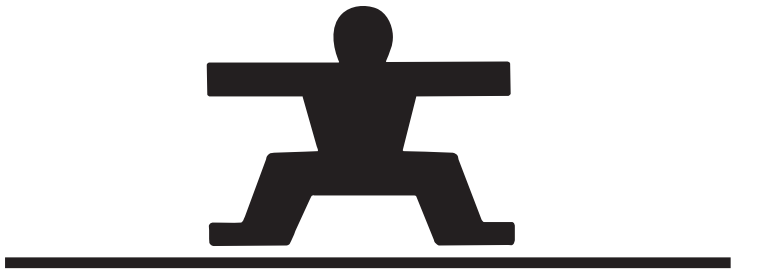


# A Not-So-Short Introduction to Networking

Andrew Starks
Tightrope Media Systems

With Contributions from:
Chris Bongaarts
Joeseph Bongaarts



December 13, 2005

# 1 Introduction

If you are not a network person by profession or hobby, this paper may serve as a helpful guide in understanding some basic concepts. We overview networking rules, conventions and issues. This guide was written with Tightrope installations in mind, but the basic concepts are universally applicable.

None of the information in this paper should be considered complete. Many difficult concepts are glossed over and oversimplified. Our goal is to provide you with useful information that may lead you to accomplishing your goals—or at the very least communicate those goals to someone who can do it for you.

# 2 The Basics: What is a Network?

A computer network is a group of computers connected by hardware and software that are able to *directly* communicate with each other.

The word 'directly' is emphasized in our opening paragraph because it is possible for computers to communicate even when they are on two different networks, but that ability does not make their networks into a single entity—they are two different networks. It is important to limit the definition to "direct communication" because we need to differentiate between the meaning of a single network and a group of networks. Often, people say 'network' when they mean "wide area network" (WAN), which is a term used to describe several networks that are interconnected and under the control of a single organization. Computers can communicate with any other computer on the WAN, but they must do so through the use of networking equipment that is able to route traffic *between* networks within the WAN. This is not the same as "one big network". Some will try to emphasize that you are speaking about a single network by using the term "local area network" (LAN), but this is no more precise than just using the word 'network'. In this document, "local area network" and 'network' mean the same thing. Much of the text clarifies the mechanics of these difference, so if you don't "get it", that's ok; you can come back to this later. :)

On most networks, all computers are considered peers. That is, there is nothing inherently special about a server in the network closet that makes it different than the computer sitting on your desk. It is the roles that we assign these computers that make them special.

These 'roles' are called *services*. They include web servers, email, file servers, time synchronization, security authentication and many others. There are thousands of services that may be provided on any network.

99% of local area networks are using Ethernet hardware and a software networking package[1], which is usually included with the computer's operating sys-

[1] . . . which is called a *stack*. . .

tem.

The key to understanding a network is to see it as one group of computers that is wired together both:

**physically** All of the computers are wired together with network cabling and can electrically 'get' to each other.

**logically** All of the computers are configured so that their network software represents them as "on the same network".

Physical connections are straightforward. In most networks today, you have Ethernet cables (CAT5e) that connect computers to an Ethernet switch, which is a box that is able to send networking singles between computers in an efficient way. When everything is connected to the same switch, you have a simple network that is just begging for the connected computers to be configured in such a way that they can communicate with each other.

The logical side of networking refers to the protocols, software and configuration that make up a network. In the post Internet Boom, almost every network in existence runs on an Internet Protocol (IP) network. IP networks have four basic components that are within the scope of this article[2]. The rest of this section is devoted to those components.

## 2.1 IP Address

If you know one thing about networks, it is probably the concept of the IP address. On any network, every network interface[3] must have a unique IP address[4]. An IP address is a number that, for the sake of readability, is broken up into four numbers separated by a period (.). Example: "`192.168.1.1`".

Each of the four numbers must be between 0 and 255. When we speak of a computer's IP address, it's important to understand that there are many addresses that are not considered valid because they are reserved for special purposes. That is, "`0.0.0.0`" and "`255.255.255.255`" are examples of addresses that are not considered valid. "`127.0.0.1`" is another special address called the *loopback address*. It represents the local computer and is effectively saying, "go to me and ask myself for something[5]".

[2] . . . and a million that lie outside it. :)

[3] The term *network interface* is more precise than saying *computer*, because a computer might have more than one network card or the connecting device might not be a computer. *Network Interface* covers any host on a network.

[4] This is probably true for every scenario that you will come across. There are special cases where a computer might have the same IP address on the same computer, usually for redundancy or performance reasons in high capacity installations. We mention it in a footnote so that when a network nerd corrects you on this point, you can say, "Yeah, I know."

[5] One of my close friends, and one of the biggest nerds that I know, has a bumper sticker that reads, *"There's no place like 127.0.0.1"*. It is at Thinkgeek.com if you want your own.

Another important concept to understand is that in an IP network, the IP address serves two purposes. The first purpose is to identify the network and the second is to identify a specific network interface on that network.

When we say "identify the network", we are alluding to the fact that the complete range of possible IP addresses are always segmented into smaller networks. That is, no valid network has an IP address range of 0.0.0.1 to 255.255.255.254[6]. We break them up into smaller chunks for manageability reasons.

