern the data patterns that the light represents.

Multimode fiber has a maximum distance of up to 3000 feet.

Single Mode Fiber has a maximum distance of up to 25 miles (40Km).

Glass has greater distance, but copper is less expensive.

Fiber also has two common connector types:

ST (Straight Tip) Connector—Developed by AT&T, this uses a BNC attachment mechanism which is popular, but mainly for multimode fiber.


SC (Subscriber or Square) Connector—Connectors are latched, so this works for either SMF or MMF. Good for 1000 matings.

Another type:

LC (Local Connector)—Small form factor, popular for Fibre-Channel and Gigabit Ethernet adapters. This connector is replacing SC connectors because of its smaller form factor.

Many connectors are Angle-Polished Connectors (APC). These are generally distinguished by some green on the connector. They have the fiber end face polished at an angle that prevents light reflecting back from its surface of the fiber strand it is facing to reverse back down the fiber strand it just traveled. This is known as insertion loss. Any optical reflections can reduce the distance that the light can travel.

Hybrid Fiber Coaxial (HFC) is a telecommunications industry term for a network that incorporates both optical fiber and coaxial cable to create a bandwidth network.



Like all cabling, network cable can get out of hand also.

Layer 2—Data Link Layer

Now that we have our network highway in place, we need to be able to use it. That is accomplished by channel coding: a stream of highs and lows are created by the NIC and convey data that one node wants to send to another. This usually includes extra information that tells the receiving node how to recover bits that have been lost, and how to synchronize with the sender.


At this level, the channel code can be thought of as a conveyor belt from one node to the next. The ends of the conveyor belt can be thought of as the device's NIC. We will conceptualize that data packets are a series of boxes, each packed within another.

NICs


Every node or device on the network has to have a Network Interface Card or Controller (NIC). This is what allows the device to send and receive information over a network cable. So for data to and from the device to get on the network, the NIC must have instructions as to how to pack the data packets and get them out of the NIC onto the network cable. This is done by a process called 'binding'.

Just as we said that you can think of the path between nodes as a conveyor belt, you can think of binding as the process of specifying which box of data is packed first and then put into another box, and then another. In more technical terms, binding tells the NIC what data protocols are in the stack above it. Every protocol on a node must be bound to one or more NICs, and every NIC to one or more protocols. The NIC puts the box with a box inside it into the final box for the trip on the conveyor belt. Today that conveyor belt is almost always an Ethernet path.


The Ethernet's job is to make sure that only one node puts a 'box' on the conveyor belt at a time. Today there is usually only a sender and receiver, so that job is fairly easy. In the past, a conveyer belt might have passed many nodes.



The original Ethernet topology. The physical network was originally coax that was "t"-ed at each node. Every node saw all traffic on the network. Ethernet meters traffic by detecting collisions between the nodes, at which time the nodes involved with the collision will wait a random amount of time before trying again. All the nodes are said to be in the same collision domain.



To ease the interconnectivity of nodes, the hub was introduced. Here, externally, there appears to be a star topology where all nodes converge on a single box, but internally and electrically a ring topology is in effect. All traffic still sees all other traffic, so they are all still in the same collision domain.



The advent of the switch allowed the sending node and the receiving node to be connected directly to each other and no other, so there is no chance of another node stepping on their transmission. The connectivity between the two nodes is set up by the destination and source bytes in the data link packet.

This final box that ends up on the conveyor belt only has the address of a node on the conveyor belt. Today, almost every conveyor belt only connects two nodes, and one end is usually a sorting center for the neighborhood. The sorting center only knows about the local neighborhood. If the local sorting center doesn't recognize the address as being in the neighborhood, it will open the box and look at the box inside to find its final destination and repack it into a new box with the address of a node that moves it closer to its destination. The Data-Link layer knows only about addresses in the neighborhood.

As we have already touched upon, early Ethernets had only a single conveyor belt that could travel in one direction or the other. Both ends of the belt couldn't send items at once. That was known as half-duplex. Today there are two conveyor belts each going in opposite directions. Both ends can send at the same time. This is known as full duplex. It obviously allows for faster transfer speeds, and provides collision prevention, as there are no collisions when there is only a node at one end and a sorting center at the other. It's 100% efficient in both directions.

Channel Bonding

Sometimes, to increase throughput even more, two or more NICs are paired together. The NICs work in parallel to increase the throughput or amount of data that can be sent or received by a particular node.

Data Link Sub-Layers

The Data Link layer is broken up into two sub-layers.

Medium Access Control (MAC)—Defines how packets are placed on the media. The physical address is defined here.

Logical Link control (LLC)—Determines what network protocol is used—usually IP. This also provides flow control and sequencing of control bits.

The Data Link layer combines bits into bytes and bytes into frames.

How does this equate to our 'box inside a box' analogy? Every NIC has a unique physical address, known as its MAC address. When we talked about the outer most box having a send to address, we were talking about the MAC address. These are the local addresses. NICs only know how to box up the data handed to it and how to address it to someone in its neighborhood; technically that neighborhood is referred to as a local net. Today the local net is usually defined by all the nodes connected to the local switch. But as we will see later switches can be split up virtually to act like two or more switches.

Most nodes today go to a neighborhood sorting center, which is the switch. Before switches, it used to be that each node connected to what is was called a hub. When all nodes in an area are connected to a central box, the IT folks refer to this arrangement as a star topology—as when drawn out, that is what it looks like.

But looks can be deceiving; although it looks like a star on paper, internally the hub forms a ring. Instead of running a network line in a round robin from node to node, each node's cable runs to a hub which has the round-robin ring inside it. This means that all the nodes connected to the hub share the same common path.

As we will see, switches used today act like a central office for a phone system, or in television terms like a television router would, connecting one point to another. Unlike a TV router where one source can be connected to many destinations, a computer switch only connects one source to one destination at a time.

Ethernet

There used to be many schemes, or protocols for passing data at the Data-Link layer. Today one has won out; Ethernet. Developed by Xerox in 1973, its original intent was to connect copier sub-systems in their large machines. The ability to connect computers made DEC and Intel early supporters, in addition to the Xerox Research Center, PARC. The 3 came up with the DIX (Digital-Intel-Xerox) standard which used coax cable and only ran up to 10 Mbps. DIX, with some changes, eventually became IEEE 802.3. Purists maintain that Ethernet refers to the original Xerox standard and not the standard used today. Novell, the original network topology, used Ethernet with significant impact, which helped make it popular. It is the data link layer where data to be sent from one host to another is framed for transport.

Ethernet uses a process called CSMA/CD (Carrier Sense, Multiple Access/Collision Detection).

Broken down this means:

Carrier Sense—A node waits for no traffic on the bus

Multiple Access—Any node can use the bus

Collision Detection—The sending node listens to what is being sent. If another node sends something at the same time the node will know that its message is corrupted.

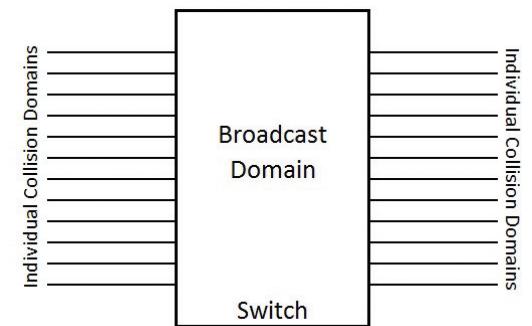
If a collision is detected, all nodes stop transmitting and each waits a random time before trying again.

The CSMA/CD method is simple, but may have collisions, thus most Ethernet networks can never reach their advertised speeds. The Data Link layer will report errors if more than 16 straight collisions occur when trying to send a frame.

With Ethernet topology the term Collision Domain (CD) is used. A CD is a set of Ethernet segments that receive all traffic generated by any node within those segments.

Before the widespread use of switches and routers, when CDs could be quite large, there was a common rule of thumb called the 5-4-3 Rule. If a node didn't detect a collision it didn't resend the packet. If the network was large enough that the last byte left

the sender before the first byte reached every node in the network, an undetected collision could occur. If two nodes communicating were at the opposite extreme edges of a network, neither node would know of a collision, and thus not resend its packet, so the data would be lost. The limiting size of a network was based on round trip signal propagation delay and inter-packet gap. The industry developed a rule of thumb to build networks of safe sizes without having to worry about delays and gaps.



Each port on a switch, and thus each node connected to a switch, is in its own collision domain.

The 5-4-3 rule says that in a collision domain, no two nodes may be separated by more than 5 segments, 4 repeaters, and 3 populated segments. Segments include from node to repeater or hub, repeater to repeater, and final repeater to destination node. Segments between repeaters without any nodes hung on them are considered unpopulated. A repeater connecting only to other repeaters doesn't count as a populated segment.

This rule today is largely ignored, because of the use of switches and routers. In fact, the terms repeater and hub were only important when the industry used passive (no active electronics) hubs. Repeaters, and their cousins, bridges, were used

to amplify the signal for additional travel, and to connect, and at the same time separate, local networks respectively. Repeaters and bridges were active (they had active electronics) devices. But the rule is still instructive as to the limitations of the Ethernet.

A NIC sends Ethernet frames (based on IEEE 802.3) consisting of 26 bytes of overhead, surround the payload data consisting of the following:

Bytes	
8	Preamble—alternating 1s and 0s to create a clock signal equal to ½ the data speed This gives the receiving NIC time to sync. The last 2 bits are 1s, which indicates stop.
6	Destination Address—This is the MAC address to which the frame is addressed.
6	Source Address—MAC address of the sender.
2	Total frame length—the length between 64 and 1500 bytes. 38-1474 Data Payload. This is the payload carried by the inner box.
4	Frame Check Sequence (CRC) for error detection.

You can see that the outer 'box' consists of timing information to determine the connection speed, destination and source addresses. It also indicates how large the contents are within the box, and includes some error checking info. At the Ethernet level it doesn't know if it's carrying email or video data.

Note: The industry often refers to this data framework not as a box but as an envelope—as the Ethernet frame does “envelop” the data it is transporting.



Each data-link packet of data is comprised of the bytes shown here. A Preamble allows the receiving device to sync with the packet. The packet provides information for the receiver on who the packet is intended for, who sent it, and how long the packet is. As shown, the payload can vary from 64 to 1500 bytes. This is what each packet of Ethernet data looks like.

A MAC address consists of six bytes (48 bits).

The first bit is a flag that indicates whether the MAC address is a broadcast or multicast address.

The second bit is a flag indicating whether this a locally or globally administered MAC address.

The next 22 bits are uniquely assigned to particular vendors.

The last 24 bits indicate a particular NIC card.

A Data-Link broadcast Layer 2 broadcast is indicated when all bits in the destination address are high, or FF:FF:FF:FF:FF:FF. This means that this packet is intended for all nodes.

What happens when a node receives an Ethernet frame?

1. CRC is checked and if correct,
2. The destination is checked—if it is a match,
3. The outer Ethernet frame is discarded and handed to the correct protocol above it. In other words, the outer box is thrown away.
4. The Ethernet-type field is checked on the inner box, to see what protocol was used at the network layer.

While the LAN on a truck may consist of a single local network, many trucks today require multiple local LANs. If the truck is part of a large venue compound, many local nets might be brought together. The simplest way to do this historically was with a Bridge. They were first used to break up collision domains (CDs). Remember that a CD usually spans all nodes connected to a hub, or in early networks, all nodes connected together on the same length of tamped coax with AUIs.

Bridge

A bridge filters traffic between local networks, also known as network segments, by looking at the MAC address. A bridge can be a stand-alone device, or built into a PC with two NICs. A bridge builds a list of which MAC addresses it detects on each port. Each port represents a network segment, and each segment is a separate CD. At first, a bridge might not know that two computers or nodes are on the same domain, and pass that frame through to other segments or domains until it sees that the second computer responds from the same segment or domain. From then on, it will know not to pass traffic on when those two nodes communicate. If a collision occurs on the far side of the bridge (in another CD) the bridge handles the collision, not the sending node in the originating CD.

Bridges do not alter the Ethernet frame in anyway. The original sender's MAC address is sent on.


If a bridge is designed to connect different protocols, such as Ethernet to Token Ring, those are referred to as Translational bridges. You won't find many of these anymore.

Bridges always forward broadcast traffic, which is traffic intended for all nodes. Network devices use broadcast addresses to discover information about a network, such as when looking for a resource on the network.

Bridges can't handle multiple routes between two different nodes. They get confused as to which collision domains the two nodes are actually in. This is known as a bridging loop. Most bridges use a technique called the Spanning Tree Algorithm to detect bridging loops and automatically disable bridge paths that form loops.

