



**H2020-MSCA-ITN-2016**

**Marie Skłodowska-Curie actions (MSCA)**

**Innovative Training Networks (ITN)**

**Grant Agreement Number 722482**

**ESSENTIAL**

***Evolving Security ScienceE through Networked  
Technologies, Information policy And Law***

**EJD**

**WP 5: Training Programme**

**LEAD PARTNER – Edith Cowan University**

**D5.7 - Fellows' Handbook for Network Architecture**



Project co-funded by the European Commission within Marie Skłodowska-Curie actions (MSCA) Innovative Training Networks (ITN)		
Dissemination Level:		
PU	Public	X
CO	Confidential, only for members of the consortium (including the Commission Services)	
EU-RES	Classified Information: RESTREINT UE (Commission Decision 2005/444/EC)	
EU-CON	Classified Information: CONFIDENTIEL UE (Commission Decision 2005/444/EC)	
EU-SEC	Classified Information: SECRET UE (Commission Decision 2005/444/EC)	

Document Version Control:		
Version 0.1	Prepared by: Craig Valli	20/March/2019



## Table of contents

<b>1. INTRODUCTION</b>	<b>4</b>
1.1. AIM OF THE HANDBOOK	4
1.2.	4
<b>2. SOME MATHS FIRST</b>	<b>5</b>
<b>3. SOME PHYSICS FIRST</b>	<b>6</b>
<b>4. THE OSI MODEL</b>	<b>8</b>
4.1. LAYER 1 - PHYSICAL LAYER	10
4.2. LAYER 2 - DATA LINK LAYER	11
4.3. LAYER 3 - NETWORK	12
4.4. LAYER 4 - TRANSPORT LAYER	14
4.5. LAYER 5 - SESSION LAYER	15
4.6. LAYER 6 - PRESENTATION LAYER	15
4.7. LAYER 7 - APPLICATION LAYER	16
<b>5. TOPOLOGY</b>	<b>17</b>
5.1. STAR	17
5.2. BUS	18
5.3. TOKEN RING	18
<b>6. TCP/IP</b>	<b>20</b>
<b>7. ANNEX</b>	<b>21</b>



## **1. Introduction**

### **1.1. Aim of the Handbook**

The aim of this handbook is to provide an easy introduction to Network Architecture primarily to the ESRs with a limited technical background. The aim is to assist the ESSENTIAL Early Stage Researchers and other students in the field of security science understand the basics of Network Architecture and to introduce them to key texts to assist their learning. This Handbook is regularly updated on the Vergaderen platform (that is, the internal communication platform used within ESSENTIAL).

### **1.2.**

The main target group of this Handbook are the ESSENTIAL Early Stage Researchers, but also any other students from the field of Security Science with a limited technical background.



## 2. Some maths first

### **Nano, Milli, Kilo, Mega, Giga, Tera or Metric Prefixes**

Nano and Milli are used for seconds when measuring intervals of packets on a network today.

Kilo, Mega, Giga and Tera are used for measuring bandwidth in bits transmitted per second.

### **Further Reading**

*[https://en.wikipedia.org/wiki/Metric\\_prefix](https://en.wikipedia.org/wiki/Metric_prefix)*

### **Binary numbers**

You should be familiar with how numbers are represented in binary. The Wikipedia entry for this is a good starting reference *[https://en.wikipedia.org/wiki/Binary\\_number](https://en.wikipedia.org/wiki/Binary_number)*

### **Hexadecimal numbers**

You should be familiar with how numbers are represented in hexadecimal. The Wikipedia entry for this is a good starting reference *<https://en.wikipedia.org/wiki/Hexadecimal>*

### **Exercise**

*Write your date of birth as a binary number and as a hexadecimal number mine is*

*100/10/11110101101 and #04/#02/#07AD*



### 3. Some physics first

#### Attenuation

Attenuation is a term that refers to any reduction in the strength of a signal. Attenuation occurs with any type of signal, whether digital or analog. Sometimes called "loss", attenuation is a natural consequence of signal transmission over long distances. The extent of attenuation is usually expressed in units called decibels (dBs). Attenuation can also be expressed in terms of voltage.

In conventional and fiber optic cables, attenuation is specified in terms of the number of decibels per foot, 1,000 feet, kilometer, or mile. The less the attenuation per unit distance, the more efficient the cable or media is. When it is necessary to transmit signals over long distances via cable, one or more repeaters can be inserted along the length of the cable. The repeaters boost the signal strength to overcome attenuation. This greatly increases the maximum attainable range of a communication.

*(ref <https://searchnetworking.techtarget.com/definition/attenuation>)*

#### Bandwidth

Is the ability for a media or medium to sustain transmission in a communication. This is typically measured in bits per second. There are physical limits that will limit the maximum possible bandwidth of transmission of data. The best is optical fibre. The following is a table of commonly used media and protocols today.

56 kbit/s	Modem / Dialup	600 Mbit/s	Wireless <a href="#">802.11n</a>
1.5 Mbit/s	<a href="#">ADSL Lite</a>	1 Gbit/s	<a href="#">Gigabit Ethernet</a>
1.544 Mbit/s	<a href="#">T1/DS1</a>	1.3 Gbit/s	Wireless <a href="#">802.11ac</a>
4 Mbit/s	ADSL1	5 Gbit/s	<a href="#">USB 3.0</a>
10 Mbit/s	<a href="#">Ethernet</a>	7 Gbit/s	Wireless <a href="#">802.11ad</a>
11 Mbit/s	Wireless <a href="#">802.11b</a>	10 Gbit/s	<a href="#">10 Gigabit Ethernet</a> , <a href="#">USB 3.1</a>
24 Mbit/s	ADSL2+	40 Gbit/s	<a href="#">Thunderbolt 3</a>
54 Mbit/s	Wireless <a href="#">802.11g</a>	100 Gbit/s	<a href="#">100 Gigabit Ethernet</a>
100 Mbit/s	<a href="#">Fast Ethernet</a>	1000 Gbit/s	1 Terabit Ethernet

The links are live Wikipedia links should you wish to start to drill some more. Some other points traditional Ethernet (wired) achieves "saturation" i.e where an unacceptable rate of

*"This project has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 722482".*



errors at around 40-60% of "bandwidth" so a 100Mbit pipe will at maximum push 60 Mbits per second or a 6 Megabyte file per second.

**Hints:** Many internet service providers have a set amount of bandwidth or access/routing to the next hop i.e the Internet backbones. *Caveat emptor* at high usage times you are in contention with ALL users of that particular network for your "share" of the bandwidth which depending on the country and the media used can see your 100Mbit connection barely be able to send 2Mbit during peak. Many ISPs will advertise their speeds that are achieved in Off-Peak times, always ask about throughput in Peak.

### Exercise

For your particular country investigate and price bandwidth for the following scenarios

- 1) Home user who works from home downloads and does remote backups ~ 1Gigabyte a day
- 2) A small geology practice who wants to upload to a maximum of 30 Gigabytes of files a day, 365 days a year to their customer

### Further Reading

Comer, Douglas (2008) Computer Networks and Internets, page 99 ff, Prentice Hall 2008.