How do we determine the part of the IP address that signifies the network and the part that identifies a specific computer on that network? We do that with the...

## 2.2 Subnet Mask

Put simply, the subnet mask's job is to split an IP address into a network address and a host address. This helps networking software determine when an address falls within the local network and when it does not.

Like an IP address, a subnet mask is a group of four numbers separated by a period. Each of the four numbers must always be between 0 and 255.

The subnet mask is a bit mask. Since computers think in "`0's`" and "`1's`", it can quickly take this number and *mask* it over the IP address. Whenever there is a "`1`" in the mask, it associates the corresponding bit in the IP address as a *"network address"*. When it sees a "`0`" in the subnet mask, the computer thinks *"computer's address"*.

Examine the table below. We can see how a computer is able to separate the network address from a specific computer's address by using a subnet mask.

| | ***Decimal*** | ***Binary*** |
|---|---|---|
| **IP Address** | `192.168.1.1` | `11000000.10101000.00000001.00000001` |
| **Subnet Mask** | `255.255.255.0` | `11111111.11111111.11111111.00000000` |
| **Network Portion** | `192.168.1.X` | `11000000.10101000.00000001.XXXXXXXX` |
| **Address Portion** | `XXX.XXX.XXX.1` | `XXXXXXXX.XXXXXXXX.XXXXXXXX.00000001` |

After reading the last four paragraphs, you should be very confused. If not, you probably don't need to read this paper. :) So let's explain things a little further.

As an example, that is technically illegal, a subnet mask of "`0.0.0.0`" would define any computer IP address as being within your network. Generally this is not a number that you will ever see. Conversely, "`255.255.255.255`" would define a network with exactly one computer on it[7]. As the example in the previous

[6] 255.255.255.255 is special, so we don't include it in the standard IP address range.
[7] Sometimes this value is used when you are using an old-time modem.

table shows us, "255.255.255.0" [8] would mean that all computers that shared the first 3 numbers of their IP address would be on the same network.

> If you are interested in the math, networking software performs a bitwise AND operation on the IP address of the local network interface and the subnet mask. It then performs the same operation using the destination's IP address to determine if they are within the same network. If not, the network communication is sent to the network's router. For a detailed explanation of subnet masks, see Wikipedia.org.

The bottom line is that the purpose of the network's subnet mask is to help computers decide if a communication request is within the network by separating an IP address into two parts: a network address and a host address. If a communication request does not fall within a computer's network, we involve the...

## 2.3 Network Router

A network router, also known as the *default gateway*, sits on the edge of your network and handles traffic *between* networks. When you are at your computer and you ask for a computer that's not in your local network, that request is sent to your router. Using magic[9], your router figures out how to get to the computer that you are asking for and the connection is complete. If it cannot, then you get an error.

Incidently, if your network did not have a router then any request that didn't fall within your subnet would be considered unreachable.

## 2.4 Domain Name System (DNS) Address

When you type an address into your web browser, you typically type a name, such as "www.trms.com". Imagine if we had to remember the IP address for every site that we wanted to visit!

DNS is a service that is able to resolve names to IP addresses. When your computer gets a name instead of an IP address, the computer listed as the DNS server will be asked to translate it.

If you have spent any time with your web browser, you'll notice that all of the names that you type end in a common suffix, like '.com', '.org', '.uk', etc.

[8] All of our examples are with either "255" or "0". Other numbers are possible, but not any number. Specifically, if you read a subnet mask in binary from left to right, all of the "1's" would be packed together followed by all of the "0's". There may be no "0's" before any "1's" in a valid subnet mask.

[9] There is a lot of 'magic' in this document, because we can't cover everything.

These are called *top level domains* and are governed by various organizations, depending on the specific domain in question.

The name that comes directly before the last dot (.) is called the *second level* domain. Examples of these would be "*whitehouse*.gov" or "*trms*.com". Second level domains are managed by various domain name registration companies including Network Solutions and Register.com. If you want to reserve one for your organization, you simple go to one of these companies, search for an available name and pay them some money. When you're picking a name, you can choose the top level domain that is most appropriate for you, naming yourself "`myfantasticname.org`" or "`myfantasticname.com`". You can even be both, if the names are available.

Some top level domains work differently than standard ones like "`.org`" and "`.com`". For example, not just anyone can get a "`.mil`" address as those are reserved for the United States Military. In the United States the "`.us`" works a bit differently for schools in that second, third level and fourth level domains are used to drill down to a specific district. For example, the Bloomington school district in Minnesota is at "`bloomington.k12.mn.us`". The fifth level domain is managed by the school for internal domains, computers and services.

If you have your own domain name, you or your organization is responsible for managing it. You have the responsibility to designate a DNS server that will manage any requests that are sent to your new domain. Your DNS server tells the world, "When people look for www.myfantasticname.org, it is at *this* IP address."

Most often, when you type an address into a web browser, there are three parts to an address, like "`www.trms.com`". The "`www`" in that address is the *third level* domain. Any domain beyond the second level domain is managed by the organization. They are used to designate specific services or computers within an organization's network. As an example, your computer on your desk is most likely named with a third and possibly fourth level domain, depending on the complexity of your network.