Switch

While Bridges are software based, switches are generally hardware based ASICs. Like a bridge, a switch keeps network segments broken apart to limit CDs. Bridges and switches are faster than routers. A switch can be viewed as a multiport bridge. Bridges and switches work at layer two, while routers work at layer three. Again, a switch is a port to port switch, only one port to one other port. Those connections will change as one network segment will need to talk to other network segments. This switching is all controlled by MAC addresses. The switch will connect a source node on one segment to the destination segment based on the destination MAC address. In many cases, switches are changing connections continuously. A switch differs from a bridge in that a bridge has only two ports and a switch can have many.




Typical switch installed. This one has 48 ports. Most connections go out to a single node. Some ports can be set up to connect to other switches, and to be administrative ports for switch set up.


Older networks were based on collapsed backbones, a central router or switch where all hubs, bridges, or switches are connected to it—the topology was that of a star or a rooted tree. Today's design is characterized by flatter architecture where all devices are in the same broadcast domain, thanks to switches. A flat network means that all hosts and resources are in the same broadcast domain.

As we have mentioned, you can confuse bridges and switches if there are multiple paths between network segments. To combat this situation upon wakeup, bridges and switches invoke Spanning Tree Protocol (STP) where each port on either device either allows traffic forwarding to other segments, or the port is set to block traffic between two segments. As each port settles into either a forwarding or a blocking state, this process is known as convergence. When all ports on bridges and switches have transitioned to either forwarding or blocking modes, the LAN is said to be converged. The goal is to leave only one path between network segments. During the convergence phase, no data will be forwarded. Convergence determines that all ports are either in the forwarding or blocking mode. STP is a data link protocol. From here on out we will concentrate on switches.

When convergence is complete, every port on every switch will be left in one of three states. It will be a root port, a designated port, or a blocked port. A root port is the port looking at a network segment that provides the shortest path back to the designated root switch. A designated port is a port “designated” to communicate with a network segment. If a switch is surrounded by four network segments, then three ports will be “designated ports” (DP), and one will be the “root port” (RP), connected to the network segment leading back to the root switch. A “blocked port” (BP) will not communicate with the network segment it is connected to, as another port on a different switch is handling that task. This is how multiple paths between nodes are prevented. Multiple paths between nodes can lead to “broadcast storms” where multiple switch ports are all trying to push the same data to a single node. This can quickly avalanche as responses from the destination are multiplied on the return path.




Here is a pre-converged LAN. Notice that any node can find multiple ways to any other node.



After convergence, you can see that some switch ports leading to other switches are blocked.

If a port fails on a switch, all the switches will re-enter the convergence mode and a port that was blocked from talking to a segment will now be the “designated” port for that segment. Often, this causes a ripple effect where a number of ports will change between the three states.



Re-converged LAN after a switch has failed.

To prepare a group of network segments to be tied together by switches, the IT administrator designates Bridge IDs (BIDs). Switches still use the BID term for each switch. Whatever switch they want to designate as the root, where all other switches point to this switch, is assigned the lowest BID value. This BID value is combined with the switch port's MAC. The BID default value is 32768, and can only be changed in increments of 4096. The switch assigned the lowest BID will be the root switch when convergence is completed. BPDU packets contain BID fields.

If a switch receives a data packet and it can't determine the output port to use, it broadcasts the frame out of all non-blocked ports except the port on which the frame was received. Once a switch determines which port to use for a particular destination, it puts that information into a table. The table is expanded as new destinations are encountered. This is called transparent bridging. When convergence occurs, those tables are erased and the process of discovery starts anew. Temporarily slow network traffic might be a symptom that a network failure or change has forced your network to re-converge.

If a node on a switch cannot communicate with another one connected to the same switch, the problem is most likely that either STP shut down one of the two ports, or that VLAN membership for the port is configured incorrectly. We will look at VLANS shortly.

So switches are either forwarding, flooding out all ports, or dropping received frames. No notification to the sender is done regardless of the switch's decision. As we will see later, lost data packets are handled at a higher layer. So to recap—the switch has three functions—Address learning, filtering, and loop avoidance. You can set switch ports to authenticate authorized MAC addresses by keeping a list generated by the administrator that is not eliminated during a re-convergence.

It should be stressed that while switches recognize frames and pay attention to source and destination MAC addresses, as well as to the ports that Ethernet frames are received and sent on, hubs do not. Hubs dutifully relay every frame received to every other Hub port. A network switch can usually simply replace a hub with no other changes. The most cost-effective way to start implementing switches is to keep the hubs, but add a central

switch that all the hubs connect to. The other advantage is that you can connect central servers to the switch so that they only receive traffic intended for them. Each node connected to its own dedicated switch port means that only two nodes are in the same collision domain. This means more bandwidth and thus throughput.

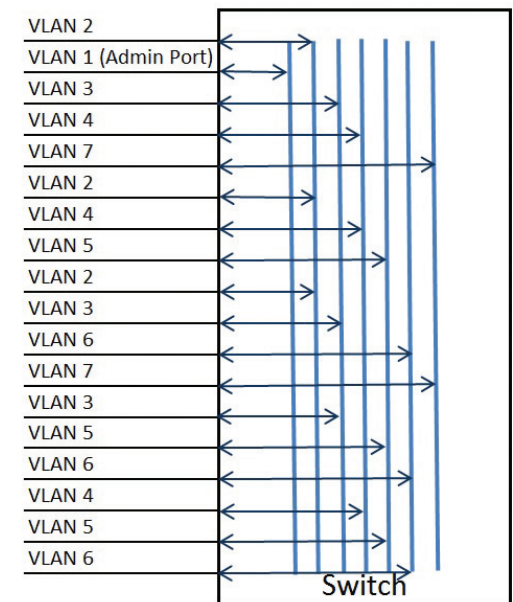
While we stated earlier that a switch connects one port to only one other port at a time, there is an exception to this. Port mirroring is used on a switch to send a copy of packets seen on one port to another port. This is also called Switch Port Analyzer (SPAN) or Port Spanning. It can cause a lot of overhead on the switch.

To add some confusion to switches, while they are generally found at the layer 2/Data-Link level, there are devices also called switches found at higher OSI layers. Some switches have ambitions of working at the next layer up, level 3/Network, and provide some rudimentary router capabilities. There are also switches that operate on layers 4–7. These provide content switching that allows load balancing. This is useful for HTTP, HTTPs, VPN traffic, especially when converging on servers and other resources.

VLANs

One powerful thing that switches allow is the creation of Virtual LANs (VLANs). VLANs allow switches to break up broadcast domains. VLANs allow you to create smaller broadcast domains within a layer 2 switched inter-network. A VLAN is treated like its own subnet or broadcast domain. In many ways, a VLAN can be considered like a workgroup.

A VLAN is a logical grouping of network users and resources connected to administratively defined ports on a switch. This process allows a single piece of hardware, the switch, to segregate different networks. This allows a truck to create control, production, and graphics networks. Neither one can see the other at the Data-Link Layer 2 level, although they can be connected at higher levels via routers, as we will see. Each VLAN is considered a broadcast domain so it must have its own subnet number, which we'll explore further when we get to the layer 3 Network-Layer.



Here a single switch has been virtually divided up into 7 switches. Nodes on VLAN 1 know nothing of VLAN 2, etc. Thus, everything associated with Replay, the production compartment, or camera control can be assigned to separate VLANs.

A switch not set up with VLANs is a single broadcast domain, but each port is a separate collision domain. The VLAN that a node is a part of depends on the switch port to which it is connected. This can be a manual process where the switch is programmed by the administrator, as to which VLAN a port is associated with.

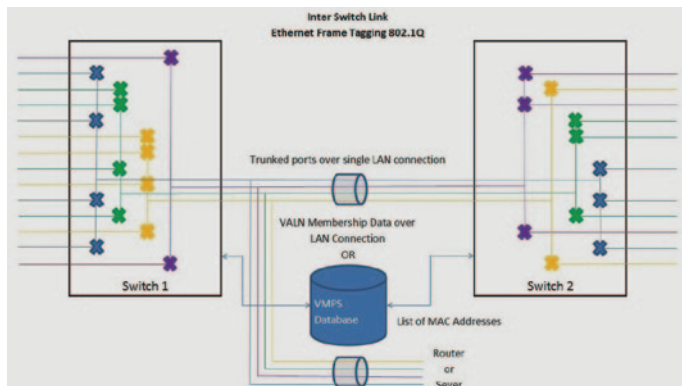
It can be done dynamically also: Dynamic VLANs look to a database to assign a node to a VLAN, based on either a MAC address, the communications protocol being used, or even on the applications run. This means that when you connect a node to a switch port, that port is automatically assigned to the required VLAN based solely on the node itself. A VLAN Management Policy Server (VMPS) can be used to do this. It automatically maps MAC addresses to VLANs.

The protocol for implementing the distribution of VLAN information is GARP (Generic Attribute Registration Protocol) VLAN Registration Protocol (GVRP), one of the longer acronyms encountered in networking, allows switches to obtain VLAN settings from a server or even other switches. GVRP allows VLAN information to be propagated from device to device.

The switch usually allows you to rename the VLAN from say VLAN2 to Graphics. Historically, VLAN 1 on switches couldn't be renamed or deleted, so many recommended that it be used for administrative purposes.

The Ethernet standard, IEEE 802.1Q to be exact, presents a way to identify which VLAN a frame is associated with, known as Ethernet Frame tagging. It inserts a field into the Ethernet frame to identify the appropriate VLAN.

VLANs can span across many switches, not just one. So today, many switches allow a form of trunking, to move VLAN traffic associated with a VLAN on one switch to another switch with ports associated with the same VLAN. A switch trunk port is a 100 or 1000 Mbps point-to-point link between two switches, or a switch and a router, or a switch to a server. A trunked port can carry up to 4094 different VLANs. Having a server connected to a trunked port means you don't need a router to have users on different VLANs able to access it.



Here switches 1 and 2 are interconnected, or trunked. The two switches act as if they are physically one switch, but as the diagram shows, VLANs can span multiple switches.

A switch port is either assigned to a single VLAN if it is an access port, or to all VLANs if it is a trunk port. Dynamic Trunking Protocol

(DTP) can operate on a port by port basis to set the switch port mode. DTP does this by negotiating with the port on the other end of the link.

Cisco implements this with what they call Inter-Switch Link, which is a way for a switch to encapsulate an Ethernet Frame with another header and CRC. The header, which is only used between switches, routers, or servers, contains VLAN ID information. This trunking is only used on Fast Ethernet and Gigabit Ethernet links.

VLANs are often confused with VPNs. VLANs are used for local internetwork to internetwork connections. VPNs are for remote (through a WAN) inter-network connections.

For security, a switch port can be set to “access mode”. Any traffic on that port will have no VLAN tagging at all. Any device attached to an “access link” is unaware of its VLAN membership. The device connected to that port assumes its part of the same broadcast domain as all the other hosts it can reach. There's no understanding of the overall physical network topology. The switch removes any VLAN information from the frame before it is forwarded out to an access-link device.

Instead of using routers, you can use VLANs with switches on the inside trusted network. Multilayer switches containing their own security features can sometimes replace internal routers to provide higher performance than VLAN architectures allow over routers.

Finally, many layer 2 switches have many layer 3 capabilities, so often you will see switches that will even have ARP (IP address to MAC address translation—we will review that in the network layer) setup. As the amount of capability that can be built into a box increases, devices that at one time would work only on one layer are becoming rarer.

Wireless

By their nature, wireless, or transmitted RF systems, can be received by any device nearby. To prevent eavesdropping, data needs to be encrypted, although you still find many who use no encryption. Wired Equivalent Privacy (WEP) Encryption was introduced around the turn of the century for wireless systems. The name implies that it is as secure as wired connections, but over the years that has been proven not to be true. WEP is wireless security at the Data-Link layer.

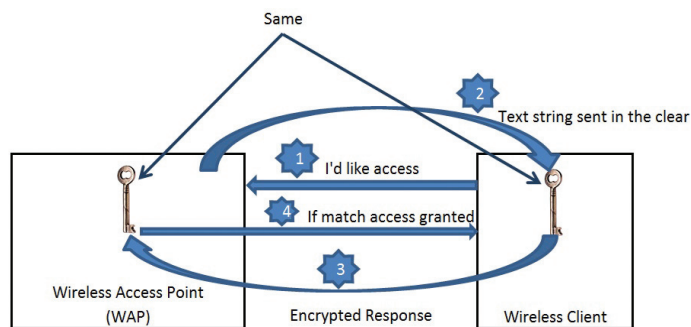
There are two methods of authentication used with WEP:

1. Open System—no credentials for authentication need to be presented to associate with a Wireless Access Point (WAP) or another host in ad-hoc mode.
2. Shared Key—A WEP key is used for authentication. A 4-way challenge-response handshake is used.
 - a. The client sends an authentication request to the access point
 - b. The access point sends back a clear-text challenge. This is the main security problem with WEP as the clear-text challenge and the encrypted result can be intercepted and reverse engineered. The challenge-response is known as a 24-bit Initialization Vector (IV).
 - c. The client has to encrypt the challenge text using the configured WEP key, and send it back in another authentication request.
 - d. The WAP decrypts the packet and compares it with the clear-text it sent. If correct, the device is allowed on the wireless network.