Forouzan, Behrouz A. (2017) *Data communications and networking*, McGraw-Hill, 2007

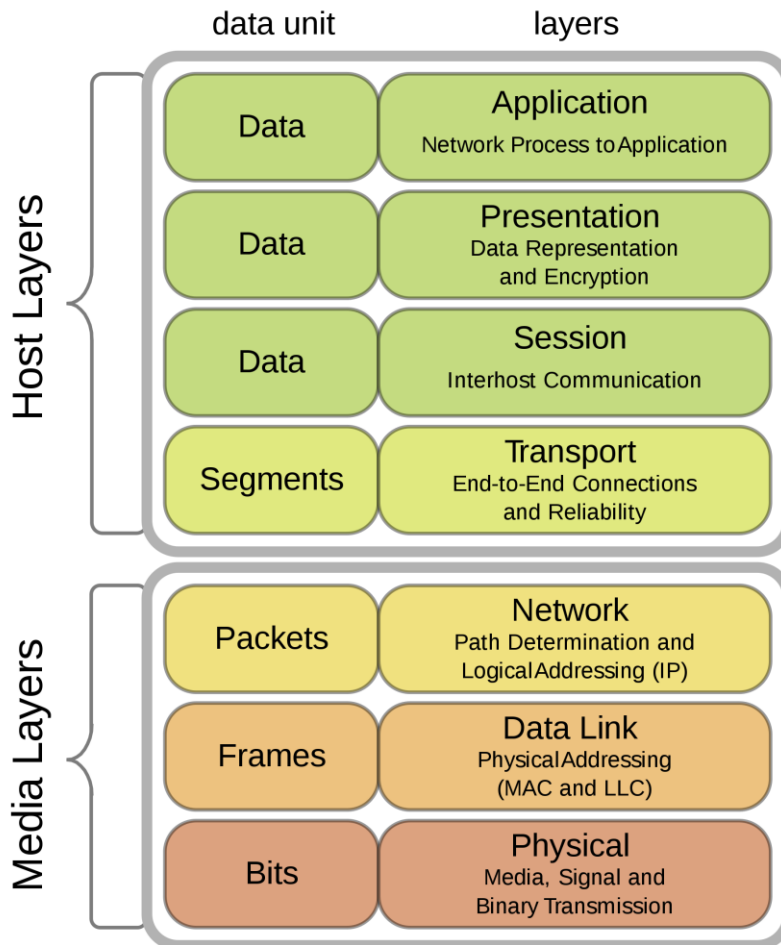
Cisco Press(2003) [Cisco Networking Academy Program: CCNA 1 and 2 companion guide, Volume 1–2](#)



## 4. The OSI Model

The model defines and codifies the concept of layered network architecture. It is a way of sub-dividing a communications system further into smaller parts called layers. A layer is a collection of similar functions that provide services to the layer above it and receives services from the layer below it. On each layer, an instance provides services to the instances at the layer above and requests service from the layer below

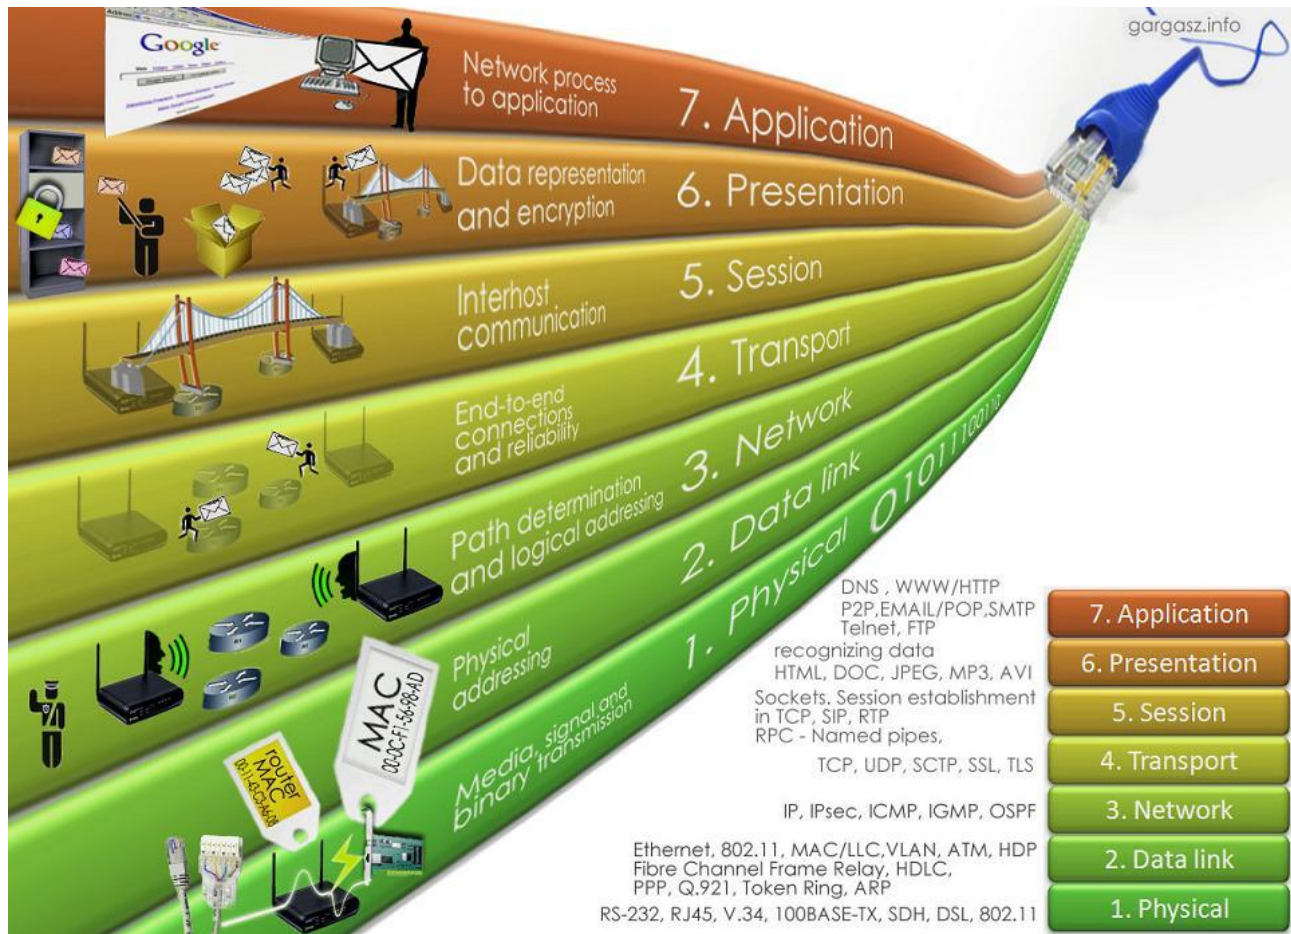
There are seven layers to the OSI model these are



*"This project has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 722482".*







*"This project has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 722482".*




## 4.1. Layer 1 - Physical Layer

### Wired Cables

This is the layer that is the “media” by which your messages flow. In the case of “wires” it converts signals into electrical pulses. Most common today is UTP (typically a blue cable with an RJ-45 connector on both ends). This supports speeds from 10Mbit to 10Gbit. The cables are male on both ends that connect to a female socket, normally wall mounted or connect into a physical Switch or Hub.



messages  
electrical  
cable with  
from  
connect to a



There are other physical layer such as RS232 which is relatively old but still has some uses today. There is also an updated 9 pin design called RS485. They have a male and a female plug. They can at best transmit at 3mbits which by todays standards is slow.

The physical wires do suffer from Electro Magnetic Interference (EMI) or EMR (R = Radiation). This can significantly disrupt communication.

### Exercise

There was in wide use initially coaxial cable for what was 10Base2 networking which has a top speed of 10mbits. Use Google image search to find some images of these cables. Also find out the function of a “terminator” in this sort of network. The maximum distance on this type of networking was 180m with no more than 30 devices.

### Fibre Optic Cables

The other is fibre optic cables these convert computer signals into light pulses. One of the main advantages is that it does not suffer from EMI. It can send signals significant distances as well and are used typically for high speed links, inter-country and inter-continental links for instance.



## 4.2. Layer 2 - Data Link Layer

Media access Control (MAC) and Logical Link Layer (LLC).

Each device at this layer has a 6 digit hexadecimal number called a MAC address the first 3 are fixed and regulated by FCC. The first three digits uniquely identify what type of device it is and who made it. This then also links to public technical documentation about the device if available. There are publically accessible resources that relate to these and the FCC has a search engine at <https://www.fcc.gov/oet/ea/fccid> The remaining three digits relate to making the device unique to allow network addressing.

Wireless networks (WIFI, 802.11) this is the physical air, with wired or fibre networks its contained in the transmission media. Hence why wireless is very hard to secure against certain cyber attacks and also is less reliable.

### Exercise

Determine MAC addresses of your network interfaces and list the manufacturer of these components. You determine your MAC address by using `ifconfig -a` (Linux) or `ipconfig /all` (Windows). Then having retrieved your MAC address use the first 3 Hexadecimal numbers and check what manufacturer made your network card.