Domain levels are like a tree where the top level is the trunk and each subsequent level is a branch that comes off of the previous level. The highest level domain might be considered the leaf, which represents a service or a specific computer. There can be any number of levels to a domain. See figure 1 on the facing page for an illustration of domain levels.

When you refer to a computer's name, you might say "`news`" or be more specific and say "`news.employee.acme.com`". The full name is known as the *fully qualified domain name*.

Domain levels higher than the second are particularly interesting to Tightrope customers who are operating within an established network, as they may provide a means of addressing your server even if there is already an existing one. In fact, this is how Tightrope operates its demonstration site, which is a full-blown Cablecast and Carousel system. If you go to "`www.trms.com`" you get our corporate web site. If you go to "`demo.trms.com`" our network hardware and configuration is able to redirect you to the demonstration system. We will explore this topic further when we talk about firewalls and port forwarding.

## 2.5 Summary of Basic Network Concepts

All computers on a network must have an IP address that is unique so that other computers can find them. All computers on a network will have the same subnet mask, which is just a number that defines the size of the network. Using math, the networking software on your computer is able to figure out if an IP address is inside or outside your network based on this subnet mask. If it is outside the network, your computer will forward your request to the router on your network that will, using magic, forward it off to the computer with which you are trying to communicate. If there is no router defined, then you can only communicate with computers within your network.

DNS servers interpret friendly names that humans understand, like "www.trms.com",

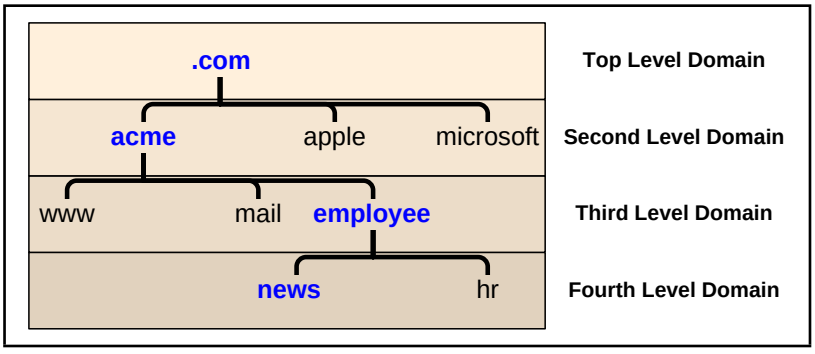


**Figure 1:** The above example illustrates how domain names are structured, where each level drills down into a specific service or computer. To get to the "`news`" computer, you would use the "`news.employee.acme.com`" domain name.

into numbers that computers understand, like "64.122.237.46". If there is no DNS server defined, then you will have to use IP addresses and not names.

# 3 Dynamic Addresses and DHCP

Imagine that you are a network administrator for a new organization that just received 500 desktop computers. Your first task is to get these computers on your network. You begin assigning each computer an IP address, subnet mask, DNS and router information. This would be extremely time consuming, but you get it done thanks to your hard work and dedication to repetitive tasks.

Now imagine that it is a year later. Some of these computers have broken and been replaced. Thirty of them got a virus and needed to be rebuilt. You added another 100 computers to one department and got rid of 30 from another...

You can see how time consuming network administration can become when you have to manually enter the networking information into a large number of computers! Enter a magic technology: Dynamic Host Configuration Protocol or DHCP.

DHCP is a service that runs on a network. If a computer is configured to use DHCP, it will seek out this service and automatically configure its network settings using the DHCP service's instructions.

Typically, if your computer uses DHCP it leases an IP address. This means that your IP address can change after a predetermined amount of time. We call these addresses *dynamic IP addresses*. If your computer is sitting on your desk and you use it for email and surfing the web, this is fine. Nobody cares about your IP address because nobody relies on your computer for any network services.

Dynamic IP addresses are a problem if your computer does have services that people need to find. For example, a web server[10] needs to be at a fixed IP address because people need to know where they can find it.

Using DHCP does not have to mean a dynamic address. DHCP can be configured to make a specific computer's address static. There are other features of DHCP that are worth exploring. Check out Wikipedia.org for detailed information about it.

The bottom line is: DHCP is a service that automatically configures networking on a computer and usually gives it a dynamic IP address. Desktop computers

[10] ...like Tightrope's Cablecast or Carousel computers...

are generally configured with DHCP and servers, like those from Tightrope, can only use this service if the address it gets is static and not dynamic.

# 4 TCP and UDP Glossed Over

Within a network, communication is happening. Computer A is saying something to computer B and they both must agree on how that communication is established and negotiated. There is a protocol to it all, like a handshake when you meet someone new.

There are two primary methods for communicating on IP networks that are within the scope of this paper: TCP and UDP.

TCP stands for Transmission Control Protocol. We don't care why it's called that, but we do care that TCP is the most common way that two computers will communicate on a network. It is the TCP in TCP/IP networks. TCP is a connection-based protocol that is able to guarantee proper communication between two computers. It is a reliable stream of data that is guaranteed to reach the destination with the same data in the same order[11].