The WEP weakness is that its initialization vector (IV), which is a pseudo-random string, is sent in the clear to the client, and it is a 24-bit key at the end of packet. The first 40-bits are the message for the client to encrypt and send back. Using RC4 encryption, the client uses the IV to randomize the text message which is actually 5-eight byte words, both by shifting the message's order and its bit-value. This reorders the text it received with the IV string and then sends it back to the WAP. The client must have the same

key the WAP has; when the WAP receives the reply back, it reviews the encrypted data and unscrambles the message; if it gets back the original data it allows the client access to send or receive one packet of payload. Those packets are usually 1500 bytes long. Another challenge-request session with a new IV is sent each time a client requests to send a new packet of data.

The 40-bits text message yields about a trillion possibilities for scrambling. Today computer processing horsepower could crack that cipher in a few seconds. A busy WAP which is constantly sending 1500 byte packets has a 50% chance of sending the same IV value after about 5000 packets. An attacker recording all traffic can use packets with the same IVs (called an IV collision) to statistically obtain an IV's cipher with two different packets using the same IV. To beef up WEP security, 128 and 256 bit challenges were developed.



The Diagram shows the four step process for gaining access to a WAP. The secret here is that both the client and the WAP have the same encryption key. When a client requests access to the WAP, the WAP issues a challenge, which is simply a text string. The client uses its key and encrypts the string and sends it back to the WAP. The WAP decrypts it with its key and if the WAP finds that it matches the original string, then access is granted.

WEP has been replaced by two more secure protocols—Wi-Fi Protected Access (WPA and WPA II) (802.11i). From 2003 onward, most WAPs used WPA. These use Temporal Key Integrity Protocol (TKIP) which replaces WEP's 40-bit encryption key that must be manually entered on WAPs and devices, and does not change. TKIP is a 128-bit key that dynamically generates a new key for each packet and thus prevents collisions.

Network Access Control (NAC)—Mainly used for wireless. This is wireless's second level of security. This is a method of securing hosts before they are allowed to access the network. One of the

most common forms of NAC is 802.1x, which supports multiple authentication schemes. NAC requires a host to join a WAP before it can get to the rest of the network.

Security Set Identifier (SSID) is the name of the wireless network and is of a Basic Service Set (BSS). SSID is a 32 character ID that represents a particular wireless network and defines the basic service set. All devices involved in a particular wireless network must be configured with the same SSID. The BSS is essentially the set of all nodes that can communicate with each other through the same WAP. Each WAP has an ID which is the WAP's MAC address. If all the WAPs are set to the same SSID, then devices can roam around freely in the same network. Doing this creates an Extended Service Set (ESS).

Temporal Key Integrity Protocol (TKIP) is an IEEE 802.11 and Wi-Fi Alliance standard. It is used in Wi-Fi Protected Access (WPA), which was introduced in 2002. TKIP was meant to replace data-link WEP security without requiring the actual replacement of existing hardware. It uses a more robust IV, and packet sequence counters to detect packets received out of order, to detect replay spoofing attacks, along with beefed up CRC checks aimed to prevent spoofing also.

TKIP in essence, wraps around the pre-existing WEP encryption key, which was way too short, and adds a much stronger 128-bit encryption. The RC4 cipher algorithm used with WEP is also used by TKIP. TKIP actually changes each packet's key. These packet keys are comprised of three things—a base key, the transmitting device's MAC address, and the packet's serial number. Each serial number is 48-bits which are augmented with a sequence number with each packet. This all acts as an initialization vector (IV). Now each and every packet will be uniquely identified, with no IV collisions over time as with WEP, and with no way of doing IV collision attacks by recycling a known key over and over to gain access to the data content. TKIP needs only one master key, as that key is used to derive other keys, whereas WEP required a number of keys for large networks.

A newer and preferred security protocol, Advanced Encryption Standard (AES) is defined in 802.11i or AES-CCMP (AES-Counter Mode CBC-MAC). CCMP is the security encryption method used by AES as opposed to RC4 in WEP. Protocols were added

and the standard became known as WPA2 (WPA2 uses AES instead of TKIP). While TKIP works with existing hardware, AES-CCMP needs new hardware, and requires more processing horsepower. AES uses 128-bit, 192-bit, or 256-bit block cipher. Basically, data is arranged into arrays or blocks, then xor (exclusively or'ed), then replaced with a value from a lookup table, then rows in that array are shifted up, or down, and finally whole columns are swapped while being multiplied by another matrix.

For 128-bit keys, 10 rounds of this shuffling occur, with 12 rounds for 192-bits, and 14 for 256-bits. The shared key provides the information required to derive and reverse the shuffling process.

The only known cracking of the AES algorithm has been via "side-channel" attacks where clues as to what processing is going on are sniffed externally. This is done by sensing things such as EMF and power patterns emanating from the system.

Another method for securing wireless links is the modulation technique used to transmit the wireless carrier. Frequency-Hopping Spread Spectrum (FHSS) is the common technique used, or the less used Direct-Sequence Spread Spectrum (DSSS) system. FHSS was partially invented by an actress from the golden age of film, Hedy Lamarr, and she and another were granted the original patent. FHSS modulates the data signal with a carrier signal that changes frequently in a random but, over time, predictable sequence of frequencies. This is also known as frequency hopping. These changes also occur over a wide frequency band, with a spreading, or hopping, code establishing the transmission frequencies used. The bandwidth is thus wider than required, but the odds are increased that two nearby WAPs can operate on the same channel (a band of frequencies) as the pattern of “sub-channels” used statistically won’t match. Manufacturers use 75 or more frequencies per transmission channel. The maximum dwell time on a particular frequency has been established by the FCC at 400ms.



This Graphic shows the spread spectrum concept. Here, two wireless signals, one represented by black, the other by red, occupy the same area of the frequency spectrum. Each column is a separate carrier, and each carrier is modulated by some modulation technique. See the Tether & Base Station section in the Camera Chapter for more information about modulation techniques. The arrows on top of the carriers represent QPSK modulation where the carrier can convey one of four states. It takes two bits in a digital word to represent four states. Here each carrier is conveying two bits. Add up a bunch of carriers associated with the same signal and you can see data throughput adds up quickly. Statistically, the odds of the two different signals interfering are small, but if they do, the higher protocol layers, which we will see shortly, will be able to fix the interference.

It should be understood that there are two parts to the spread-spectrum transmission. The first is which frequency or sub-channel to use within the channel, and secondly, there is the modulation technique used. The most common modulation types known are AM or FM. Neither of these is used; the “constellation” type techniques are used, such as QPSK, or some flavor of QAM. These allow multiple states, representing more than one bit to be conveyed per baud. Baud means one cycle of the carrier.

Direct-Sequence Spread Spectrum (DSSS) uses spread-spectrum also, but most of the time all carriers within the channel are quiet as a sub-channel frequency only “pops-up” for a short period, much shorter than the bit-rate of the data would suggest. When the spectrum is viewed across the whole channel it appears as noise.

IEEE 802 Subcommittees

- x.1 LAN/MAN Overview and architecture
- x.1 LAN/MAN bridging and management (higher layer LAN protocols)
 - x.1s Multiple spanning tree
 - x.1w Rapid reconfiguration of spanning tree
 - x.1x Port-based Network Access Control
- x.2 Logical Link Control (LLC)
- x.3 CSMA/CD access method (Ethernet)
- x.3ae 10 bit Ethernet
- x.4 Token passing bus access method and physical layer specifications
- x.5 Token ring access method and physical layer specifications
- x.6 Distributed queue dual bus (DQDB) access method and physical layer specifications (Metropolitan Area Networks)
- x.7 Broadband LAN
- x.8 Fiber optic
- x.9 Isochronous LANs (standard withdrawn)
- x.10 Interoperable LAN/MAN security
- x.11 Wireless LAN medium access control (MAC) and physical layer specifications
- x.12 Demand-priority access method, physical layer, and repeater specifications
- x.13 Not used

- x.14 Cable modems (proposed standard withdrawn)
- x.15 Wireless personal area network (WPAN)
- x.16 Wireless metropolitan area network (Wireless MAN)
- x.17 Resilient packet ring (RPR) access

Power over the Ethernet (PoE)

The Ethernet LAN connection can provide power to a device so that it doesn’t need a local power source. Devices such as IP Phones, Wireless Access Points, IP Gateways, PDAs, etc. often look to receive power over the network cable attached to them. Switches usually have an internal setting that limits the amount of power it will deliver to devices connected to it. It must be understood that power for a device is not routed over multiple hops; it is the last hop device that supplies that power.

Switches have an absolute power rating when it comes to supplying power to devices. Often that maximum power available to power devices connected to a switch’s port is under 500W, with individual ports allowed to deliver only a fraction of that, usually under 20W. You can usually set a limit under that amount if you are worried about heat, power supply load and their effects on reliability. Any devices added to a port that exceed these settings simply will not be powered.

User Authentication

Managed switches often provide a method of authentication for traffic entering a network. A switch port facing a remote ingress point, say from the Internet or any outside network traffic, can be set up to require authentication. Thus, some ports will require authentication to get on your local LAN, but once on your LAN, the inward looking ports will not require authentication.

The simplest way to have a port authenticate is to store password data in a database in the switch itself.

Another way, which you won't find on a truck LAN often, is to use separate authentication servers known as RADIUS and TACACS+.

Radius stands for Remote Authentication Dial In User Service. While dial-in service isn't found much today, it also applies to DSL service. A switch port connected to a DSL line will have Point-to-Point Protocol, carrying authentication traffic with username and password information, and often additional information about the client. Here is where the switch acts as a layer 3/network layer device. The switch sends the authentication information in a UDP (we will look at those in the Layer 4/Transport section) to a RADIUS server which checks to see if the authentication information is correct.

The authentication information is often sent in the clear, but this is considered a bad security choice. Often there is a three step authentication process, where access is requested, the authentication server sends a string of text in the clear, called a challenge, and the requesting device then responds with a "hashed" string combining both challenge string and the password. Some servers are set up to challenge not only at first contact, but also periodically.

TACACS+ (Terminal Access Controller Access-Control System Plus (plus means an upgrade of TACACS)) was developed by Cisco and is considered a more robust, but also more complicated security system. TACACS+ has separate servers handling separate parts of the authentication process. Unlike RADIUS, TACACS can also manage user privileges. TACACS+ requires that all authentication be encrypted.

Also on a port that requires authentication, the switch will usually let you specify how to handle different traffic, such as telnet, Secure Telnet (SSH), and Secure HTTP, and which authentication method to use.

Truck Switch Setup

Setting up and organizing the IP Networking infrastructure in the truck is a combination of configuring the switches in the truck

along with a small router, which we will look at shortly. You will often hear about a managed switch. Simple switches can simply be hooked up and they will route traffic based on MAC addresses, and not much more. A "Managed" switch will allow VLANs to be setup, security to be implemented, and much more. The trade off is that there can be extensive setup and management duties for these switches.

Up until now, we have looked at switches as separate entities. They don't have to be; often switches can be configured in what is known as a stack configuration. They are hooked together and configured as if they are a single large switch. One of the switch frames in the stack is made the master, and another is specified as the backup. It is typical that up to six or eight individual switches can be configured into one "super" switch. If more switch frames are added than the stack can support, the stack may partially work or not work at all.

When stacking switches, two of the ports on each frame, usually fixed and not selectable, are used to cascade from one switch to the next. While the looping from frame to frame will work as a chain, it is usually desirable to run a cable from the bottom frame back to the top stack frame, and put the frames in a "loop" configuration. This allows the other frames to continue to work as a single unit if a frame fails—because traffic can be sent in both directions around the loop in order to avoid the failed frame. The loop topology allows a failed frame to be pulled out and replaced without affecting the rest of the stack. It is also recommended that backup power supplies are used at least for the master and backup frames.

All configurations for the stacked switches are stored as configuration files on the master and backup frames. This means that a new frame inserted into the stack to replace a failed one, will not have to be configured manually but will automatically be setup via the master. Also, if the replacement frame is larger or smaller, port wise, the replaced frame will be setup based on the number of ports that the failed unit had. The stacked switches will usually indicate on their front panel if they are the master or backup frame, and what their assigned stack number is.

Not only can stacks be split, either by failure or by forcibly disconnecting them, you can merge stacks together by extending

the loop to include another stack of switch frames. If this is done, the masters of each stack will negotiate, to determine which master stays the master of all. This is usually done by which switch has been running the longest, or by which has the lowest stack ID number, or MAC address. You are also able to select which frame is the master.