### 4.3. Layer 3 - Network

Layer 3, the network layer, is the layer where routing takes place. The protocols and technologies in Layer 3 enable network-to-network communications. A core technology in Layer 3 is a router that routes packets between two networks. A Layer 3 switch is simply a Layer 2 device that also does routing (a Layer 3 function). Another key aspect of routers is that each interface on a router has its own IP address, because each of those interfaces is on a different physical networks.

More technically Layer 3 provides the functional and procedural means of transferring variable length data sequences from a source host on one network to a destination host on a different network, while maintaining the quality of service (QoS) requested by the transport layer (in contrast to the data link layer which connects hosts within the same network). The network layer performs network routing functions, and might also perform fragmentation and reassembly of transmissions, and report delivery errors.

Routers operate at this layer—sending data throughout the extended network and making the Internet possible. This is a logical non-hierarchical addressing scheme whose values are chosen by the network engineer. A number of layer-management protocols belong to the network layer, including routing protocols, multicast group management, network-layer information and error, and network-layer address assignment. The network layer can have at least three sub- layers: 1) Sub-network access that considers protocols that deal with the interface to networks, such as X.25; 2) Sub-network dependent convergence, used when it is necessary to bring the level of a transit network up to the level of networks on either side; and 3) Sub-network independent convergence, which handles transfer across multiple networks, the best example of which is CLNP, or IPv7 ISO 8473. (It manages the connectionless transfer of data one hop at a time, from end system to ingress router, router to router, and from egress router to destination end system. It is not responsible for reliable delivery to a next hop, but only for the detection of erroneous packets so they may be discarded. In this scheme, IPv4 and IPv6 would have to be classed with X.25 as subnet access protocols because they carry interface addresses rather than node addresses

So much of what a network administrator does is dealing with IP addresses and subnetting, routing protocols, firewall rules and [Access Control Lists \(ACLs\)](#), and many types of [Quality of Service \(QoS\)](#) that is all enabled by Layer 3 technologies.

Functions of the network layer include:

*"This project has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 722482".*



### Connectionless communication

For example, [IP](#) is connectionless, in that a data packet can travel from a sender to a recipient without the recipient having to send an acknowledgement. Connection-oriented protocols exist at other, higher layers of the OSI model.

### Host addressing

Every host in the network must have a unique address that determines where it is. This address is normally assigned from a hierarchical system. For example, you can be :

"Fred Murphy" to people in your house,

"Fred Murphy, 1 Main Street" to Dubliners,

"Fred Murphy, 1 Main Street, Dublin" to people in Ireland,

"Fred Murphy, 1 Main Street, Dublin, Ireland" to people anywhere in the world.

On the Internet, addresses are known as [IP addresses](#) (Internet Protocol).

### Message forwarding

Since many networks are partitioned into subnetworks and connect to other networks for wide-area communications, networks use specialized hosts, called gateways or [routers](#), to forward packets between networks.



## 4.4. Layer 4 - Transport Layer

The [transport layer](#) provides the functional and procedural means of transferring variable-length data sequences from a source to a destination host, while maintaining the quality of service functions.

The transport layer controls the reliability of a given link through flow control, [segmentation/desegmentation](#), and error control. Some protocols are state- and connection-oriented. This means that the transport layer can keep track of the segments and re-transmit those that fail delivery. The transport layer also provides the acknowledgement of the successful data transmission and sends the next data if no errors occurred. The transport layer creates segments out of the message received from the application layer. Segmentation is the process of dividing a long message into smaller messages.

OSI defines five classes of connection-mode transport protocols ranging from class 0 (which is also known as TP0 and provides the fewest features) to class 4 (TP4, designed for less reliable networks, similar to the Internet). Class 0 contains no error recovery, and was designed for use on network layers that provide error-free connections. Class 4 is closest to TCP, although TCP contains functions, such as the graceful close, which OSI assigns to the session layer. Also, all OSI TP connection-mode [protocol](#) classes provide expedited data and preservation of record boundaries. Detailed characteristics of TP0-4 classes are shown in the following table:[\[7\]](#)

Feature name	TP0	TP1	TP2	TP3	TP4
Connection-oriented network	Yes	Yes	Yes	Yes	Yes
Connectionless network	No	No	No	No	Yes
Concatenation and separation	No	Yes	Yes	Yes	Yes
Segmentation and reassembly	Yes	Yes	Yes	Yes	Yes
Error recovery	No	Yes	Yes	Yes	Yes
Reinitiate connection <sup>a</sup>	No	Yes	No	Yes	No
Multiplexing / demultiplexing over single <a href="#">virtual circuit</a>	No	No	Yes	Yes	Yes
Explicit flow control	No	No	Yes	Yes	Yes
Retransmission on timeout	No	No	No	No	Yes
Reliable transport service	No	Yes	No	Yes	Yes

<sup>a</sup> If an excessive number of [PDUs](#) are unacknowledged.



An easy way to visualize the transport layer is to compare it with a post office, which deals with the dispatch and classification of mail and parcels sent. A post office inspects only the outer envelope of mail to determine its delivery. Higher layers may have the equivalent of double envelopes, such as cryptographic presentation services that can be read by the addressee only. Roughly speaking, [tunneling protocols](#) operate at the transport layer, such as carrying non-IP protocols such as [IBM's SNA](#) or [Novell's IPX](#) over an IP network, or end-to-end encryption with [IPsec](#). While [Generic Routing Encapsulation](#) (GRE) might seem to be a network-layer protocol, if the encapsulation of the payload takes place only at the endpoint, GRE becomes closer to a transport protocol that uses IP headers but contains complete Layer 2 frames or Layer 3 packets to deliver to the endpoint. [L2TP](#) carries [PPP](#) frames inside transport segments.

Although not developed under the OSI Reference Model and not strictly conforming to the OSI definition of the transport layer, the [Transmission Control Protocol](#) (TCP) and the [User Datagram Protocol](#) (UDP) of the Internet Protocol Suite are commonly categorized as layer-4 protocols within OSI.

#### 4.5. Layer 5 - Session Layer

The [session layer](#) controls the dialogues (connections) between computers. It establishes, manages and terminates the connections between the local and remote application. It provides for [full-duplex](#), [half-duplex](#), or [simplex](#) operation, and establishes procedures for checkpointing, suspending, restarting, and terminating a session. In the OSI model, this layer is responsible for gracefully closing a session, which is handled in the [Transmission Control Protocol](#) at the transport layer in the Internet Protocol Suite. This layer is also responsible for session checkpointing and recovery, which is not usually used in the Internet Protocol Suite. The session layer is commonly implemented explicitly in application environments that use [remote procedure calls](#).

#### 4.6. Layer 6 - Presentation Layer

The [presentation layer](#) establishes context between application-layer entities, in which the application-layer entities may use different syntax and semantics if the presentation service provides a mapping between them. If a mapping is available, presentation protocol data units are encapsulated into session protocol data units and passed down the [protocol stack](#).



This layer provides independence from data representation by translating between application and network formats. The presentation layer transforms data into the form that the application accepts. This layer formats data to be sent across a network. It is sometimes called the syntax layer.<sup>[8]</sup> The presentation layer can include compression functions.<sup>[9]</sup> The Presentation Layer negotiates the Transfer Syntax.

The original presentation structure used the [Basic Encoding Rules](#) of [Abstract Syntax Notation One](#) (ASN.1), with capabilities such as converting an [EBCDIC](#)-coded text [file](#) to an [ASCII](#)-coded file, or [serialization](#) of [objects](#) and other [data structures](#) from and to [XML](#). ASN.1 effectively makes an application protocol invariant with respect to syntax.

## 4.7. Layer 7 - Application Layer

The [application layer](#) is the OSI layer closest to the end user, which means both the OSI application layer and the user interact directly with the software application. This layer interacts with software applications that implement a communicating component. Such application programs fall outside the scope of the OSI model. Application-layer functions typically include identifying communication partners, determining resource availability, and synchronizing communication. When identifying communication partners, the application layer determines the identity and availability of communication partners for an application with data to transmit. The most important distinction in the application layer is the distinction between the application-entity and the application. For example, a reservation website might have two application-entities: one using HTTP to communicate with its users, and one for a remote database protocol to record reservations. Neither of these protocols have anything to do with reservations. That logic is in the application itself. The application layer per se has no means to determine the availability of resources in the network.