We say that TCP communication is *connection-based* because before data is transmitted, a connection is established. Just like picking up a telephone to call your neighbor, communicating with TCP means that you establish the connection first, then you begin communicating. When you're done 'talking', you hang up, or drop the connection. Because there is a constant connection, the two communicating computers are able to monitor the packets of information, ensuring their correctness and that they arrived. The down side is that these connections have a significant amount of overhead that each communicating computer must deal with.

By contrast, imagine that you are leaving your office and before you shut the door you yell out, "Bye! Lock up before you leave!" Did anyone hear you? Who knows!?! You left before you got a response.

UDP is much like yelling in the dark. It stands for User Datagram Protocol and using this method, the computer sending the data packages the information into a datagram and sends it into the network. You have to rely on *hope* that it will get to its final destination[12] because Unlike TCP, UDP does not require a connection before communication can occur. Furthermore, it does not make any promises about "if", "when" or "in what order" each UDP message will arrive. If a UDP datagram doesn't make it to its destination, your networking software will not return an error.

[11] But not necessarily immediately, as it uses retransmission to achieve those guarantees.

[12] Most services that use UDP, such as DNS, implement their own retransmission strategy, so it's a bit more complicated than "yelling in the dark". :)

UDP datagrams usually contain a return address so, if a response is expected, the computer that receives the datagram knows where to send it.

UDP has very low overhead and is therefore very popular to use for services where success does not need to be guaranteed, like synchronizing your computer's clock or streaming audio on a Voice over IP (VoIP) call[13].

We talk about these two communication types because there are important limitations when dealing with UDP and network address translation, a topic we explore in section 7 on page 12.

# 5 Network Ports

As we learned a couple of sections ago, every computer on a network must have a unique IP address. It is kind of like a temporary serial number in that it identifies one specific computer on one specific network.

But there are many things to do on a network and they can all happen at the same time! We may know where a computer is, but how do we address the service that is on the computer that we want?

Think of a cable box. A cable box sits on top of your television and it is addressed by your cable company using its serial number (IP address). When you turn your TV on and flip through the channels on your cable box, you receive different television shows (services) from the channels (ports) that your cable company provides.

Networks are similar in that communications are handled on ports, which are like a channel. When you ask for a web page, you basically say, "Hey computer, I need to ask port 80 to give me your home page and send it back to me at IP address 208.40.80.2 on port 51,589", where 51,589 is any arbitrary port number that your computer has available. The server responds back with, "Hey, so you want to talk?" Your computer says, "Yup!". The connection is then established[14].

Some port number assignments are governed by The Powers That Be[15]. Port 80 is HTTP (a.k.a. the web), port 21 (both UDP and TCP) is FTP, mail is port 25—there are thousands of services that are available on a network. Some are famous and are always expected on a specific port. Others are arbitrarily assigned by their designer and may conflict with someone else's choice.

[13] If a chunk of audio doesn't arrive in the right order, the receiver can buffer it or drop it since there is no time to retransmit. In VoIP, low latency is much more important than complete accuracy.

[14] Nerds call this process *The Three Way Handshake* or *SYN-SYN/ACK-ACK*, which reminds this author of the movie *"Mars Attacks!"*

[15] The Internet Assigned Numbers Authority

There is a lot more to ports and how they work. For our purposes, it's enough to know that IP addresses are used to locate computers while ports are used to locate services on those computers and to facilitate multiple connections between different computers at the same time.

It is important to understand ports because the topic will come up when we explore NAT and firewalls in later sections.

# 6 Private and Public IP Addresses

There are hundreds of millions of computers and hundreds of thousands (if not millions) of networks in operation throughout the world. The geniuses that invented the Internet back in the 70's never imagined that everyone and their mom would be using it. The result is that there are not nearly enough IP addresses to accommodate the number of devices that are using the Internet.

To alleviate this problem, The Powers That Be[16] decided to reserve three blocks of addresses for private networks:

**10.0.0.1 - 10.255.255.255**

**172.16.0.1 - 172.31.255.255**

**192.168.0.1 - 192.168.255.255**

History dictates that not only are these private address blocks differently sized, they are also segmented differently. For example, the "`192.168.x.x`" block is most often represented as 256 different networks with 255 IP addresses each, with a subnet of "`255.255.255.0`". Since these addresses are private, there is nothing written in stone about how you segment your network. It's just what other networking types would expect to see when they look at your network configuration.

Someone that decides to use these IP addresses can do so without any coordination with an outside organization. That is because by their very definition, routers on the Internet will not view these as addresses that they can route. That is why they are called *non-routable* or *private* IP addresses.

Addresses that are outside the range of those listed above are considered *routeable* or *public*. If you arbitrarily choose a public IP address and then connect it to a router which is connected to the Internet, you will create problems and your network will not work correctly.

People creating networks in their own organization almost universally use private IP addresses. Back in the early days of the Internet, this was thought

[16] Again, The Internet Assigned Numbers Authority

of as a second-class solution. Now reality has set in, addresses are scarce, security concerns abound and we all have private IP addresses on our desktop computers[17].