It should be noted that it is usually preferable to let the master assign IDs when merging stacks; in some cases, the losing master in a merged stack will shutdown if it has a manually assigned ID. When two stacks are merged, no two frames can have the same ID. The combined configuration files are now only found in the winning master and backup frames.

If a switch is inserted which is set to run as a standalone unit, it will continue to do so and the master/backup will treat that frame as a failed unit and route traffic around it.

Managed switches also have setups to mitigate the effects of Broadcast Storms by blocking the offending traffic after it occurs for a set time, and restricting that traffic for a period of time. You can restrict traffic on individual ports to a list of MAC addresses. This is often called an Access Control List (ACL), and can be created manually, or the switch can build the list and then lock it at some point.

A managed switch also lets you set the throughput capabilities of each port, so you can restrict the amount of traffic on a port by port basis, by instructing a port to advertise its throughput capability. The crossover settings on each port can also be set. Connector wiring on switch ports is generally "crossed" where receive pins are transmit pins, etc. You can set the port to treat its port like the switch it is, or like a port on an end device. This is not often necessary anymore, as most NIC cards can sense if the

transmit and receive wires are crossed and act accordingly.


Ports can also be set to work together to increase bandwidth between devices. This is called Link Aggregation Group (LAG) management. There is a protocol defined that orchestrates this, Link Aggregation Control Protocol (LACP). LAGs can be used to ensure path redundancy. Ports assigned to a LAG can't be part of a VLAN, and there can't be auto-negotiation enabled. They must be a full duplex set, and generally the rest of the port's parameters must match each other.

Authentication setup is also a necessity on managed switches. You can simply elect to have none, or use RADIUS or TACAC+ security that was already explained. This needs to be done for every port. If RADIUS or TACAC+ is used, the authentication server to be used must be identified. A switch which has had a port programmed to challenge a user is in turn challenged for authentication by the authentication server. So the switch's credentials must also be entered.

Layer 3—Network Layer

This next OSI layer going up the ladder is referred to as the Network Layer. Layer 3 devices need to locate a specific network along with a real or virtual device or node on that network. Layer 2 or Data-Link addressing only needed to find specific nodes or NICs on the local network.

The main Network (layer 3) protocols are IP (Internet Protocols), ICMP (Internet Control Message Protocol), RARP (Reverse ARP), and Proxy ARP. The ICMP and ARP protocols support IP.



Here you can see that IP header and trailer, the IP "box" or "envelope" is wrapped around the layer 2 Ethernet data.

The Layer-2 envelope or box is wrapped in the IP box next. The first part of that IP data box is the IP Header (20 bytes) structured as follows:

Bytes	
4	Version
4	Header length
8	Priority and type of service
16	Total length
16	ID
3	Flags
13	Fragment Offset
8	Time to Live
8	Protocol
16	Checksum
32	Source IP address
32	Destination IP address

The payload that follows this header is the envelope or box from the next layer up, the Transport Layer

IP Addressing

The IP address is a 32-bit number that represents a particular Node. This address is either comprised by network and host, or by network, subnet and host. This is a two or three layer structure on hierarchical address instead of flat scheme. In the two layer scheme, the highest order bits represent a network, while the lower order bits represent a particular node within a network. In a three layer approach, the previous statement is still true but the number of bits used for network and particular node addressing is specified.

To make matters more confusing, the number of bits allotted for each can be via a few simple rules or, as we will see, to stretch the use of a limited number of addresses, the percentage that represents networks and particular nodes on that network can be varied to almost any combination. The important point to keep in mind is the simple top level premise; some bits describe the network and some describe the node on a network.

There are theoretically 4.3 billion unique possible addresses available with the 32-bit addresses scheme although only about 250 million can be actually assigned.

The 32-bit IP address is grouped into four 8-bit octets. Each octet can represent a number between 0-255. Actually this grouping wasn't for the equipment's sake but rather for human readability. 124.10.1.18 is easier to read than 01111100.00001010.00000001.00010010. It could have split up into hex values, where the previous value would read 7C.0A.01.12, but using all decimals was deemed easier to read. It's interesting that the same wasn't decided in dealing with MAC addresses. But as we will see shortly, it might have been easier in some cases.

IP address assignment started with the first group or octet referring to the network. While used with the ARPANET in the 1960s, 256 networks might have been enough. It was soon realized, however that that was minuscule as to what would soon be needed. So the standards' bodies soon developed what they called a Classful address. What this meant was that IP addresses would be divided into class, A, B, and C.

Class A—This would keep the original addressing structure; First octet contained the network, all other bits describe the node. To a machine, a class A address was easy to identify; the first bit of the address was always 0. That meant that only the remaining 7 bits were left to define a class A network. 7-bits yield 128 possibilities. Since digital systems tend to start at zero instead of one, class A network addresses have IP addresses where the first octet has a range between 0 -127. That leaves 16,777,216 possible addresses on each of those networks. Only the largest entities have class A addresses. You'll find a lot of the original big players such as AT&T, HP, GE, IBM, Xerox, MIT, Apple, Level 3, a number of U.S. government

entities, and the U.K. hold class A addresses. Some that have disappeared, like DEC, have been inherited by others, HP in this case.

Class B—A machine would look at the first bit and know this address can't be a class A because the first bit is a one. It then looks to see if the second bit is a 0 or 1. If the second bit is zero it knows it is dealing with a Class B address. This gives a first octet with a range from 128-191. Class B addresses use the first two octets to identify the network. But again, it doesn't have all 16 bits to use for unique networks, as the first two were used to identify the address as class B. So a class B address can identify 16,384 addresses. That leaves 65,536 addresses for nodes on each network.

Class B addresses point to many universities, other large players that weren't there at the beginning like Cisco, Tektronix, Sun, individual armed services, NASA, most defense contractors, a myriad of government agencies around the world, large IP service providers, and most hi-tech companies of any size.

Class C—As with class B addresses, the machine would see that the first bit is one. In this case, the second bit is one also, so it confirms that the third bit is zero. If so, then this is a class C address, which gives the first octet a range of 192-223. A class C address uses the first three octets, minus the first three identifying bits, so the remaining 21-bits from the first three octets yield 2,097,152 network addresses, allowing 256 nodes on each of those networks.

Class	1st Bit	2nd Bit	3rd Bit
A	0	X	X
B	1	0	X
C	1	1	0

First three bits in IP address determine the IP address class.

We mentioned decimal versus hex representation before; if used on IP addresses, Class A addresses would have the first hex value less than 8, Class B a value of 8, and class C a value of B.

Because of the demand for IP addresses, there is a series of addresses in each of the three address classes that are reserved for private use, which means is that there may be many nodes around the world sharing the three blocks of addresses below.

Class A	10.0.0.0–10.255.255.255	16,777,216 addresses
Class B	172.16.0.0–172.31.255.255	1,048,576 addresses
Class C	192.168.0.0–192.168.255.255	65,536 addresses

Many home routers use the class C reserved addresses. Why don't these addresses, used repeatedly, conflict with others using the same addresses? There is another technology that comes into play at the network layer to keep the usage by others separate, known as Network Address Translation (NAT) which we will look at a little later.

To make the most use of the available addresses, modern network layer equipment allows bit-by-bit deployment of which bits are used for network addresses and which for individual node addresses. This is implemented via the network mask. Most often when you type ipconfig from a DOS command line you will see the mask with a value of 255.255.255.0.

Let's look at a few examples: An IP administrator would write the 255.255.255.0 mask as

xxx.xxx.xxx.xxx/24

A specific example:

124.10.1.18/24

This means that although this is a class A address (1st octet < 128) treat it, network wise, as if it were Class C; using the first 28 bits for the network or subnet.

124.10.1.18 in binary form =

01111100.00001010.00000001.00010010

A 255.255.255.0 mask in binary form =

11111111.11111111.11111111.00000000

A machine would take the mask and would do a Boolean AND with the address getting:

01111100.00001010.00000001.00000000

Notice that the final 8 bits are now all 0

Converting this number back to decimal yields:

124.10.1.0

This is the network address part, known as the subnet.

Thus the 124.10.1.0 subnet can have 256 addresses on it.

Now the same example with a more complicated subnet: An IP administrator would write the 255.255.255.240 mask as

xxx.xxx.xxx.xxx/28

To get wonky; using 28 bits out of the 32 bits for masking is known as a 240 mask. Again using the same IP: 124.10.1.18/28. Again, this is a class A address but we'll treat it network wise, as if it were Class C; using the first 28 bits are the network or subnet.

Thus not only the first, second, and third octets, but ½ of the fourth are used to indicate the network or subnet address.

Again 124.10.1.18 in binary form =

01111100.00001010.00000001.00010010

A 255.255.255.240 mask in binary form =

11111111.11111111.11111111.11110000

Again a machine would take the mask and OR the address getting:

01111100.00001010.00000001.00010000

Notice that the second to last bit is now 0

Converting this number back to decimal yields:

124.10.1.16

This is the network address part, known as the subnet.

What this means is that subnet 124.10.1.16 can have 16 separate nodes on it, with 124.10.1.18 being one of those addresses.

A 240 mask allows 16 subnets in the network range of 124.10.1.0 at 16, 32, 48 etc., and not to confuse, each of those 16 subnets can have 16 nodes on them.

Let's introduce another subtlety for subnets now. In the above example the 124.10.1.16 subnet extends up to 124.10.1.31, which as we pointed out, contains 16 possible nodes. But in reality, there are only 15 nodes possible. That is because the upper most address on a subnet is reserved for broadcasts, or messages intended for every node on the subnet. For this subnet that address is 124.10.1.31.

Let's look at one final example: 192.20.2.21/30:

With 6 of the last 8 bits of the octet devoted to subnet addressing, this equates to a 252 mask and would show up as a 255.255.255.252 mask entry in ipconfig.

```
11000000.00010100.00000010.00010101
    (binary address)
11111111.11111111.11111111.11111100
    (mask)
11000000.00010100.00000010.00010100
    (subnet binary)
    192.20.2.20
    (subnet address)
```

This is a class C with the 1st, 2nd, 3rd, and most of the 4th octets devoted to specifying the subnet. So 192.20.2.21 is on the 192.20.2.20 subnet. Thus, there are subnets every 4 addresses.

This is an extreme example, and as we just mentioned, the top address in the subnet is reserved for broadcasting. Often, the lower address in the subnet is reserved for the gateway, leaving only two addresses for nodes on the subnet.

Rule of thumb: when setting up a small network, use a class C address. It is easiest for most people to understand and configure. This gives you one network with up to 256 hosts.

Classless Routing

We will look at routers shortly, but the various address schemes leads to the concept of classless routing. If addressing does not conform to classic A, B, and C IP addresses, and subnet schemes using the corresponding subnet mask, then they are considered Classless. With Classless routing, different devices can have different masks as we have seen in the several examples above.

DHCP/BootP

So each node, or host as IT folks call them, has an IP address. How do they get those addresses? You can manually assign them to each host, but then you have to keep track of which host has which address, and you must make sure that two or more hosts don't end up with the same IP address, which quickly gets tedious. Plus PCs, mainly laptops, come and go. Each new production team that shows up on a truck expects to use your network for their machines. You need an automatic way to hand out and re-use IP addresses. Not only IP addresses, but each host also needs to know how to reach out to the outside world, from whatever subnet it is on; its "gateway" so to speak. As we've just seen, the device needs to know what the subnet addressing scheme is.

Originally, when people were connected to large systems, such as Unix based systems, with workstations that had no appreciable storage of their own, they used a system called BootP, which stands for Bootstrap Protocol. The local workstation still had to have a local disk drive to load the BootP software, but over time NIC cards came with embedded BootP firmware. BootP told the workstation where to go to find its bootable image (OS and user files) and what its IP address was. BootP required a server that kept the list of available IP addresses and who they were assigned to.

DHCP (Dynamic Host Configuration Protocol) is almost exclusively used today. DHCP differs from Bootstrap (BootP) in that BootP assigns an IP address to a host, but the host's hardware address must be entered manually in a BootP table. DHCP is dynamic BootP. What BootP could do that DHCP can't, is to send an OS that a host can boot from. DHCP can't do that; it runs on top of BootP protocol on port 67 (ports come in the Layer 4/Transport section).

The DHCP server provides an IP address, the default gateway (router), two DNS (Domain Name Service) server addresses, and a network mask, along with the lease time, or how long the address is good for. When DHCP hands out an IP address, it usually provides a lease for its use for 8 days by default. It is up to the client to renew the lease. We'll look at DNS in a bit.

DHCP enables individual client machines on an IP network to configure all their IP settings (IP address, default gateway, the subnet mask, and the DNS server) automatically. For a simple home or truck router its internet port would be a DHCP client looking to the ISP (Internet Service Provider) for its IP address. For the LAN ports out the other side (in the truck itself) it would be a DHCP server, and the truck nodes would be clients looking to the local router for their IP addresses. DHCP can share a set of IP addresses with a greater number of temporary users. A DHCP client will always ask for the last IP address it had from the DHCP server.