---





## 5. Topology

### 5.1. Star




The star network topology is the most commonly used form today due to the proliferation of Ethernet networks utilising UTP cables. Large networks are a series of "stars" connected together. The star topology has many advantages for use including low cost, separation of failure, easy to expand and administer are just a few of the advantages. They are also relatively easy to expand limited only often by the media's ability to broadcast the transmission.

## 5.2. Bus



This was one of the first network types used. It is still in significant use today as a network topology normally used for local networks e.g automotive and aerospace use CanBUS networks to supply control systems networks within the platforms. One of its core problems is that if the "bus" is interrupted or not closed it will not function at all.

## 5.3. Token Ring



Token Ring is still used environments where error-free transfer/transmission is required or where timing is highly sensitive. Hence why you see Token-Ring still being implemented in many industrial based process systems. It is relatively expensive to deploy due to higher hardware input costs per seat

and is a problematic to troubleshoot due to its complexity, this sees its adoption not widespread, but it should be noted is still used for critical applications.

### **Further Reading**

[https://en.wikipedia.org/wiki/Token\\_ring](https://en.wikipedia.org/wiki/Token_ring)

*Castelli, Matthew (2002). Network Consultants Handbook. Cisco Press. ISBN 978-1-58705-039-8.*

*Gallo, M., & Hancock, W. M. (2001). Networking Explained. Digital Press.*

(Gallo & Hancock, 2001)



## 6. TCP/IP

There are several core texts to read and also get as your personal reference collection if you are to be working in the areas of networking. These books are technical but are basically the defacto bibles of TCP/IP if you are a non-IT graduate recommend you acquire, read and keep Kozierok (2005) first.

Comer, D. E. (2013). *Internetworking with TCP/IP Volume One* (6th ed.). Pearson.

Pyles, J., Carrell, J. L., & Tittel, E. (2016). *Guide to TCP/IP: IPv6 and IPv4* (6th ed.). Cengage Learning.

Less technical and more accessible are the following:

Kozierok, Charles M. (2005) *The TCP/IP Guide: A Comprehensive, Illustrated Internet Protocols Reference*, 1st Edition, No Starch Press

Leiden, Candace & Wilensk, Marshall (2009) *TCP / IP For Dummies* 6th Edition

(Pyles, Carrell, & Tittel, 2016)(Comer, 2013)

Comer, D. E. (2013). *Internetworking with TCP/IP Volume One* (6th ed.). Pearson.

Gallo, M., & Hancock, W. M. (2001). *Networking Explained*. Digital Press.

Pyles, J., Carrell, J. L., & Tittel, E. (2016). *Guide to TCP/IP: IPv6 and IPv4* (6th ed.). Cengage Learning.



## 7. Annex

Now proceed to read The No Sweat Guide to Network Topology from Auvik and other papers made available to you on Vergaderen



The No-Sweat Guide to

# Network Topology



KEVIN DOOLEY



## ABOUT KEVIN DOOLEY

Kevin has more than 15 years of experience as a network engineer. He has designed and implemented several of the largest and most sophisticated enterprise data networks in Canada and written several highly regarded books on networking for O'Reilly and Associates, including *Designing Large-Scale LANs* and *Cisco IOS Cookbook*. Kevin holds a Ph.D. in theoretical physics and numerous industry certifications.

## ABOUT AUVIK NETWORKS

Auvik is a cloud-based software-as-a-service (SaaS) application that makes it dramatically easier for small and mid-sized businesses to build and manage their IT networks. With Auvik, dozens of normally complex and time-consuming tasks are reduced to simple, automated steps. And it's all done through a beautiful and easy-to-use interface.

Copyright © 2015 Auvik Networks Inc. Some rights reserved.

You may give away or pass along this e-book for free as long as it is left completely intact, unaltered, and is delivered via this PDF file.

# Contents

---

- WHY TOPOLOGY MATTERS** . . . . . 4
  - What you'll learn in this guide . . . . . 5
- BASIC NETWORK TOPOLOGIES** . . . . . 6
  - Network layers . . . . . 6
  - Ethernet, bus and switched . . . . . 7
- NETWORK SEGMENTATION** . . . . . 9
  - VLANs and trunks. . . . . 11
  - Routing . . . . . 12
- HIGH AVAILABILITY** . . . . . 14
  - Device-level high availability. . . . . 15
  - Layer 2 network high availability . . . . . 15
  - Layer 3 network high availability . . . . . 17
- HOW TO DRAW NETWORK DIAGRAMS** . . . . . 18
  - Common network diagram symbols . . . . . 18
  - Drawing Layer 3 diagrams . . . . . 20
  - Drawing Layer 2 diagrams . . . . . 21
  - Drawing Layer 1 diagrams . . . . . 22
  - Drawing combined-layer diagrams . . . . . 22
  - How to map an existing network . . . . . 24
- USEFUL TOPOLOGIES** . . . . . 29
  - Single switch . . . . . 29
  - Single switch with Internet connection . . . . . 31
  - Touchdown segment . . . . . 33
  - Three-tier . . . . . 35
  - Web-tier . . . . . 37
  - Hub-and-spoke . . . . . 42
  - Ring . . . . . 44
  - Internal firewalls . . . . . 47
- HELPFUL RESOURCES** . . . . . 49



# Why Topology Matters

---

As a network administrator responsible for the care and feeding of a network, it's vitally important you have an extremely detailed understanding of your network topology.

Without this information, even basic troubleshooting can be unnecessarily difficult. You'll find, if you haven't already, that troubleshooting is *much* easier if you have detailed and up-to-date network documentation.

- When one device is having trouble communicating with another one, you need to know how the packets are supposed to get there.
- You need to know whether there are firewalls in the middle, or whether there are particular bottlenecks where you might be suffering from congestion problems.
- You need to understand any single points of failure where a simple device or link failure could have catastrophic effects.

There are several other good reasons to maintain your network documentation as well.

If you ever need to change anything, to add a new switch or a new link to a remote office, you need good documentation to understand how the traffic will flow through these new pieces of network. This will help ensure the new network is stable and efficient.

Another important reason to document your network is that you might need to show it to an auditor. PCI auditors always want to see your network topology diagrams. They need to know where the firewalls and servers are and all the different ways that somebody could get into the network.

Even if you never need to submit to this kind of external audit, if there were ever a serious problem with your network and you needed to either defend the design to management or bring in external consultants, having good and current diagrams will always help your case.

There are two ways to obtain a deep understanding of your network.

The first is to be some sort of super genius who can keep it all in her head at once. The trouble with this approach, though, is that it makes you so completely indispensable that you can never be promoted, which would be a shame for somebody so brilliant.

The second way, which has always worked better for me, mostly because I'm not even an intermediate genius, is to document everything.

## WHAT YOU'LL LEARN IN THIS GUIDE

*The No Sweat Guide to Network Topology* covers two main topics:

- How to document your network topology
- Some of the most common and useful topologies that you'll run into

By the time you get to the final pages, you'll be able to map your network both logically and physically. You'll understand the differences between the different logical and physical layers and how they relate to one another. And you'll know several of the most common network topologies, along with where and when to use them.

### A note on mapping methods

You can map your network manually using pencil and paper, or with a drawing program like Visio. You might also use automated software tools to help you. It's obviously a lot easier with the software tools, but even the best software needs a little bit of help.

For example, it's very hard for any software to know how the cables have been run or where the patch panels are. Yet these are often the technical details you need when you're troubleshooting a problem or connecting a new device. On the other hand, automated tools are extremely good at sorting out Layer 2 and 3 connections between devices.

The process of documenting a network is highly personal. I'll show you how I like to do it, and I'll try to explain why I think my methods are useful, but they're certainly not the only way.

Ready? Let's get started.

# Basic Network Topologies

---

This guide uses the concept of network protocol layers in every section, so it's worth taking a moment to explain what they mean.

## NETWORK LAYERS

There's a standard seven-layer model called the OSI model that every networking book mentions. The idea is that higher-layer protocols sit on top of lower-layer protocols. But reality is both more complicated and simpler than this model suggests.