The trouble with a private address is that you cannot communicate with anyone outside your network. If I try to go to "`www.trms.com`" and my computer uses a private IP address, the remote computer will not know where to respond to because the routers that are in between our two computers will not allow traffic to go back to a private IP address.

If this is true, then how is it that we all have private IP addresses and we are still able to communicate on the Internet? That question leads us to our next section...

# 7 Network Address Translation

Network Address Translation (NAT) is a magic technology that makes our Internet possible. Without it, there would be a lot fewer computers on the Internet and that would be a Bad Thing. But what is NAT?

NAT is a feature of a router[18] that enables computers inside your network that are using private IP addresses to communicate with computers on the Internet.

When you ask for something that is outside your network, your computer goes to the router. A router using NAT, in turn, completes the request *on your behalf*. The destination computer then establishes a connection with your router and your router is responsible for marshalling the packets back to you.

Your computer thinks that it is communicating directly with the destination computer and the destination computer thinks that it is talking to you from your router's address[19].

Because TCP communications are based on a connection and that connection is basically established by information that is stored in the packets of data that are a part of that communication, NAT works[20].

NAT does not work very well with UDP communication because there is no connection, and thus no handshake with which to base future communications upon. Fortunately, some newer routers are able to make educated guesses about

[17] For fun, open up a command prompt on your desktop computer. Type "`ipconfig`". Chances are that your address falls within one of the ranges listed above!

[18] You'll remember that a router routes IP traffic to remote computers. A router is at the edge of your network and links it to other networks.

[19] Incidently, the destination computer sees you as your router, but if they're aggressive enough, they can mine your private IP address out of the packet. That's how online poker games know if you're cheating! :)

[20] ...mostly. :) There are times when NAT fails even with TCP. Examples include certain kinds of FTP, MSN messenger file transfers and others.

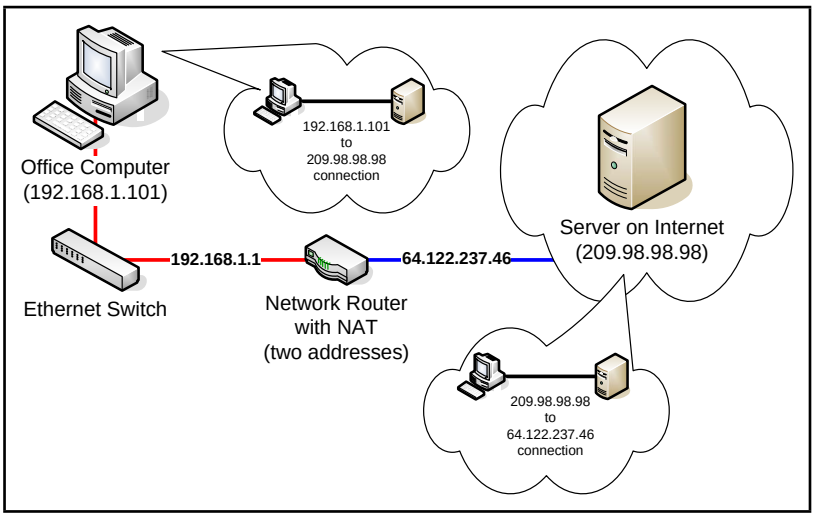


**Figure 2:** A router with NAT will use at least two addresses. In this example they are ("192.168.1.1") for the local address and ("64.122.237.46") for the public IP address. Local computers address it from the "192..." address. Computers on the Internet 'see' your computer as though it has the "64..." address. When you communicate with a computer outside of your network, your router uses NAT to make the connection on your behalf.

where incoming UDP datagrams should be sent[21] based on outgoing UDP datagrams that your computer recently sent. Educated guesses and occasional failures are an accepted part of UDP-over-NAT because the spirit of UDP is to be unreliable.

> Despite the modern UDP-enabling features of NAT equipped routers, there are many administrative and protocol headaches associated with NAT. Since NAT came after the design of the Internet, it wasn't a part of the underlying design of IP networking. As a result, it violates basic assumptions embedded into the design of the Internet to accomplish its amazing feats.
>
> It is for this reason that we all look forward to a day when a new version of IP networking, called IPv6, is able to finally retire the current one.

The key thing to remember about NAT is that there are two connections: you with the router and the router with the destination. The router makes a connection *on your behalf* and forwards all return packets back to your computer. For the most part, it works just like if your computer were directly connected to the destination computer.

# 8 Firewalls

A firewall is a device or a feature on a router that is able to block specific IP traffic based on a set of rules. Firewalls were traditionally installed on the edge of your network, but have become an important feature of operating systems and software firewalls are now found on many desktop computers.

To understand the significance of firewalls, we need to acknowledge a troubling fact: *all non-trivial computer systems have bugs.* It doesn't matter what platform you operate on or if you are up to date with your "Microsoft Patches". There are bugs in your system and many of them can be exploited to gain access to your network resources.

One of the central purposes of a firewall is to block access on your network so that bad people can exploit fewer bugs. :) Why expose a service to the Internet that might be a "way in" when you are not even using it? In fact, it doesn't even need to be a bug to allow access. What happens if you have a web server running on your computer that isn't configured properly? It is just waiting for someone to "configure it for you".