Windows clients that are configured for DHCP but unable to access a DHCP server, will default to a special IP address starting with 169.254.x.x. Do IPCONFIG/RENEW at the DOS prompt to see if "DHCP Server Unavailable occurs". Use the "IPCONFIG/RELEASE" and then "IPCONFIG/RENEW" commands when changing to a new DHCP server.

A DHCP client requests an IP address via a broadcast request. The DHCP Discover Message is as follows:

Client sends out a broadcast at both layers 2 and 3 to discover the DHCP server.

Layer 2 message is all Fs—FF.FF.FF.FF.FF

Layer 3 messages is all 1s—255.255.255.255 (which means all networks and hosts)

DHCP communication is connectionless—it uses UDP (covered under the Transport Layer)

One or more DHCP servers will offer IP addresses. The requesting client will accept one and any other offers will be withdrawn.

Routers (covered later under Network Layer) must be set up to forward a DHCP broadcast from a client out to other router ports as a unicast to one of more DHCP servers for redundancy.

ICMP

ICMP (Internet Control Message Protocol) are message packets that are generated by devices at the network layer. They generally carry error messages. If a router, for example, can't forward an IP packet, an ICMP message is sent back to the sender. A router will also send out a 'buffer full' message via ICMP. The innocuous Ping command uses ICMP echo requests. A header indicating an ICMP packet will be found immediately after the IP header.

IPSec

IP Security (IPSec) was designed for providing authentication and encryption over the internet. It works at the layer we're looking at now, the Network layer (layer 3) and secures all applications that operate in the layers above it. Because of the large support it enjoys, it is the standard for VPNs on the internet today. IPSec uses an Authentication Header (AH) and Encapsulating Security Payload (ESP). The AH has no encryption and is used for authentication only. The AH header replaces the IP packet header. ESP provides both authentication and encryption abilities.

IPSec has two modes; transport mode and a tunneling mode. The transport mode creates a secure IP connection between two hosts. The data is protected by authentication and/or encryption, but in this mode, a tunnel is not created.

The tunnel mode encapsulates the complete packet within IPSec. ESP is used to give both authentication and encryption, so it is more commonly used to create secure tunnels. In tunnel mode a hacker cannot even see what transport protocol is being used.

IPSec provides end to end security and allows a device to roam from one network to another without dropping the connection. IPSec is an encryption protocol that works with both IPv4 and IPv6.

ARP

Address Resolution Protocol (ARP) is used to resolve network layer IP addresses into Data-Link MAC addresses. ARP finds the hardware address of a host from a known IP address. DNS uses ARP. The host, or local computer, keeps a local ARP list also, but you shouldn't see addresses in the ARP table for a given interface that aren't members of the same IP subnet as the interface. You can add entries to the ARP table using the ARP -s command at a DOS prompt. This entry stays until the device is rebooted. ARP -d will delete entries.

A DNS server provides name resolution using ARP protocol and provides a destination IP address and MAC address. A host sends out an address resolution request in what is known as an ARP packet. A DNS server will receive the request and provide a response with the corresponding Data-Link address. ARP requests and responses stay within the restraints of local networks, as a router will not pass them.

In most trucks, the local router talking directly to an ISP modem or a venue's router is not only your DHCP provider but also your DNS server.

DNS

The Domain Name Server (DNS) system replaced the host file that used to be kept and updated on every node on the networks in the early days. The DNS server provides name resolution using ARP protocol and provides the destination IP address and MAC address. While it is a level 3/Network protocol, it uses port 53 on the transport layer. It is the network phonebook.

Top level DNS servers handle .Com, .Org, .Gov, etc. domains. Second level servers handle 2nd level names such as the actual company name. Sub-domains, such as newyorkoffice.company.com are handled by an ISP or the company's own DNS server.

DNS servers keep a list of Canonical Names (Cnames). ftp.somesite.com is the Cname for www.somesite.com.

The DNS server has 'A', PTP, Cname, and MX records. 'A' records translate host names into IP addresses. PTP (Pointer) records translate IP addresses into host names. MX (mail exchange) records keep email addresses.

Windows Servers and UNIX/Linux come with built-in DNS server software. Many DNS local servers are cache-only, which means they will resolve names on the internet for the network but are not responsible for telling other DNS servers the names of any of its clients. Authoritative DNS servers actually hold the IP addresses and names of systems for a particular domain or domains of the local network in special storage areas called Forward Lookup Zones. The DNS server also keeps Reverse Lookup Zones, where URLs are determined from IP addresses.

Network Address Translation

As you will recall, there is far more demand for IP addresses than separate, individual addresses available. So there needs to be a way to separate internal use of IP addresses from the rest of the universe, a way to translate private IP addresses to public addresses. So Network Address Translation (NAT) was developed. NAT setup and maintenance is a concept that is usually out of the truck engineer's scope, but you must be aware of it.


Besides slowing the depletion of IP addresses, NAT is useful for network migrations and mergers when different networks have duplicate IP addresses, for server load sharing, and for creating virtual servers. This situation is quite possible when a truck is part of a large venue. ISP providers change from venue to venue, and you don't want to have to change the internal address scheme each time. NAT is absolutely required when connecting to the internet without globally unique IP addresses, which no truck will ever have.

NAT is usually run on what is called a boarder router. Most trucks don't have routers, so again a router would most likely be encountered in context of a large venue, or a device the venue itself would maintain for outside vendors as a gateway to the internet.

NAT has a couple of disadvantages. This address translation introduces path delays, there is loss of end to end IP traceability, and some applications won't run over NAT.

NAT allows what is called Address overloading. This is a type of Dynamic NAT that maps multiple unregistered IP addresses on the inside to a single registered IP address to the outside internet by using different port numbers. This is known as Port Address Translation (PAT). As an example, if four nodes on a truck are trying to reach the same web server back at Fox, all four nodes would get the same IP address on the internet side, but each would get assigned a different source port number. However, they would all have the same destination port number, 80 in the case of an http web server. The web server and your NAT server would know where to direct return web traffic by the source port numbers, but keep in mind that on the truck side of the NAT, each node would still have separate IP addresses.

As with DHCP and DNS in most trucks, the local truck router handles NAT duties.



Here you can see that five truck network nodes have Class C addresses (192.0.1.x) and all are accessing a remote web server, using HTML services, so they all have port 80 destinations. The router converts all five nodes to the same Class B address (168.0.10.15). This address was most likely assigned to the router by the internet provider. All five reach the web server, which responds back to each of the five by using the five different port numbers. When the response gets back to the truck router, it converts the outward facing IP address back to the local IP addresses based on the return port number. It should be pointed out that the local truck nodes don't know about the 168.0.10.15 IP address, or the 845x port numbers—all they know is that each reached out to a web server, and each got back its own response from the server.

Routers

Routers are usually something that will be out of the truck engineer's scope of work. But the engineer absolutely must be aware of their general use, and how they affect the larger network topology that often will surround a single truck. In addition, there are IT folks who have extensive training in dealing with these devices, and the engineer at least needs to know some of the acronyms and terminology used. Just as there is a large industry, training and certifying IT professionals, a very large subset of that is mastering the subtleties of router control and management. When an IT professional talks about network policy, the bulk of that is enacted at the router/network layer, while some part of it might occur at the switch/data-link layer and layers above the network layer.

Routers are used to provide physical segmentation between subnets. They also provide connectivity to VLANs on different subnet LANs. While routers won't pass ARP requests, they need to initiate them. Each port on a router is usually connected to separate subnet. For a host to reach separate subnets, they need to pass through a router. Routers can also be set up to discriminate against traffic based on IP addresses, port addresses, which port the traffic is arriving on or leaving to, and other parameters.

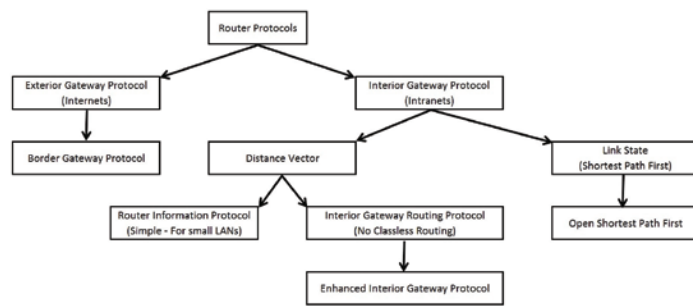
While we haven't talked specifically about ports yet, as that is a layer 4 or Transport concept, routers, considered layer 3/Network devices, also work on the next layer up.

Routers will pass DHCP requests if set up to do so. If a host transmits 255.255.255.255 (layer 3) and FF:FF:FF:FF:FF:FF (layer 2) and if the DHCP server is not on the same LAN segment a router will see the address as a broadcast. The router will look at the port number set on the transport layer (for DHCP port 63) and know that this is a DHCP/BootP request, and will unicast the request out the proper router port to the DHCP server on the proper LAN segment.

Many types of hardware can be used as DHCP servers, which include routers. That means there wouldn't need to be a separate DHCP server. Remember, if all you have is a wireless WAP/router that is most likely acting as your DHCP, while also performing DNS and NAT duties.

Another important aspect about Routers is that they can provide Quality of Service (QoS) levels for specific types of traffic.

There are a number of philosophies as to how routers are organized and work. We will briefly look at the main ones.



Family of router protocols

EGP

There are two basic overriding router philosophies. The first is Exterior Gateway Protocol (EGP), which works outside or between one or more Autonomous Systems (AS). An AS could be the NOCs LAN/WAN or all the trucks and equipment that make up the Super Bowl. There is only one EGP approach, which is Border Gateway Protocol (BGP).

BGP

BGP is used by many ISPs or by really large corporations. Small to medium networks do not use this protocol. BGP allows multiple paths to ISPs, and is used to connect multiple AS together. AS simply means a network or set of networks that are under the same administration entity. The AT&T network, MIT, your ISP are AS. Inside your truck, or compound is an AS. This protocol uses classless routing.

As we saw earlier, Classless routing protocols extend the standard Class A, B, or C addressing scheme (/8, /16, /24) by using a subnet mask or mask length to indicate how routers must interpret an IP network ID. Classless routing protocols include the subnet mask along with the IP address when advertising routing information. Subnet masks representing the network ID are not restricted to those defined by the address classes, but can contain a variable number of high order bits. Such subnet mask flexibility enables you to group several networks as a single entity in a routing table, significantly reducing routing overhead.

The BGP is the core routing protocol of the Internet. It maintains a table of IP networks or 'prefixes' which designate network reachability among various AS. It is described as a path vector protocol.

BGP is considered a Distance Vector (DV) protocol. DV finds the route with the fewest hops. Vector means the direction to the desired network. BGP makes routing decisions based on path, network policies and/or rule sets. BGP routers share path and other metrics with all nearby BGP routers. These metrics are stored in router tables. When all DV routers have routing table information from nearby routers integrated into their tables, they are considered converged. As high level paths are added or disappear, routing table re-convergence will ripple across various parts of the internet.

IGP

The second basic routing philosophy is called Interior Gateway Protocol (IGP). Whereas EGP was concerned with routing between various AS, IGP is concerned with routing within an AS. IGP routing is what would be in play within a venue compound if routers are in use. There are two routing protocols used with IGP. Distance Vector (DV) which we looked at briefly and Link-State (LS).

DV

There are three DV routing protocols used by IGP. They are Routing Information Protocol (RIP & RIPv2), and IGRP.

RIP

RIP is very common, as it is the simplest to implement. RIP was invented in the late 1970s and early 1980s. It is the oldest of the routing methods. Its basic metric is to use hop counts. This value is usually defaulted at 15. This means that as packet is passed from router to router, each router increases the hop count metric by one. When a router receives a packet that has its hop count at 15 it discards it, thus preventing "loops". That router sends an update to all other routers in the network that the particular route was unreachable. The router will start an internal timer before it

will accept information from other routers that a destination has become reachable, and to try that path in the future.

The drawback in using the hop count metric is that this limits the size of the network that can be implemented using RIP. In addition, RIP didn't allow for Classless routing, it only understood the traditional Class A, B, and C addresses. Techniques used by RIP to maximize routing efficiency include that as RIP routers advertise available routes among routers, a RIP router remembers where it learned of a new route, and will never advertise that route back to the router that it received the information from. This also stops loops from occurring.

RIP does not respond rapidly to changes; they can flood the network with info as they update. Most routers know how to do RIP, but OSPF is replacing many RIP interior routers.

Another aspect of RIP that limits network size is that individual routers update each other, the default is usually every 30 seconds. Network traffic can become large from update traffic as router count is increased. Again, when all DV routers have information from other routers in their tables, they are considered converged.