Real networks use things like VPN tunnels where a Layer 3 network (or sometimes a Layer 2 network) sits on top of another Layer 3 network with a different logical topology. So although it's important to understand the layers, it's also important not to be too pedantic about them.

When talking about network topology, we're mostly interested in the bottom few layers.

**Layer 1 is the physical layer.** For network design purposes, this means the things you can touch: the cables and the equipment.

But it technically also includes the electrical and optical signalling properties. Layer 1 defines the properties of the cables that are necessary to carry the signals. Layer 1 also worries about wireless signalling where that's used.

As network designers, we need to make sure we get the right cables and we need to watch out for distance limitations. And of course we need to make sure the right devices are physically connected to one another.

**Layer 2 is the logical link layer.** Most modern networks are based around some sort of Ethernet. We have everything from 10Mbps to 100Gbps flavors of Ethernet. There are LAN and WAN Ethernets. And even though there are also a lot of other WAN protocols, like DSL

and T1, most WAN providers will prefer to deliver it to you using an Ethernet handoff. So those other Layer 2 protocols usually aren't your problem.

The other important Layer 2 protocol to consider is Wi-Fi. I think of Wi-Fi as simply a flavor of Ethernet because it uses very similar signalling. Obviously there are some very important radio frequency (RF) propagation and isolation factors to worry about with Wi-Fi, but those are actually Layer 1, not Layer 2 issues.

**Layer 3 is the network layer.** Throughout this book that will mean either IPv4 or IPv6. In 2014 or 2015, you'll probably be building your network to support IPv4 though I strongly recommend keeping IPv6 in mind as a future option. Fortunately, a good IPv4 network topology will generally support an IPv6 network perfectly well, although the converse might not be true.

For reasons lost to the mists of history, Layer 2 packets are called frames while layer 3 packets are called packets.

**Layer 4 is called the transport layer.** It's where things like the familiar Transmission Control Protocol (TCP) and User Datagram Protocol (UDP) are defined.

We talk about common protocols like HTTP and HTTPS as operating on TCP port 80 and 443 respectively. DNS is UDP port 53.

Things like TCP, UDP, and ICMP have special IP protocol numbers, defined in the Layer 3 part of the IP packet header. Then these higher layer protocols like HTTP and DNS are defined by their port numbers within the Layer 4 part of the header.

FOR REASONS LOST TO THE MISTS OF HISTORY, LAYER 2 PACKETS ARE CALLED FRAMES WHILE LAYER 3 PACKETS ARE CALLED PACKETS.

**Layers 5 through 7 contain dynamic session information for particular applications.** Since this book is focused on network topology, we won't be discussing these top layers.

## ETHERNET, BUS AND SWITCHED

Ethernet is a **bus architecture**, which means all of the devices are connected to it as peers. In a bus, all communications between all of the attached devices are visible to one another, and only one device can talk at a time.

This is how Ethernet was originally designed, and it's still how Wi-Fi works. However, it's not very efficient. If A wants to talk to B, it has to first wait for C to be finished talking to D. It's a

fine system if there aren't many devices on the bus, but as the number of devices grows, they wind up having to wait longer and longer for a chance to talk.

An **Ethernet switch** is a way to achieve the same architecture without all the waiting. In a switched Ethernet architecture, every device talks to just one other device, the switch. The devices are all organized using a small **hub-and-spoke** or **star** connection model.

If A wants to talk to B, it first sends its packet to the switch. The switch reads the destination Ethernet MAC address in the frame received from A, and looks it up in its internal MAC address table to find the interface that holds that address. Then the switch forwards the frame across its own backplane to the destination interface.

The other important advantage to a switch, besides eliminating waiting, is that each individual interface can operate at **full duplex**. Full duplex means the devices connected to a switch can send and receive Ethernet frames at the same time. So not only does the network not have to wait for every other device to stop talking before it sends a frame, it doesn't even need to wait for the switch to stop talking.

A switch is clearly much more complicated than a simple bus, but it also provides much more efficient use of bandwidth. If the backplane is fast enough, the switch can forward Ethernet frames between all connected devices simultaneously in all directions without waiting.

There are two important internal switch architectures. The switch can either operate as a **store-and-forward** or a **cut-through** switch.

Store-and-forward means the switch reads each full Ethernet frame and puts the information in a memory buffer before forwarding it to the destination interface.

A cut-through switch, on the other hand, reads only enough of the Ethernet frame's header to know where it needs to go. It then immediately starts forwarding that frame out the destination interface.

Cut-through switches have lower latency (the time it takes to forward a packet from one port to another). However, the other big difference is that when a store-and-forward switch receives a frame with errors, it can drop it, while a cut-through switch has already started transmitting. Fortunately, Layer 2 errors are fairly rare.

**FULL DUPLEX  
MEANS THE DEVICES  
CONNECTED TO  
A SWITCH CAN  
SEND AND RECEIVE  
ETHERNET FRAMES  
AT THE SAME TIME.**

# Network Segmentation

---

As we saw in the previous section, breaking a Layer 1 Ethernet bus into a whole lot of switch ports improves network efficiency. Communications on the network can take place simultaneously instead of sequentially.

Segmenting an Ethernet bus into small, switch-port-sized network domains is also important for security reasons. It turns out to be a bad idea that every device can see every packet, even the ones that aren't addressed to it.

Similarly, there are both efficiency and security reasons for segmenting Layer 3 networks.

Let's consider the efficiency reasons. Even with a switched Ethernet network, there are two important classes of frames that are automatically forwarded to every port.

The first are so-called **unknowns**. These are frames whose destination Ethernet MAC address is not in the switch's MAC address table. In other words, the switch doesn't know where to send the frames, so it has no choice but to forward them out every port, hoping somebody will be able to handle them.

For most types of normal devices, the only way the switch is able to populate its MAC address table is by listening to every port. Each time it sees a new source MAC address in a frame, it puts that MAC address into the table and assigns it to the interface that received the frame.

Devices that are new to the network, and ones that don't talk frequently, won't be in the table. (For basic housekeeping reasons, old MAC addresses are removed from the table if they haven't been seen in a long time.) The switch must "flood" these frames out all interfaces.

The other type of frame that must be forwarded out all interfaces are **broadcasts**. These are frames sent to the special Ethernet MAC address FF:FF:FF:FF:FF:FF. The most common types of broadcast frames are ARP and DHCP requests.

**THERE ARE TWO  
IMPORTANT  
CLASSES OF  
FRAMES THAT ARE  
AUTOMATICALLY  
FORWARDED TO  
EVERY PORT.**

**Address Resolution Protocol (ARP)** is a special protocol that allows the end devices on an IP network to find one another. Device A wants to send an IP packet to device B, but it's never sent a packet there before. So it sends out a special broadcast frame in ARP asking for the destination MAC address that corresponds to the IP address. The switch floods this frame to all devices on all ports and the one that owns that IP address responds. Then the IP session can start.

**Dynamic Host Configuration Protocol (DHCP)** is how devices that are new to the network (or just rebooted) can automatically learn important information like what IP address and mask they should use, and where the DNS server and default gateway are.

Naturally, the new device doesn't know anything about the network yet, so it has to use the same broadcast destination MAC address to request this information. The frame gets forwarded out all switch interfaces and the DHCP server responds.

Both ARP and DHCP frames are relatively rare compared to the rest of the application packets flying around a network, but it's important to notice that they get forwarded to everybody. Every device on the network has to receive these frames, look at them, and decide whether they need to do anything.

As the number of devices on a network grows, the number of broadcast frames grows. If you put enough devices on the same network, you start to run into serious efficiency problems once again. That's why we segment the network into smaller **broadcast domains** using functions like VLANs and routing.

One of the interesting features of IPv6 is that it uses multicast instead of broadcast for these types of functions. Multicast packets are only forwarded out the ports that explicitly say they want to receive them. It's much more efficient because devices don't need to see things that don't interest them.

The second important reason for segmenting your network is security. It's much easier to control communications between devices if the traffic must pass through a common control point where it can be examined. We tend to use special devices called firewalls for these control points.

Although there are firewalls that can operate at Layer 2, these are generally reserved for special situations. Layer 3 segmentation is easier to understand. And as I always say, things that are easier to understand are easier to support.