Firewall rules may be set for incoming traffic and outgoing traffic. Incoming traffic refers to traffic from outside the network coming into your network. Out-

[21] . . . using a feature called UDP NAT Traversal

going traffic refers to requests made from within your network to the outside. Incoming rules help protect your computer from attacks. Outgoing rules help protect everyone else from your computer should it fall victim to a virus or hacker.

The most common type of rule that a firewall will follow is a port rule. These rules simply block traffic, either incoming or outgoing, on a specified port. For example, your firewall may block all traffic on TCP port 23, which is commonly used for the telnet service.

Other rules might be based on traffic that comes from specific IP addresses. Your firewall may be configured to block all incoming traffic on port 23, except when it comes from an IP address that is from a remote location within your organization.

### 8.1 Dire Warning About Firewalls

Many people view a firewall as *the answer* to their security issues. They are not. In fact, firewalls are really just a small (but important) part of an overall security strategy.

Firewalls do not address the human elements of security, like 'phishing' scams and viruses that come through email attachments. Furthermore, they do not stop attacks on the services that you do expose on the Internet.

Also, there is something called the "chewy middle" of your network that can completely negate a firewall's effectiveness. This is where someone within your network shows some initiative by installing a $40 wireless access point and provides anyone within a reasonable proximity to your network full and unencrypted access to the inside.

## 9 Port Forwarding

Port forwarding is a feature of a router that is typically used in conjunction with NAT and firewalls. This feature forwards incoming traffic on specified ports to addresses that are inside the network. An example of port forwarding might be, "Forward all incoming traffic on TCP port 80 to the computer at "`192.168.1.3`".

Port forwarding is at the heart of many networks that provide services on the Internet. It is a fantastic way to partially shield a computer behind a firewall while allowing specific traffic through to services that are running on that computer.

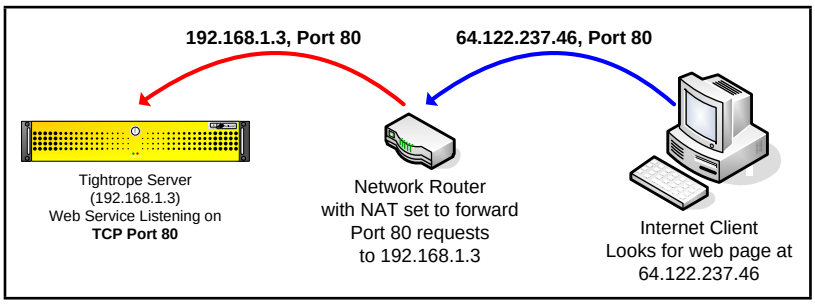


**Figure 3:** In this example, a computer on the Internet is looking for a web site that you are hosting behind a router configured with NAT and port forwarding. It looks for the site at your router's external IP address and the router forwards that request to your web server.

# 10 Virtual Private Network

A Virtual Private Network (VPN) is a method of extending a private network to computers that are located on a remote network through the use of a public networking infrastructure. When you use a VPN from home to get into your work's network, you are making a connection *on top of* your Internet connection. You are, in effect, dialing into your work's network over the Internet.

This new connection is an encrypted link that effectively joins your home computer to your work's network. This connection includes its own IP address, subnet mask, DNS address, etc. It is a new connection in every sense of the word. Your home computer now has all of the rights and privileges that you would normally have if you were on your computer at work. Typically, this means that you can print to a printer at work, see all of the network resources that are normally available, etc.

Incidentally, we say that VPN is a tunneling protocol, in that you tunnel through the Internet and into a remote network. VPN wraps all of your communications into encrypted packets and unwraps them at the "front door" of the network that you are accessing.

# 11 How Do I Access Cablecast or Carousel From Home?

This is the most common sales and support question that we get at Tightrope Media Systems. Because our products are web based, people assume that they will be able to access them from home. This assumption is correct, provided the network is configured to let it happen!

There are three ways to gain access to your Cablecast or Carousel system from the Internet, each with their own security and accessibility tradeoffs.

## 11.1 Option 1: Hang It Out On the Internet

If you give your computer a public IP address, people will be able to access your Tightrope server. Using this method, you will not only be able to update your system, but others will be able to access the public web interface of Cablecast and Carousel, giving them the ability to view messages and see your Cablecast system's schedule.

The down side is that your server is hanging out on the Internet, ready to be infected with the latest worm or controlled by the next bored teenager.

Because of the insecurity of this option and the relative ease of the alternatives, security professionals suggest that you seek alternatives to exposing your servers directly to the Internet.

## 11.2 Option 2: Use Port Forwarding

The most popular method of putting your Tightrope server on the Internet is to use port forwarding. Your server sits inside your network on a private IP address and traffic that comes in on a specific port is forwarded to your server.

The upside to this configuration is that any type of access that you need on your Tightrope server can be accommodated, including anonymous access to the system's public web interfaces.

On an existing network, this configuration might be difficult to accomplish because of security concerns and existing configurations. Specifically, if your network administrator forwards a port to your Tightrope server, there is still no guarantee that the server will not become infected or compromised. Once it is compromised, the rest of the computers on your network are in jeopardy.