Originally, each RIP router transmitted full updates every 30 seconds. In early deployments, routing tables were small enough that the traffic was not significant. As networks grew in size, however, it became evident there could be a massive traffic burst every 30 seconds, even if the routers had been initialized at random times. It was thought, as a result of random initialization, the routing updates would spread out in time, but this was not true in practice. It has been shown that without slight randomization of the update timer, the timers synchronized over time and sent their updates at the same time. Modern RIP

implementations introduce deliberate variation into the update timer intervals of each router.

There is a newer version of RIP known as RIPv2. This newer version allows classless routing. RIPv2 also allows routers to know if traffic is external (via EGP) or internal (IGP).

IGRP

Interior Gateway Routing Protocol (IGRP), invented by Cisco, overcomes RIP hop count limit. IGRP has a default hop count of 100, but can be set as high as 255. It also adds more metrics to determine the "cost" of any particular route by considering metrics such as a path's available bandwidth, the transit time, current path traffic, and even packet size. But IGRP also can't be used in a classless routing situation, as it has no field in its tables for a subnet mask.

EIGRP

Enhanced Interior Gateway Routing Protocol (EIGRP) is also a Cisco proprietary routing protocol loosely based on their original IGRP. This only works on Cisco routers. EIGRP was developed for enterprise wide routing environments. EIGRP is an advanced distance-vector routing protocol, with optimizations to minimize both the routing instability incurred after topology changes, as well as the use of bandwidth and processing power in the router. Routers that support EIGRP will automatically redistribute route information to IGRP neighbors by converting the 32 bit EIGRP metric to the 24 bit IGRP metric.

LS

LS (Link State) is a shortest path-first protocol. LS Routers each create 3 separate tables.

The first table lists directly attached neighbors, built using "Hello" packets. A second table determines the topology of the entire network, which is built using Link State Advertisement (LSA) or Link State Packets (LSP), listing every available network or neighbor use. Unlike DV routers, which only know about their neighbors, LS routers know about their neighbors, and tell the

whole network about those neighbors.

LS routers then run a series Shortest Path First (SPF) tests and generate a third table, which along with the first two tables form the actual routing table. LS routers know more about the internetwork than any DV protocol.

OSPF

Open Shortest Path First (OSPF) is the main LS approach. LS routers continually monitor their neighbors by sending tiny messages, called hellos, and share detailed info, called link state advertisements. If a connection is lost or created, the routers share this information with their neighboring routers. OSPF is a dynamic routing protocol for use in IP networks. As an IGP protocol, it also operates within a single AS. OSPF works within a single routing domain or AS. It gathers link state information from available routers and constructs a topology map of the network. OSPF detects changes in the topology, such as link failures, very quickly, and converges on a new loop-free routing structure within seconds. The topology built among OPSF routers determines the routing table presented by the network's edge router, an Internet Layer above it. OSPF supports classless routing, an unlimited hop count, and OSPF routers converge very quickly. As with all engineering feats it comes with trade-offs; it is more costly and more difficult to troubleshoot when things go wrong.

Truck Router Setup

Most routers found on trucks today will combine many functions into one. While we will look at VPNs later in the application layer, let us say that Virtual Private Networks all secure communications between separate LANs by creating encrypted "tunnels" through the internet so that you can use the internet like it was a dedicated link between various points. Many routers allow a number of VPNs to be established between it to fixed LANs, such as to the truck vendor's headquarters and the truck client's NOC, and also to a number of mobile users, such as those using laptops.

VPNs can be setup between two VPN capable routers (called a Gateway-to-Gateway VPN Tunnel as routers are referred to as Gateways) or between a VPN router and a PC running a VPN client, or that has an OS with a built in IPSec Security Manager or

3rd party VPN application software. IPSec is the security protocol that VPN uses.

Not only are routers layer 3 devices, but they also have layer 2 switch capabilities. While only offering a few switch ports as outputs, they naturally become the root switch for other switches on the truck. These ports are now operating in the gigabit range. The router allows setup of VLAN, which is a layer 2 service.

These routers also offer QoS features that allow for consistent video and voice traffic through and out of the truck, if desired. These routers also act as DHCP, NAT, and Proxy servers, can serve up the time to clients via Network Time Protocol (NTP), and offer built in firewall protection. A part of the firewall protection offered by these routers is via signature protection. This is a file inside the router that looks for malicious known data patterns indicating an attack, and blocks them. This file can contain thousands of rules and must be updated periodically as new threats come along.

Firewalls often have default settings that limit Java and ActiveX programs from running, and cookies from being stored, which could greatly restrict the use of many internet sites. Conversely, truck side users might be able to get to websites you want to restrict, if the router allows users to access proxy servers on the web.

Sometimes, the truck router will be assigned a dynamic IP address from the upstream provider, either from the venue itself, or from the current ISP in use. On occasion, the upstream provider will offer a static IP address that will have to be entered manually. Most routers default to using DHCP to obtain a temporary address. If your connection to the ISP is via DSL, this means your first hop towards the internet is to the phone company, which usually uses what is known as a Point to Point Protocol over the Ethernet (PPPoE).

Remember that NAT prevents the truck-side LAN IP addresses from being seen on the internet side of the router, and all that is seen is the IP address communication with the internet side of the router. But most routers allow even the internet side IP address to be invisible to the outside world.

To set up routers, most have internal web servers that present setup pages to a browser. The setup browser/PC is on the truck side. Many routers use private class C IP addressing, and often their IP addresses are either 192.168.0.1 or 192.168.1.1. So to protect the router from unwanted “tweaking”, keep the router’s password safe, and put the router administrator port on a VLAN restricted to key truck personnel. Some ISPs will want the hostname and domain name fields entered. Your domain name is the top level, and next level parts of your address. Top levels are .com, .tv, .net, etc. so your domain name would look like “trucksareus.com”, or “tvonwheels.tv”. Your host name is your sub-domain name such as “ourunit10”. So the complete internet name for a truck might be “ourunit10.trucksareus.com”. If the “ourunit10” sub-domain isn’t used and www.trucksareus.com is used, then “www” is the hostname.

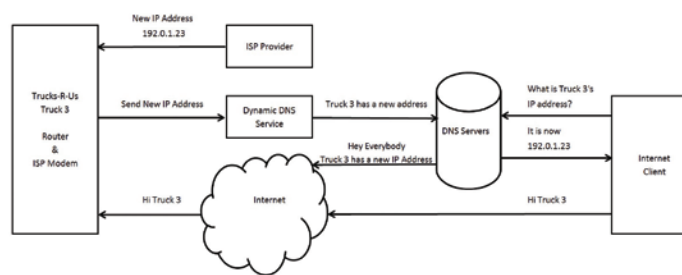
Most routers on trucks will operate in the “Gateway” mode. Here all devices on the truck side share the same IP address that the router has obtained from the internet or WAN side. If the router is set to the “Router” mode, then its NAT functionality is disabled as it expects another router to be upstream of it. Also, all nodes towards the truck side must now have static IP addresses. The router can also be set to use RIP to build and keep router tables in the “dynamic routing” mode, or you can build your own routing tables in “Static Routing” mode.

Most routers must also be told to enable inter-VLAN packets to be allowed between VLANs on different subnets. Many routers have internal diagnostics that allow the router to ping other devices, to run Traceroute tests, and even run diagnostics on the CAT cable connected to its ports.

There is another issue unique to remote trucks: You usually move from one ISP provider to another if your connection to the outside world is via some flavor of landline. While your truck might have a wireless WAN provider that is piggybacked onto a cell network, so you are always behind the same ISP, if the venue or client is

providing your connection, it’s harder for anyone on the internet to find you, because you are not registered via the new ISP. While you can reach out from the truck and request web pages, email and other services, anyone looking to initiate a connection with you won’t know how to find the temporary ISP IP address currently assigned to you.

The answer is called Dynamic DNS (DDNS), a service that is usually provided by a third party. A router that supports this allows it to automatically determine the IP address it has been assigned and to contact the DDNS service provider with that information. The provider then uses that information to update backbone DNS servers, which propagate out to the rest of the internet. Many routers have clients built in for using some of the most common DDNS providers. If you have a router that doesn’t have the built in client, providers will provide a client to run on a PC.



When the truck arrives at a new location and receives a new IP address from the local IP provider, and if a Dynamic DNS (DDNS) service provider is subscribed to, the new IP address is sent, if the truck’s router is setup to do so. The DDNS makes the DNS servers that serve the internet aware of the new address so that when someone from outside the truck reaches out to the truck they will find them. Remember that the truck can always reach out on its own using NAT and PAT to initiate contact. But if someone from the outside wants to initiate contact, a DDNS type service is required.

Layer 4—Transport Layer

It is at this 4th OSI level that end to end communications are addressed. Up until now we’ve been concerned with getting data packets from point to point, but not communications from a sender to its final intended recipient. The transport layer has the protocol that does that. It can both deliver to the end user and then let the sender know it arrived, or just make its best effort at delivering and if it fails it fails. In terms of old fashioned mail, you can request a return receipt on delivery, or hope for the best. But like snail mail, the fact that you didn’t request a return receipt doesn’t stop you from calling the intended recipient and asking if they received your mail. In the IT world you can go to a layer above the Transport layer to check on message receipt.



Continuing with our ‘box within-a-box’ or ‘envelope’ analogy, at the transport layer our IP frame is surrounded by a TCP frame, with is the predominant protocol in use at this layer.

But in the IT world, message receipts are the norm, and the protocol that has won out in that task is called Transport Control Protocol (TCP). With it, all transport of data packets have the sender receiving return receipts from the receiver. If the sender does not receive a receipt back it resends the lost packet(s).

Today the only reason you might see a different transport layer protocol, like Novell’s old SPX, is for security. If hauling packets over longer distances, a path might be setup where a router unwraps TCP packets and re-inserts the data into SPX packets. Any TCP sniffing along the way won’t understand the SPX wrappers. At the far end, the packets are rewrapped into TCP.

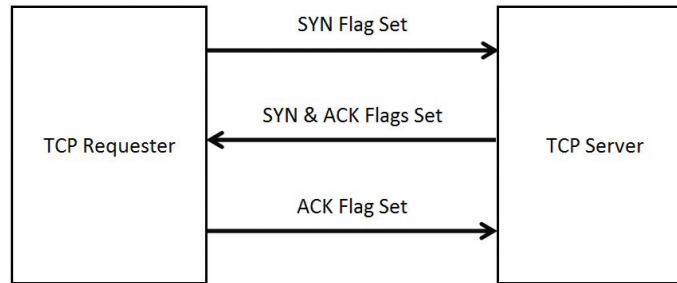
Since TCP is in the layer above IP, and they are almost always used together, you will often see the two protocols written together, TCP/IP. These two protocols are bound together. You will hear IT folks talk about protocol stack binding. This is what is meant. In the network setup sections of the OS you tell the OS what upper layer protocol is handed to, and wrapped in, and what protocol is in the layer below it. Actually, in most cases today, the stack binding is more correctly TCP/IP/Ethernet.

TCP

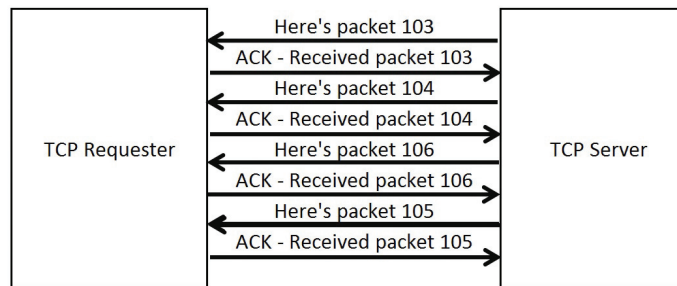
Because TCP provides return receipt to the sender, it is referred to as connection oriented. TCP creates a virtual circuit via a 3-way handshake setup. It uses sequencing, acknowledgements, and flow control. This includes full duplex traffic, error checking, and windowing flow control; this comes with high overhead.

To set up a TCP communication session, a three-way handshake is employed. The sender directs a packet to the intended recipient, which is often aimed at a service provider, such as a server. In that packet, a flag (a single bit) referred to as a SYN flag, is set. The recipient or service provider responds with both the SYN and ACK flag/bits set. The client sends back a packet with just the ACK bit set. A connection is now established. Data transfer can now commence. Every packet sent has a packet count that increments with each packet sent. Every time the recipient receives a packet, it returns an ACK to the sender, along with the packet number received. The sender keeps track of the packets received, and after a timeout, will resend packets that haven't been acknowledged. This means that packets are not always received at the recipient in the order they were originally sent.

The sender and recipient will also negotiate how many packets will be sent before the sender expects to receive an ACK back.