## VLANS AND TRUNKS

All modern Ethernet switches have a feature called Virtual Local Area Networks (VLANs). VLANs give you the ability to create groups of interfaces that are able to communicate together.

Devices in the same VLAN can send Ethernet frames directly to one another. If a device in a particular VLAN sends a broadcast frame, it's received by all the other devices in the same VLAN, but not by devices in different VLANs.

Sometimes a VLAN is also called a broadcast domain.

Suppose you have a switch whose interfaces are chopped up into two different VLANs and you want to connect this switch to a second switch that has the same VLANs. You want all of the devices on VLAN 100 on switch A to talk to all of the devices on VLAN 100 on switch B. And you also want the VLAN 200 devices on both switches to talk to one another.

There are two ways to make the connection. The obvious way is to simply configure a port on each switch on VLAN 100 and a different port on each switch on VLAN 200, then connect these ports together.

This will work, but it doesn't scale well. If you want a lot of VLANs, you won't have very many interfaces left for the devices they connect to. So we use something called a **trunk** instead.

A trunk is a single connection between two switches that contains a bunch of VLANs. We use a special part of the Ethernet protocol called 802.1Q to make the connection.

The 802.1Q protocol works with a very simple extension of the standard Ethernet protocol that adds VLAN information to the header. Now all of the frames for VLAN 100 are tagged for this VLAN and all of the frames for VLAN 200 are tagged for that VLAN.

Further, if one of the switches has a third VLAN, 300, you could use the same trunk to carry this VLAN to the second switch. Then you could use another trunk to a third switch to carry VLAN 300, even though the one in the middle doesn't have any interfaces configured for VLAN 300.

There's a special VLAN in every trunk called the **native VLAN**. When a switch receives a frame without 802.1Q VLAN tags on a trunk interface, it assumes the frame was intended for

VLANS GIVE YOU  
THE ABILITY TO  
CREATE GROUPS  
OF INTERFACES  
THAT ARE ABLE TO  
COMMUNICATE  
TOGETHER.



the native VLAN and forwards it accordingly. By default, most switches use VLAN 1 as their native VLAN, but this is configurable.

The trouble with the native VLAN is that, if I can somehow trick the switch and inject an untagged frame on the trunk interface, the frame will wind up on the native VLAN. In principle, you could use this technique to hop packets from one VLAN to another, perhaps violating a security segregation rule.

For this security reason, when it comes to native VLANs, I always follow a simple rule: Leave the native VLAN as VLAN 1 (the default value), then disable VLAN 1 on all switches. This way, even if somebody is able to hop a packet onto the native VLAN, it doesn't do any good.

## ROUTING

There are two ways to interconnect VLANs. You can interconnect them at Layer 2 using a bridge, or you can interconnect them at Layer 3 using a router.

The Ethernet switches we discussed in the previous section are a form of bridge. When you use a bridge, Ethernet-connected devices on one side can send Ethernet frames directly to devices on the other side. If I use a bridge to interconnect two VLANs, then it's as if there's only one VLAN. The devices talk directly.

Routing is a Layer 3 concept that forwards packets based on IP addresses, not Ethernet MAC addresses. To understand routing, we first need to talk about subnetting.

## Subnets

Every time you create an IP network, you define a subnet. A subnet is simply a way of grouping IP addresses that will be treated together. They'll generally be located on a common VLAN.

Every subnet needs a router of some sort to allow packets to leave the subnet and go to another. The router has a routing table that tells it how to direct packets between subnets.

A subnet is defined using a network number and a mask. The IP address of every device in the subnet has two parts, the network portion and the host portion.

Consider a simple example like 192.168.1.15 with a mask of 255.255.255.0. If you think of the mask in binary notation, each 255 is actually a string of eight 1s and no 0s. Everywhere you see a 1 is part of the network, and everywhere you see a 0 is a host.

**THERE ARE  
TWO WAYS TO  
INTERCONNECT  
VLANs: AT LAYER 2  
USING A BRIDGE OR  
AT LAYER 3 USING A  
ROUTER.**

In our example mask of 255.255.255.0, the first three bytes (192.168.1) are all network and the last byte (15) is the host portion.




Figure 1 - Subnet mask

All devices in the same subnet can talk to one another directly. In our example, that means everything from 192.168.1.1 to 192.168.1.254 can talk to each other. They can still talk to devices that aren't in this range, but the packets must be sent to the router, which will direct them to the right destination network.

In most cases, you'll define a single **default router** for each subnet. If a device wants to send a packet to 192.168.2.95, for example, it will compare this destination address to its own address and mask and see that the destination address is on a different network.

The only thing the device will do differently when the destination is in another subnet is to put the packet into an Ethernet frame whose destination address is the MAC address of the default router instead of the MAC address of the destination device. Then the router looks in its routing table to see where to send the packet next.

# High Availability

---

One of the most important considerations in any network design is **high availability**. The rule of thumb is that for anything important, any single device or any single link should be able to fail without affecting the application.

If you have an application on your network that needs high availability, you need to also make sure the network itself can recover. Nothing works if the network doesn't work.

In practice, of course, the network takes some time to recover from a lot of failure modes.

I used to make a practice of calculating the net failure probability and the expected system availability from different levels of redundancy. I don't do that anymore.

The problem with that way of approaching redundancy is that it assumes most of your failures will be hardware failures. In truth, most failures are caused by human error. A lot are due to software bugs. Even more are the result of configuration errors.

The thing that causes humans to make errors (assuming the humans involved are making an honest effort) is complexity. The more complicated something is, the more likely somebody will make a mistake with it.

**THE MORE  
COMPLICATED  
SOMETHING IS,  
THE MORE LIKELY  
SOMEBODY WILL  
MAKE A MISTAKE  
WITH IT.**

Complicated software has more bugs. Complicated IT infrastructure has more configuration errors. And complicated environments are harder to troubleshoot when something does go wrong.

All of this is to say that adding extra devices and links with complicated redundancy mechanisms could make your network less stable. However, with that caution, high availability is an essential feature of any mission-critical network.

## DEVICE-LEVEL HIGH AVAILABILITY

By device-level high availability, I mean the server itself is able to support some sort of high availability. This could take many forms. It could involve redundant power supplies, redundant network connections, or even fully redundant devices. These are all good strategies with different real and theoretical benefits.

The thing most likely to fail in most servers is the power supply. Network devices are somewhat less likely to fail, but if a network device like a switch fails, it affects a lot of devices. Full device redundancy is obviously best of all, particularly if the redundant devices connect to different network switches. We call this N+1 redundancy, meaning you have one more than you need.

## LAYER 2 NETWORK HIGH AVAILABILITY

There are two popular Layer 2 network high availability mechanisms, and dozens of less common ones. The first is link redundancy using Link Aggregation Control Protocol (LACP), and the second is Spanning Tree Protocol (STP).

### Link Aggregation Control Protocol

LACP is a way of logically bonding two or more (usually to a maximum of eight) physical links into a single **channel**. In Cisco jargon, this is called EtherChannel.

The two devices interconnected with LACP dynamically load balance traffic among the physical links. There are many different load-balancing algorithms, but they generally involve some mix of Layer 2 and 3 source and destination addresses.

The various load-balancing mechanisms mean that any one session will always take the same physical link. So if you make a channel out of four Gigabit Ethernet links, you get a theoretical maximum of 4Gbps aggregate throughput, but any single session can only be, at most, 1Gbps.

The nice thing about LACP is that it provides redundancy as well as load balancing. If one of the individual links fails, the two devices shift traffic to one of the remaining links automatically, usually in a couple hundred milliseconds.

The biggest problem with LACP is that it's designed to allow connections from one device to one other device. So by itself it doesn't help you if one of the devices fails.

**THE NICE THING ABOUT LACP IS THAT IT PROVIDES REDUNDANCY AS WELL AS LOAD BALANCING. ITS BIGGEST PROBLEM IS THAT IT DOESN'T HELP YOU IF ONE OF THE DEVICES FAILS.**

There are a few multi-chassis LACP implementations, however. Cisco has a feature called Virtual Switching System (VSS), which is available only on the large 6500 series switches. The company's Nexus series switches have a different feature called a Virtual Port Channel (VPC) that performs a similar function.