The most effective way past this obstacle is to place your Tightrope server inside your network's DMZ, which actually does stand for Demilitarized Zone. A DMZ is a tiny network with a firewall on either side. One firewall separates the DMZ from the Internet and the other separates the DMZ from the internal network. That way, if your Tightrope server is infected, only the computers within the DMZ are at risk.

The second obstacle is that your network may be using the port that you need. There are three ways around this problem:

1. Use a new public IP address for your Tightrope servers. This will eliminate the port conflict because your server will be the only computer using that port on that address.

2. Use a firewall that supports named forwarding and add a DNS entry for your Tightrope system's web server. That way, if someone is looking for "`carousel.mydomain.org`" the firewall/NAT server will forward them to the Carousel machine instead of your main web server.

3. Use a different port on your Tightrope server by changing it in Internet Information Server. Instead of using port "`80`" for web access, you could use port "`8080`". See section 12.2 on page 20 for information about how to do this.

## 11.3 Option 3: Use VPN

This is a limiting option because all of the public web features of Tightrope's system will be unusable, given that there is no anonymous access from the Internet to your Tightrope server. It is secure, however, because you are simply using the same VPN access that you would use for your regular network access. If your organization already uses VPN, you don't even need to involve your IT department.

## 11.4 The "Forget the IT Department" Option

Instead of fighting your network administrators, you may be able to buy an inexpensive connection of your own.

DSL or cable modem connection might cost you only 30 to 80 dollars per month and provide a method to access your computer from outside your network.

Using this method, you would access your system from your desk at work by going out on the Internet through your regular network and back into your building through the new connection that you purchased for your Tightrope System. You access your Tightrope servers as you would any other computer on the Internet.

The biggest limitation is that you cannot connect any machine within your building's network to any Tightrope server. For Carousel machines, this is not a significant limitation unless you are uploading large video files. For Cablecast installations with video servers, this is a problem because you will not be able to transfer video files into the server using Windows Networking.

Another consideration has to do with the type of Internet connection that you purchase. It is very common for cable companies to block certain ports on your network. Also, obtaining a static IP address is often difficult or expensive to accomplish.

For ways around these limitations, see the next section...

# 12 Avoiding The Tyranny of Cable Modem Providers

If you are on your own or find it impossible to put your Tightrope system on the Internet through your IT department, getting your own connection might be the best option.

When an organization wants to put out a web presence, they will purchase a connection that is designed for the task. Traditionally this has meant a T1 connection, which is very expensive.

A great alternative to a T1 is a DSL or cable modem connection, which can be had for less than 40 dollars per month, in some cases. Unfortunately, these connections are designed for consumers and as such lack static IP addresses. In some cases the Internet Service Provider (ISP) will even block incoming access to common ports, like TCP port 80, in an attempt to stop you from hosting web sites with your connection.

Sometimes your provider will have a business version of their services which will provide you with everything that you need to host a web site. If your ISP does not offer business class service or it is cost prohibitive, you have one last option.


Many Tightrope customers have a working relationship with their cable provider. The policies in place might be more applicable to high traffic sites than it is to your situation and you may be in a position to ask for some flexibility.

**BEFORE YOU CONTINUE READING THIS SECTION:**
Call your cable provider and ask them if the following steps are acceptable to them. You do not want to circumvent their policies only to suffer their wrath when you find they've canceled your account and you are facing possible legal action. Get any negotiated exceptions to their policies *in writing!*

***Tightrope Media Systems does not condone nor advocate wanton violation of your ISP's acceptable use policies!***

## 12.1 Dynamic DNS

The first order of business is figuring out how to find your Tightrope server from outside your network. To do this, you use a technology called Dynamic Domain Name Service (DDNS). DDNS providers offer the same services as DNS

except that they are able to track your dynamic IP address using special software that you install on your Tightrope server. This service is generally very inexpensive.

The down side to this service is that if your IP address changes, your site may be down for a short period of time. This is because the software might not discover the switch right away.

To find a dynamic DNS service provider, simply 'Google' the term *dynamic DNS*. You'll find many from which to choose. :)

### 12.2 Change Your Port Number

Your ISP may choose to block incoming data packets on specific ports. You can get around this by changing the port on which your Tightrope system's web server is listening.

1. Right-click on "**My Computer**" on your Tightrope server's desktop.
2. Select "**Manage**"
3. Expand the "**Services and Applications**" Branch. Expand the "**Internet Information Services**" and "**Web Sites**" branches.
4. Right-click on "**Default Web Site**" and click properties. (figure 4 on the next page)
5. Under the "**Web Site**" tab, find the "**TCP Port**" field. (figure 5 on page 22)
6. Enter a port that is higher than 50,000. These are called *unregistered* or *user ports* and are unlikely to conflict with another application on your computer.
7. When you access your Tightrope server, you will have to designate the port number by entering a ":" then the number after the address. Example: "`carousel.mydomain.org:8080`" or "`192.168.1.3:8080`".

## 13 Time Synchronization, UDP and NAT

We spent so much time talking about NAT, UDP and TCP in this guide because there are situations where you will want to use UDP through your router and will have problems doing so. The most common situation is where you want to synchronize a Tightrope server's clock to "Internet Time" using a service called Network Time Protocol (NTP).