The initiator of TCP communications starts what is known as a 'three way handshake'. First, the initiator sends a packet with the SYN flag (same as a single particular bit) set. The Server acknowledges with both the SYN and ACK flags set. The initiator then responds with an ACK flag, and the communications session is set up.



Communications proceeding between the requestor and server. The server sends a numbered packet, which the requester acknowledges, and the requester also sends a request for the next packet. The transfer of information proceeds along an orderly path, with requests for the next packet, and the packets arriving. Notice in this graphic, that sometimes a packet will arrive out of order. Since all are numbered, the receiver is able to put them back in order. If the receiver doesn't receive a requested packet within a fixed time, it will ask for it again. The data transfer can be very "bursty", with the server serving packets to the requester until its buffers are full. The requester then stops asking for packets until it has processed what it has already received.

TCP Header—24 bytes	
Bits	
16	Source Port
16	Destination Port
32	Sequence #
32	Ack #
4	Header Length
6	Reserved
6	Code bits
16	Window size
16	Checksum
16	Urgency
32	Options

UDP

There is actually a second protocol in addition to TCP that is still in constant use, UDP. Where TCP is considered connection-oriented, UDP, or User Datagram Protocol (UDP) creates a Connection-less connection. A common transport/layer-4 protocol, UDP has no return receipt. It's often referred to as a "ship and forget" connection. This method is un-subsequenced, and has no virtual circuit like TCP, but has low overhead, No ACKS, and no flow control.

Why use UDP when TCP is available? Primarily for reasons of overhead; when sending media, especially video, the sheer amount of data often precludes the luxury of making sure the previous data got to its location before shipping more. Just as video has been transported since its inception, it goes into a pipe and hopefully enough arrives intact to re-create the original.

As we alluded to earlier, you can often use layers higher in the OSI stack to discern if data is missing. Often this is done by adding extra data, in the form of forward error correction, to be used to determine what didn't make the trip intact.

UDP Header—8 bytes:	
Bits	
16	Source Port
16	Destination Port
16	Length
16	Checksum

Here you can see that an UDP header has a lot less overhead than TCP headers.

Ports

So at this point, we have laid the groundwork to get data from one point to another reliably. How does the receiver know what to do with that data? Is it video, or email, is someone requesting a web page, or is it just trying to ship a file from one location to another? In other words is this media, SMTP, HTTP, SNMP, or FTP traffic? We specify the type of traffic by adding an additional tag to the address, a port number. This is analogous to adding an apartment number to a snail mail address for an apartment building, or connecting flight numbers on your flight out of Rochester. While everyone on the plane going to an airlines hub in Chicago will arrive for a while at O'Hare, it is the second flight number that specifies where each person on the first flight will finally end up.

In the IT world, the port number signifies to the OS what software is expected to handle the data. The most common that many are familiar with, is port 80. If data arrives with 80 as its destination port number, that data will be directed to a web server running on the box. Port 21 belongs to FTP software; Port 53 indicates that someone's requesting info from a DNS server; Port 67 indicates a machine looking to be assigned an IP address. There can be tens of thousands of port numbers. Some software/services will respond to multiple port numbers.

Common layer 4 ports:			
21	FTP	80	HTTP
22	SSH	110	POP3
23	Telnet	119	News
25	SMTP	123	NTP
53	DNS	143	IMAP4
67	DHCP/BootP	161	SNMP
69	TFTP	443	HTTPS

If you're a web server back at ESPN, you're receiving Port 80 web page requests continuously. How does it keep track of requests from Grand Rapids, Fort Meyers, and Boise? While the requesting machine would use port 80 as its destination to alert the web server on the receiving machine that it's after web services, the sending machine would pick a random, non-reserved port number as its sending port. The receiving web server would then associate that random port as the person in Fort Meyers, and another random received port number as another requester, in Boise. The web server doesn't know, or care, if you're from Boise or Rochester, it's only the IP address, and the TCP port its paying attention to.

In addition, when you connect to a web server to download a page, that session is not persistent. Once the server has delivered the requested page, it severs the current relationship. Often, to sort out random requests coming in that happen to have selected the same port numbers, either internal to the web server, or often separately in another box, there will be a proxy server that will translate duplicate requesting port numbers so that the web server can cope. It should be easy to see why a determined hacker would be able to flood a web server with web requests, known as a Denial of Service attack, and bring a web service to its cyber-knees.

An important fact regarding TCP and UDP sockets: a TCP port is distinct from a UDP port with the same number.

Layer 5—Session Layer

This is the OSI layer, the 5th up in the stack, which is concerned with how an application talks over the network. Port numbers are used to route lower network traffic to the right application at this layer. The main interface between applications and all the network services below is the Network Basic Input/Output System (NetBIOS). NetBIOS is a set of Application Programming Interfaces (API). APIs are calls that an application can make to move data into the network stack below, or extract data out of that stack. APIs can be thought of as a defined tool set for programmers writing applications. The OS vendor, Microsoft in most cases, has written these interfaces so that the application developer doesn't need to worry about the intricacies of network software below. Most GUIs you have up on the screen in front of you are part of an application that is making calls to the network API. It is the API that talks with the TCP/IP or UDP/IP code below in the Transport Layer.

NetBIOS has been around almost as long as the IBM PC. It was originally 3rd party software, and was originally used to talk to an early Ethernet LAN competitor, Token Ring topology. Eventually, Novell adopted NetBIOS for their IPX/SPX protocol. Today it is TCP/IP or UDP/IP that NetBIOS talks to. This demonstrates how the OSI layer concept can isolate the upper layer from lower ones. NetBIOS manages connections based on the names of the computers involved. Microsoft abandoned NetBIOS with Windows 2000.

Microsoft supported two distinct session layer protocols. We just discussed NetBIOS; the second, Sockets, is almost exclusively used today and is often called WinSocs. You can think of a socket as just that, a connection that an application can make to the OSI stack below. Microsoft kept the use of NetBIOS for a long time, so applications

would not have to be rewritten. Whereas NetBIOS used a name and an extension, Sockets uses an IP address and a port number. Sockets are a new set of APIs that have the same concept as NetBIOS did.

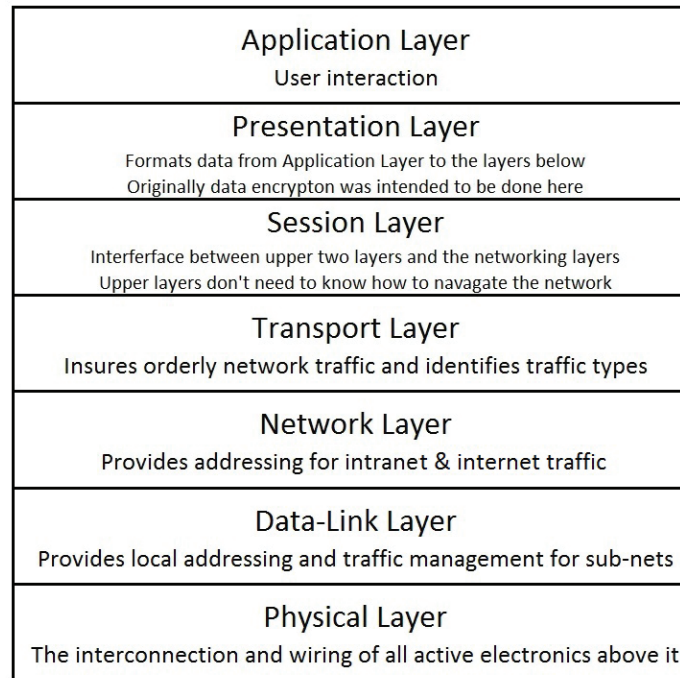
Layer 6—Presentation Layer

We are now on the 6th OSI layer. This layer can be important; its main intent is to translate data between applications and the services below this layer. Encryption was intended to be done here, and still is, but data encryption today is found on any of the top five OSI layers. The Presentation layer is expected to format data it receives from the application layer for handoff to the services below. When an application is compiled from a high level language into machine byte code the compiler often works at this layer also.

Layer 7—Application Layer

As the name implies, this is where the software resides that gives computers their utility, and what users interface with. As we mentioned in the session layer, applications call APIs, both in the presentation and down into the session layers. APIs are a special set of commands that allow applications to request services from the operating system. APIs makes it easy for programmers writing applications to access the network without knowing any of the details of the network.

Many protocols are implemented at the application layer. HTTP protocol is a common application layer protocol, as are SMTP, POP, SNMP, etc.



Servers

An important distinction at the application layer is between clients and servers. Microsoft has client, or workstation, and server versions of their OS. Workstations can act as servers but lack some of the more powerful features of the server versions. Servers place greater priority on remote requests than on local requests made at the server.

As we've discussed, you can run across many types of servers, like DNS, DHCP, etc. Often many functions such as DNS and DHCP can be combined into a single box, such as a wireless router.

Servers are used to organize, and provide services to local networks: from DHCP that hands out addressing to nodes/hosts, to DNS acting as the network's phone book, or to defining Autonomous Systems (AS). The concept of an AS as set up by Microsoft is the Domain. This is a set of resources and hosts that are treated and administered as a single system. The basic building block is a Domain Controller (DC), which often has a backup. A host authenticating with a DC is granted a certain set of

user rights across the extent of the domain. Each user can be given a different set of rights. Rights give you access to services and other hosts. This allows a single login and authentication to acquire all of a user's access and rights.

These rights and privileges are stored in directory services, which store information about a network's resources, including users, applications, files, and printers. When Windows 2000 was launched, directories' services were then renamed Active Directory. A domain server running Active Directory has to also be running DNS, as they can't be separate services.

A domain is given a name. The name of the truck could be the name of the domain.

If a LAN has no machines running a server grade OS, then each system on the network must act as its own server. This means users must have an account on each system they want to access. In a domain based network no one has a local user account, all users get domain user accounts that give them access to the network with a single login. These domain accounts are set up by software on the server system.

The difference between servers and clients that are part of a workgroup, versus ones that are part of a domain, can often be confusing. With Workgroups, all computers are peers; no computer has control over another. Each computer has a set of user accounts. To log on to any computer in the workgroup, you must have an account on that computer. There are typically no more than twenty computers in a workgroup. A workgroup is not protected by a common password. All computers must be on the same local network or subnet.

With Domains, one or more computers are servers. Network administrators use servers to

control the security and permissions for all computers on the domain. This makes it easy to modify, because the changes are automatically made to all computers. Domain users must provide a password or other credentials each time they access the domain; if you have a user account on the domain, you can log on to any computer on the domain without needing an account on that computer. You can probably make only limited changes to a computer's settings because network administrators often want to ensure consistency among computers. There can be thousands of computers in a domain and the computers can be on different local networks.

Proxy Server

Another common type of server is a Proxy Server. A proxy server is either a complete computer system or an application program running on a server that acts as an intermediary for requests from clients seeking resources from other servers. A client connects to the proxy server, requesting some service, such as a file, connection, web page, or other resource, available from another server. The proxy server evaluates the request according to its filtering rules. For example, it may filter traffic by IP address or protocol. If the request is validated by the filter, the proxy provides the resource by connecting to the relevant server and requesting the service on behalf of the client; it acts as a proxy for the requester. A proxy server may optionally alter the client's request or the server's response, and sometimes it may serve the request without contacting the specified server. In this case, it 'caches' responses from the remote server, and returns subsequent requests for the same content directly.

A Proxy server can translate addresses like a NAT service, but does so in a very different way. Besides just translating IP addresses, it also can translate one port number to another. This means that applications like email and web browsers must be proxy aware in that they must be able to change their standard ports to whatever the proxy server uses. Changing ports increases security. Instead of port 80 for web pages it can be some other port, which the proxy will change to the right port before sending out.

Proxy servers can also be used to cache commonly used web pages, and can also be used to filter web access and services such as IRC chat. Unlike a NAT device, a Proxy server works at

the application layer. These are sometimes known as application gateways. Proxy Servers work slower than NATs because they translate IP addresses and port numbers both. A Proxy Server is considered part of the firewall in a network. It keeps the machines behind it anonymous, mainly for security.

Firewalls

Firewalls are usually a combination of hardware and software. The hardware part is usually a router, but it can be a computer or dedicated piece of hardware (a black box) with 2 NIC cards in it. One NIC connects to the public side and one connects to the private side. The software part scrutinizes each incoming and outgoing packet and rejects any suspicious ones. They can also permit, deny, encrypt, and proxy all traffic that flows through. They can be used to create different zones within an intranet.

Firewalls can stand between the inside network and the outside internet, between the local network and a resource, such as a server, or simply between different network segments.

Firewalls maintain Access Control Lists (ACL). These filter traffic based on source and destination IP addresses, protocols in use, source and destination port numbers, and packet type.

Stateful Firewall—this keeps track of connections/sessions via TCP's 3-way handshake. These firewalls tend to be a bit slower at establishing connections, but after a connection is established, they are usually faster because they just have to check the session state table for the connection instead of comparing the packet against all the relevant ACLs.