VSS solves the problem of distributing the channel between two switches by putting a very high-speed link between them and treating them as if they were a single switch. The controller module in one switch becomes the brain for both devices.

VPC takes a different and more complicated approach. Once again, it requires a fast link between the switches. Then the two switches exchange state information about the individual links in the channels they share.

## Spanning Tree Protocol

The other popular Layer 2 redundancy mechanism is Spanning Tree Protocol. STP works across the whole broadcast domain on all switches. Spanning tree solves two problems: It allows redundant links and redundant switches, and because that would cause loops, it also prevents loops.

Spanning tree works by logically rebuilding the switch-to-switch links as a tree. One switch is designated as the **root bridge**, which is the base of the logical tree. Then the switches figure out the shortest paths they can all take to reach the root bridge. Every link that isn't on one of those shortest paths gets disabled to remove the loops.

Spanning tree does all of this using a combination of individual link costs and bridge priorities. Usually you can leave the link costs at their default values and just play with the bridge priorities.

Select the switch you want as the root bridge and give it the lowest priority. Select another switch that should become the root if the first one fails and give it the second lowest priorities.

Setting your root bridge and backup in this way generally gives you good results, but you might have special reasons for wanting to force certain links to be active and others to be disabled. In that case, you can further manipulate the protocol by manipulating port costs and bridge priorities. But that's more advanced network traffic engineering and beyond the scope of this short guide.

**SPANNING TREE  
SOLVES TWO  
PROBLEMS:  
IT ALLOWS  
REDUNDANT LINKS  
AND REDUNDANT  
SWITCHES, AND  
BECAUSE THAT  
WOULD CAUSE  
LOOPS, IT ALSO  
PREVENTS LOOPS.**

## LAYER 3 NETWORK HIGH AVAILABILITY

There are two main ways of getting high availability at Layer 3. On network segments that have a lot of normal end devices like servers or workstations, or even embedded systems like IP phones, printers, or other appliances, the most important thing to make redundant is the default router for the segment.

This is done using a router redundancy protocol such as Hot Standby Routing Protocol (HSRP), Virtual Router Redundancy Protocol (VRRP), or Gateway Load Balancing Protocol (GLBP).

All of these protocols work in a similar way. They all present a virtual IP address that's used by the other devices in the subnet as their default gateway.

The virtual IP address can move around between the different physical routers. This way, if there's a failure of one of the routers or any of their upstream links, another one can automatically take over.

The other important Layer 3 redundancy mechanism is a routing protocol. Routing protocols such as Routing Information Protocol (RIP), Open Shortest Path First (OSPF), Enhanced Interior Gateway Routing Protocol (EIGRP), and Border Gateway Protocol (BGP) are used to distribute the routing tables among the different routers in the network.

One of the most important things to remember about network high availability mechanisms of all types is that you have to monitor them.

Suppose an automatic failover mechanism has engaged to protect your network and keep everything working transparently following a critical link or device failure. If you don't know the failure has happened, you can't fix it. Then you're left with the mistaken belief that you have a highly available network.

# How to Draw Network Diagrams

---

Before we discuss specific useful network topologies, I want to talk about network diagrams because I want to use diagrams to illustrate the topology types.

Good network diagrams are not hard to make, but I find them distressingly rare. Even network engineers with years of experience often make diagrams that are jumbled and hard to understand.

The important thing is to be clear in your own mind about what information you're trying to convey. It's better to draw several diagrams that show different aspects of the same network than to try to put everything on one sheet of paper.

It's a good idea to make a separate diagram for each network protocol layer. I usually start with a Layer 3 diagram to show the routing and IP subnets. Then I make a Layer 2 diagram showing the switch connections, trunks, and LACP channels. After that comes a Layer 1 diagram showing physical layouts of the devices.

In more complicated networks, I also do diagrams showing traffic flows, routing protocol distribution mechanisms, VPNs, and other important aspects of the network design.

It's important to draw each of these as separate pictures because they show different things. Combining them only confuses the information and makes the drawing harder to understand.

## COMMON NETWORK DIAGRAM SYMBOLS

There are some useful shorthand symbols to use in network diagrams. These are my rules, and they aren't universal, but I find them useful, particularly for drawing network topologies on a whiteboard.

My Layer 3 network devices are **circles**. My Layer 2 devices are **rectangles**.

**Clouds** are used to summarize parts of the network that aren't important for the purposes of the diagram. This could mean the Internet or a WAN or it could even mean a collection of internal network segments like user VLANs.

**Triangles** represent multiplexer devices, which used to be more common in network diagrams but aren't typically seen in enterprise network designs anymore. So I've started to use triangles for IP phones.

Everything else is a **square** or a **rectangle**, intended to represent a generic box.

Each symbol also has a specific marking on it that indicates exactly which type of device it is. When I draw these figures on a whiteboard, a router has an **X** on it and firewalls have a **>**.

If you're creating network diagrams using a drawing program or your map is being automatically generated by a software tool, the symbols will be much more elaborate.

	Hand-drawn	Visio	Auvik
Router			
Firewall			
Switch			
Load Balancer			
VLAN or IP Segment			
Cloud			

Figure 2 - Network diagram symbols



Note that all the diagrams you'll see in this guide are shown as I might draw them with Visio using Visio's symbols. For interest's sake and so you can see the differences, on page 32 I've also included a second version of a diagram, shown as the Auvik tool would map it.

I mention symbols you might use when drawing on a whiteboard because that's where a lot of good network designs start. A group of people gather around the board and figure out what needs to talk to what and how best to accomplish that.

Typically you'll be deploying the network to support some sort of application. And typically it will involve some client and server devices. It might involve an Internet connection and perhaps some firewalls. In a large organization, there could be a separate person responsible for each of these areas, which is why the whiteboard is such a useful place to start.

## DRAWING LAYER 3 DIAGRAMS

I always start with Layer 3 diagrams, which show the IP subnets and all Layer 3 network devices like routers, firewalls, and load balancers. The Layer 3 diagram must show all of the important network segments and subnets and how they're interconnected.

Layout is important. I like to show the layout so that it represents the flow of traffic in a broad sense. For example, if I have a bunch of servers being accessed by a group of users, I'll try to put the user network segments on one side of the picture and the servers on the other side.

Similarly, if I want to show how a LAN connects to external networks like the Internet, I group the external networks all on one side or at the top of the picture.

Or, if the point of the picture is to show a WAN with a large number of remote offices connecting to the same network, I'd probably show the connecting WAN in the middle of the picture and the various remote sites around the edge of the page.

Another layout consideration should be evident in all the diagrams you see in this book: I always draw my network segments either horizontally or vertically. About the only time I use a combination of vertical and horizontal is when I want to show a fundamental difference between the functions of the segments.

For example, I might draw all of my workstation and server segments horizontally, but then draw a special common network management segment vertically down one side of the page. This makes it immediately obvious that the management segment is special.

**LAYOUT IS  
IMPORTANT IN A  
LAYER 3 DIAGRAM. I  
LIKE TO SHOW THE  
LAYOUT SO THAT  
IT REPRESENTS THE  
FLOW OF TRAFFIC IN  
A BROAD SENSE.**

The Layer 3 diagram should show any high availability mechanisms and redundant network components or redundant paths. It's customary to show router redundancy protocols as an **elongated ellipse** that covers the router links included in the high availability group.

The other important thing about Layer 3 diagrams is that they should *only* include Layer 3 objects. I don't want to see switches in a Layer 3 diagram, for example. I don't want to see any kind of indication of trunk links on a Layer 3 diagram either.

You can show a switch on a Layer 3 diagram only if it's a Layer 3 switch, and then only because it functions as a router. Including Layer 2 objects like a switch in a Layer 3 diagram is confusing, particularly in more complicated pictures.

Another useful thing to put into a Layer 3 diagram is organizational boxes. If there are security zones or interesting groupings of users by function or servers by application, put them together on the picture, put a box around them, and label the box clearly. It's then easy to see the exact network path those users take to reach their servers.