NTP uses UDP port 123. Some organizations will have an NTP server running within their network and configuring a Tightrope server to use it is trivial.

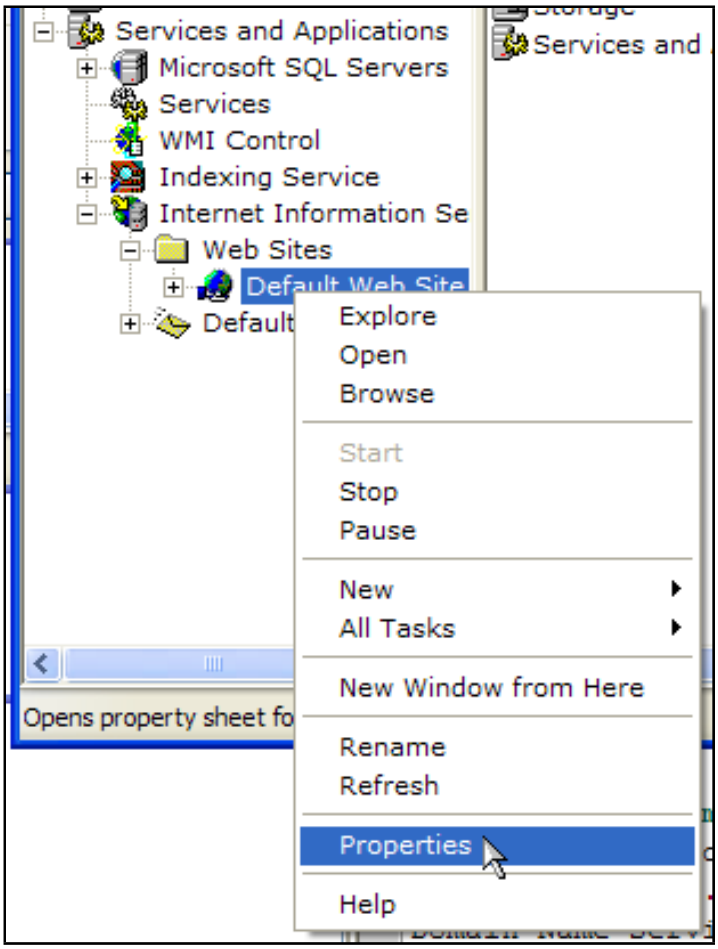


**Figure 4:** Navigating to Internet Information Services

If there is no NTP service available, you will have to configure your router to forward all NTP traffic to your Tightrope server, or enable UDP-NAT Traversal, which is able to guess at the destination of incoming UDP packets based on recent outgoing traffic.

On inexpensive consumer cable modem routers, simple versions of UDP-NAT Traversal are often enabled by default. More expensive router/firewall/NAT combinations require some configuration.

# 14 Summary

In this paper we were able to cover quite a bit of ground. You should now have a basic grasp on the following concepts:

- IP addresses, subnet masks, routers and DNS
- DHCP and Dynamic IP addresses
- The difference between TCP and UDP and why it matters
- A basic understanding of network ports
- Understand private and public IP addresses

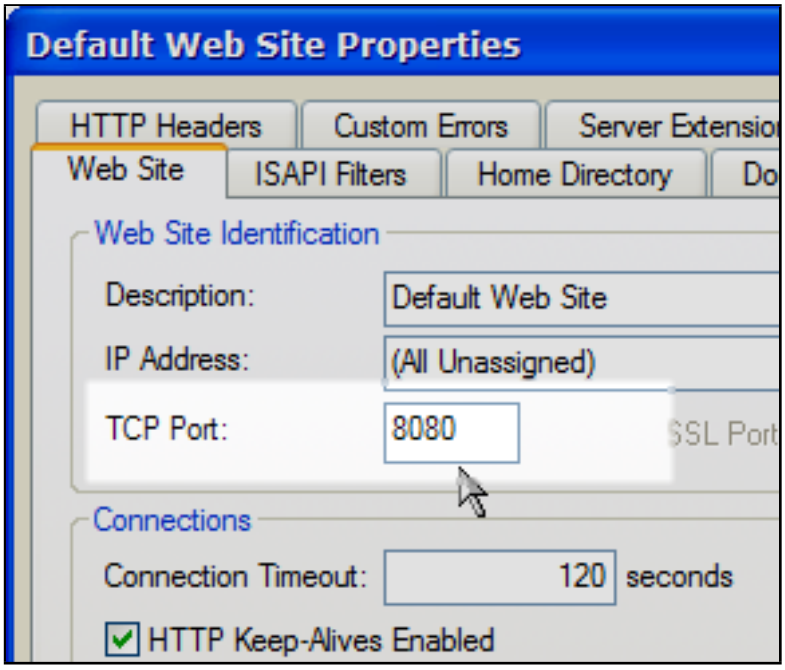


**Figure 5:** Changing your port number.

- Network Address Translation
- Firewalls and port forwarding
- Virtual Private Networking
- The various options for getting your system on the Internet