Stateless Firewall—this only looks at individual packets and applies them to an ACL. It doesn't care about the context the packets are in. It's simpler and requires less memory. These are best used on an internal network where there are fewer security threats.

Many firewalls will use Dynamic Packet Switching, which ensures that packets they forward match sessions initiated on their private side by something called a dynamic state list or state table, which keeps track of all communication sessions between stations from inside and outside the firewall.

This list changes dynamically as sessions are added and deleted. Thus, only packets for valid and current sessions initiated inside the firewall are allowed to pass.

Firewalls can also use content Filtering:


Attachments—all .exe files are blocked; content-encoding; email headers; language phrases; proximity of words to each other; URLs, and finally Bayesian—or the probability that some combination of the above might spell trouble.

Many firewalls are implemented at the Application level because they are more aware of data use. They not only know the network and transport layer contexts but also the application context such as FTP, SNMP, HTTP, etc. A drawback is that these firewalls are slightly slower. Application firewalls can scan actual payloads for malware and other content.


A host based firewall is installed on a host and only protects that machine; this software based firewall isn't usually as robust as a dedicated firewall.

When a network has resources that must be accessed from the outside, a DMZ zone is usually setup. This is done either with a single computer or router or with two machines. The purpose is to allow outside users, along with inside users, access to resources such as email, ftp, web, DNS servers, and often a proxy server, but still keep outside users out of the internal network.

In the single firewall approach, there are three NIC cards, one outward facing, one inward facing, and one port to the resources between the inside and outside networks. Traffic from the outside is only allowed to the resources in the DMZ. Traffic from the inside is also routed only to DMZ resources. Any web traffic generated from inside must go through the proxy server in the DMZ. Any traffic with malicious intent must get through these DMZ servers. If designed and maintained correctly, that should not happen. The bad guys may cause havoc in the DMZ but shouldn't be able to get to the inside.



Conceptual view of the DMZ firewall topology. Outside and inside users meet only inside the DMZ. Inside users who want to communicate outside the firewall must negotiate through the Proxy.



The way a DMZ is usually physically setup: A port facing outward and a port facing inward, and a port into the DMZ. For simplicity, we show two ports into the DMZ, but often, all devices inside the DMZ are on the same subnet or port.

For extra security, dedicated firewalls between the outside and the DMZ, and the DMZ and the inside can be employed. This would require hackers to breach two sets of firewalls. As an aside, the inside and outside firewalls are usually running over the same physical cabling, as many of the servers employed in the DMZ don't use separate NICs for outside and inside traffic.

VPNs

Virtual Private Network (VPN) is a tunneling protocol. It allows a host to traverse an insecure network, such as the internet, and become what appears to be local to a remote network.

There are three types of VPNs:

Remote Access VPNs—Remote users connect to secure corporate networks

Site-to-site VPNs—Intranet VPNs. These allow corporate remote sites to connect to the corporate backbone securely over a public medium like the internet, instead of requiring a more expensive WAN connection such as frame relay.

Extranet VPNs—Allow an organization's suppliers, partners, and customers to be connected to the corporate network in a limited way for business-to-business (B2B) communications.

While a VPN is at the Application layer, it uses the network/layer 3 or transport/layer 4 insertion. While there are a number of methods for securing VPN traffic, the most common is one we looked at earlier, IPSec. Others are Transport Layer Security (SSL/TLS), Secure Socket Tunneling Protocol (SSTP), Point-to-Point Protocol (PPP), and Secure Shell (SSH).

To facilitate many VPN connections, you'll sometimes find a VPN Concentrator (VPNC). A single device like a VPNC will handle a very high number of VPN tunnels, into the hundreds or even thousands. A VPNC will be used for remote-access or site-to-site VPN. These were specifically developed to address the requirement for a purpose built, remote-access VPN device. Encryption for a remote-access VPN through a concentrator is usually handled by IPSec or SSL, and user authentication can

be achieved via Microsoft's Active Directory, Kerberos, RADIUS, RSA and digital Certificates. Many VPN Concentrators also have a built-in authentication server and allow ACLs to be implemented through them.



Server Farm

Home Port Support

As we mentioned earlier during earlier chapters, namely Setup and Logistics, a truck will often have wireless connectivity that is carried over a cell phone service. This service is usually with something called WiMax, which is essentially a city wide Wi-Fi to create WANs. Often this wireless service is implemented with an "AirCard", a wireless modem that is plugged into a computer via a USB or PCMCIA connector.

Some truck vendors have software running on a PC in the truck that will phone home when the truck is parked and powered. Software at the truck vendor's headquarters is starting to resemble rudimentary NOCs. Software at these NOCs can be linked to operations software supporting a database so data coming from the truck can be correlated with its job details. By identifying who is assigned to the truck that day, it can send fault notices via text, email or a cell call, to contact the folks who may be back at the hotel, if something is wrong. If the truck is on site and has a day off or is dark over the weekend, it will continuously monitor systems and check for alerts for any faults, power outages, or security breaches.

Third parties, such as equipment vendors, can log into the truck systems, either via applications or web pages, and be granted permission to work on particular equipment in a specific truck. Some truck vendors can also talk on the truck's intercom, as most use VoIP these days. They have control of a video/audio router buss, which can be used for diagnostic feedback, or to call up any of the truck's internal security cameras, or even a small wireless camera on board to help with diagnostics or parts identification.

Conclusion

No other area of truck activity has changed as dramatically as computer networking. Almost the entire workflow that comprises all aspects of remote television production is controlled by software running on computer hardware connected via network technology. The amazing thing about modern network technology is that only the top layer of the OSI stack, namely the application layer, is particular to a television production truck. The bottom six layers are almost totally generic. Despite this, the learning curve for the average truck engineer has been steep. As we have pointed out earlier, it's important to master this area of technical endeavor, or the long term prospects for truck employment are slim.



Almost every piece of equipment in these two racks is masquerading as television gear but at their heart they are PCs.

Network Troubleshooting

Most trucks don't travel with network test equipment to facilitate network troubleshooting. Luckily, there are a number of basic command line functions in Windows that facilitate basic troubleshooting. Let's look at a few.

NSLOOKUP

NSLOOKUP allows you to query a name server and find out which name resolves to which IP address. The first time it is typed, it displays the name and IP address of the default DNS server.

The prompt then changes to just ">" (no path).

If you now type a domain name, it will continue to give you the same DNS server and its IP address and then the domain's server name and the IP addresses associated with domain name.

Set type=mx to lookup a mail server, if there is one

type=a for A records

type=any for all records.

Type ? at the > prompt to see all possible commands associated with NSLOOKUP

To check and see if your DMZ firewall is allowing port numbers you need to use, you can try to telnet to the DMZ server using the desired port number, and see if it responds.

DNS Issues

If you're having problems reaching URLs or mail servers etc., on a particular machine, first try flushing the local DNS cache using ipconfig/flushdns. Ping can be used to see if the local DNS server is resolving IP addresses by pinging a known URL. If an IP address is provided, try pinging the same site with its IP address. If it is found, this tells you that you might have DNS problems. Next, try checking your DHCP settings with ipconfig as it might be handing out wrong DNS server information.

arp -a To see DNS entries in your computer's ARP table

Try using the IP address instead of UNC to get around any DNS issues.

Network Stack

The network protocol stack could be uninstalled from a machine but the browser will still be able to view a local HTML document.

Often, it is helpful to be able to change the order in which a computer uses and lists its available NIC cards. To change the NIC Binding Order:

1. Click Start, click Run, type ncpa.cpl and then click OK, or open the control panel and select Network Connections. You can see the available connections in the LAN and High-Speed Internet section of the Network Connections window.
2. On the Advanced menu, click Advanced Settings, and then click the Adapters and Bindings tab.

3. In the Connections area, select the connection that you want to move higher in the list. Use the arrow buttons to move the connection.

To change settings in the network stack:

cmd line enter ncpa.cpl or right click on network adaptor and select the IP4 stack, then Properties.

Here you can choose to use DHCP IP addresses or a manually supported one. If the Advanced button is selected, you can set a default gateway address, or DNS suffix to use (currentvenue.com), or set WINS default addresses and how NetBIOS is used.

TCP/IP diagnostic procedure

1. To see if the NIC is working, ping 127.0.0.1 or ping localhost.
2. To determine your computer's IP address and gateway type "ipconfig".
3. Ping some local machines using their IP addresses and DNS name.
4. Use "net view" to see local area network systems.
5. Make sure that the net mask is correct by entering "ipconfig".
6. Try renewing the DHCP lease by entering "ipconfig /renew all".
7. Determine your computer's name by going to Start and Run and enter "msconfig"
8. Select the "Tools" tab and then "System Information" and the "Launch" button.
9. Typing "hostname" will also give you the computer's name.

7. Run `netstat` which will show current network connections to your system.
8. Run `netstat -s` which will show you network statistics. This is good for seeing if you are sending and not receiving or vice versa.
9. If you can't get to the internet, see if you can ping the router. The router has two interfaces: one to the local net and one facing the internet. If you can't ping the far side of the router, then something is wrong with the router.
10. Use `tracert` to trace route over the internet to a location on the internet. `Tracert` is a Windows command prompt (Microsoft version of `Traceroute`). Enter `tracert <DNS name or IP address>`

Network Activity

Displays active TCP connections, ports on which the computer is listening, Ethernet statistics, the IP routing table, IPv4 statistics (for the IP, ICMP, TCP, and UDP protocols).

`netstat -a` Displays all active TCP connections and the TCP and UDP ports on which the computer is listening.

`netstat -e` Displays Ethernet statistics, such as the number of bytes and packets sent and received. This parameter can be combined with `-s`.

`netstat -n` Displays active TCP connections, however, addresses and port numbers are expressed numerically and no attempt is made to determine names.

`netstat -o` Displays active TCP connections and includes the process ID (PID) for each connection. You can find the application based on the PID on the Processes tab in Windows Task Manager. This parameter can be combined with `-a`, `-n`, and `-p`.

`netstat -p <protocol>` Shows connections for the protocol specified by Protocol. In this case, the Protocol

can be `tcp`, `udp`, `tcpv6`, or `udpv6`. If this parameter is used with `-s` to display statistics by protocol, Protocol can be `tcp`, `udp`, `icmp`, `ip`, `tcpv6`, `udpv6`, `icmpv6`, or `ipv6`.

`netstat -r` Displays the contents of the IP routing table. This is equivalent to the `route print` command

`netstat -s` Displays statistics by protocol. By default, statistics are shown for the TCP, UDP, ICMP, and IP protocols. If the IPv6 protocol for Windows XP is installed, statistics are shown for the TCP over IPv6, UDP over IPv6, ICMPv6, and IPv6 protocols. The `-p` parameter can be used to specify a set of protocols.

`netstat <interval>` Re-displays the selected information every Interval seconds. Press CTRL+C to stop the re-display. If this parameter is omitted, `netstat` prints the selected information only once.

`net view` to see all local network hosts.

`net view <host>` to see shared resources on a local network host.

The PathPing tool is a route tracing tool that combines features of Ping and Tracert with additional information that neither of those tools provides. PathPing sends packets to each router on the way to a final destination over a period of time, and then computes results based on the packets returned from each hop. Since PathPing shows the degree of packet loss at any given router or link, you can pinpoint which routers or links might be causing network problems.

NetBIOS

`Nbtstat` is designed to help troubleshoot NetBIOS name resolution problems. When a network is functioning normally, NetBIOS over TCP/IP (NetBT) resolves NetBIOS names to IP addresses.

`nbtstat -c` shows the contents of the NetBIOS name cache, which contains NetBIOS name-to-IP address mappings.

`nbtstat -n` displays the names that have been registered locally on the system by NetBIOS applications, such as the server and redirector.

`nbtstat -r` displays the count of all NetBIOS names resolved by broadcast and by querying a WINS server.

`nbtstat -rr` sends name release packets to the WINS server and starts a refresh, thus re-registering all names with the name server without having to reboot.

`nbtstat -s` list the current NetBIOS sessions and their status, including statistics.

Routing

Sometimes you will need to manually tell routers how to route your traffic. That is done using the route command. Route is used to view and modify the IP routing table.

Route Print displays a list of current routes that the host knows.

Route Add adds routes to the table.

Route Delete removes routes from the host's routing table.

Routes added to a routing table are not made persistent unless the -p switch is specified. Non-persistent routes only last until the computer is restarted or until the interface is deactivated. The interface can be deactivated when the plug-and-play interface is unplugged (such as for laptops and hot-swap PCs), when the wire is removed from the media card (if the adapter supports media fault sensing), or when the interface is manually disconnected from the adapter in the Network and Dial-up Connections folder.

For two hosts to exchange IP datagrams, they must both have a route to each other, or they must use a default gateway that knows a route between the two.