In more complicated network designs, I often use a base Layer 3 diagram showing the VLANs, routers, and firewalls. Then I create several other diagrams to lay over the base diagram. I might have an overlay diagram showing the routing protocol design, another one showing VPNs, and still another showing key application data flows, if that's an important consideration.

## DRAWING LAYER 2 DIAGRAMS

Layer 2 diagrams show Layer 2 objects like switches and trunks. They include critical information like which VLANs are included in which trunks, and they show spanning tree parameters like bridge priorities and port costs. In many cases, this is too much information to show easily, so I generally use callout boxes to hold some of the information.

Unlike Layer 3 pictures, Layer 2 diagrams don't need to be laid out in any special way. The most important thing is to keep the picture clear.

If two devices are intended to provide redundancy for one another, then their positions on the page should be related. They should either be located beside one another or in parallel locations on opposite sides of the picture.

**UNLIKE LAYER 3 PICTURES, LAYER 2 DIAGRAMS DON'T NEED TO BE LAID OUT IN ANY SPECIAL WAY. THE MOST IMPORTANT THING IS TO KEEP THE PICTURE CLEAR.**

If there are different link speeds, they should be indicated in the diagram. I usually show link speed with the thickness of my diagram's connecting lines. The faster the link, the thicker the line.

Sometimes I also use color to indicate special properties of different physical links. For example, I might make fiber optic cables red and copper cables blue. (Technically the cable type is Layer 1 information, but because it doesn't tend to cause confusion in the picture, it's alright to include in your Layer 2 diagram.)

## DRAWING LAYER 1 DIAGRAMS

I usually use Layer 1 diagrams to show physical connections between devices, but they're also useful for showing cabinet layouts.

Layer 1 diagrams should show port numbers and indicate cable types. In a network that includes many different types of cables, such as fiber optic cables, Category 5 / 6 / 7 copper cabling, and so forth, it's useful to give each cable type a different color.

If there are patch panels, particularly if you want to document how patch panel ports map to device locations and switch port numbers, this information belongs on the Layer 1 diagram.

And if there are different link speeds, you might want to give them different line weights, as described previously for Layer 2 diagrams.

Another type of diagram that's often useful in data center designs is a cabinet layout. It's a diagram that shows exactly what you would see when looking at the front (and sometimes also the back) of the cabinet. A cabinet layout is helpful when you need to tell a remote technician how to find a certain piece of equipment.

## DRAWING COMBINED-LAYER DIAGRAMS

There's one very special type of diagram in which it's possible to combine Layer 2 and 3 in a single picture. Such a combined diagram is sometimes useful if you have combined Layer 2 and 3 switches and you need to show the relationship between these layers.

**A COMBINED  
DIAGRAM IS  
SOMETIMES USEFUL  
IF YOU HAVE  
COMBINED LAYER  
2 AND 3 SWITCHES  
AND YOU NEED  
TO SHOW THE  
RELATIONSHIP  
BETWEEN THESE  
LAYERS.**

The diagram above illustrates a case where showing the relationship between layers is important. The Layer 2 / 3 switch is taking part in redundancy protocols at both Layer 2 and Layer 3.

In the example, four layer 2 / 3 switches are used to interconnect two VLANs. In the inset, which is a pure Layer 3 picture, the arrangement looks highly redundant. But the mixed-layer picture shows you that the switches are in fact cascaded one after the other with no link redundancy at all.

Look in particular at the devices labeled A and B. On the pure Layer 3 diagram, these devices appear to be just one hop apart with four redundant paths. However, the mixed diagram shows that packets from A to B actually have to pass sequentially through all four switches.

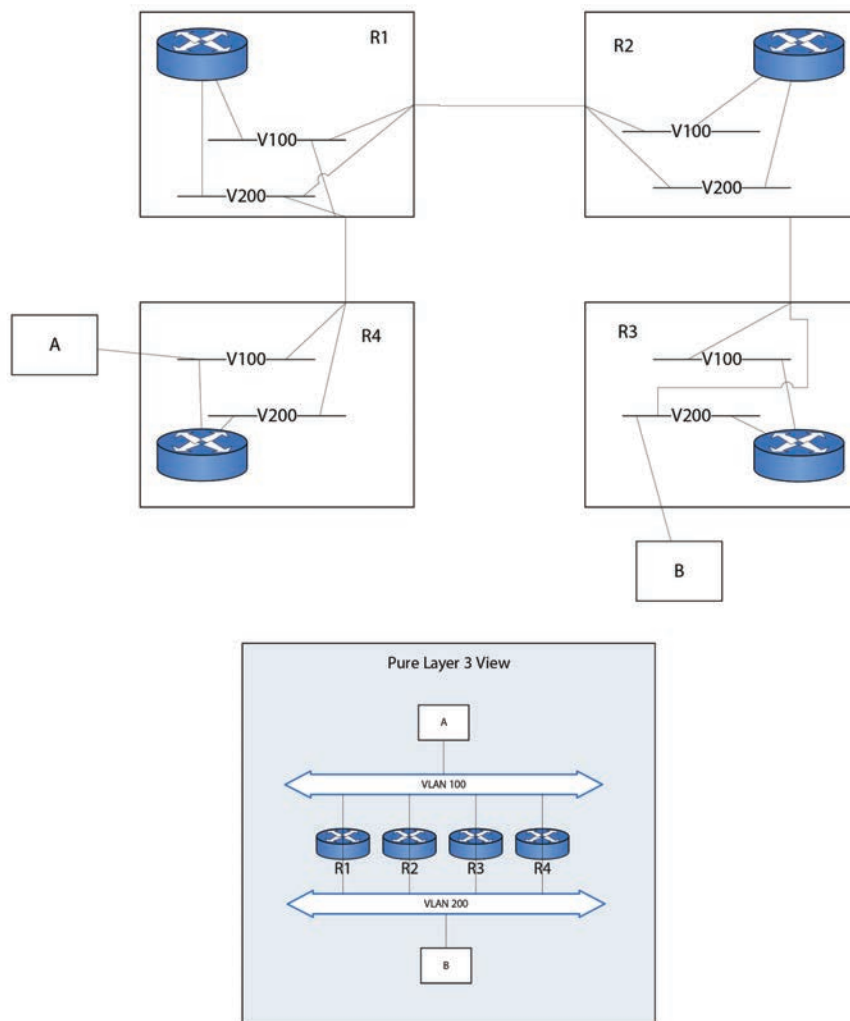


Figure 3 - Combined Layer 2 / 3 diagram

A combined-layer view is also useful when thinking about things like HSRP configuration. Which switch will be the default gateway for each VLAN? And, related to this, will the packets from A to B take the same path as the packets from B to A? None of these details appear in the pure Layer 2 or the pure Layer 3 picture.

Instead, we show the relationship in a combined-layer diagram by drawing boxes for the Layer 2 switch with the VLANs inside it, connected to the Layer 3 router, also inside the switch. The VLANs are connected to trunk interfaces to another Layer 2 / 3 switch.

Note that while this diagram can show the interaction between the layers, it doesn't make either the Layer 2 or Layer 3 network design terribly clear. I'd actually draw all three as separate diagrams, each showing a different important aspect of the network design.

## HOW TO MAP AN EXISTING NETWORK

Unless you already have a really good set of network documentation, the best place to start learning about network diagrams is by mapping your own network. I always start with pencil and paper, and I expect to draw and redraw the picture many times.

The first time I draw a diagram, it will generally be at Layer 3. I'll log into one device that I think is in the middle or the routing core of my network, and I'll look at its routing table.

On Cisco switches and routers, the command to see the routing table is "show ip route". On Cisco firewalls, it's just "show route". Most other vendors have a similar command.

The output of the "show route" command gives you a list of routes. Each route is defined by a route **prefix**, which defines the subnet range, and a **next hop**, which defines how to get to that subnet. The next hop will either be something like **local** or **connected**, or it will be the IP address of the device that takes care of forwarding packets to the destination.

I draw a small circle to represent the forwarding device, and around it I draw short lines to indicate the locally connected subnets.




Figure 4 - Step 1 of mapping an existing network

Then I look at all of the next hops that are IP addresses instead of directly connected interfaces. Pretty much by definition, to be useful as a next hop, these devices must be on directly connected networks. Place these devices on the appropriate segments, as in Figure 5.

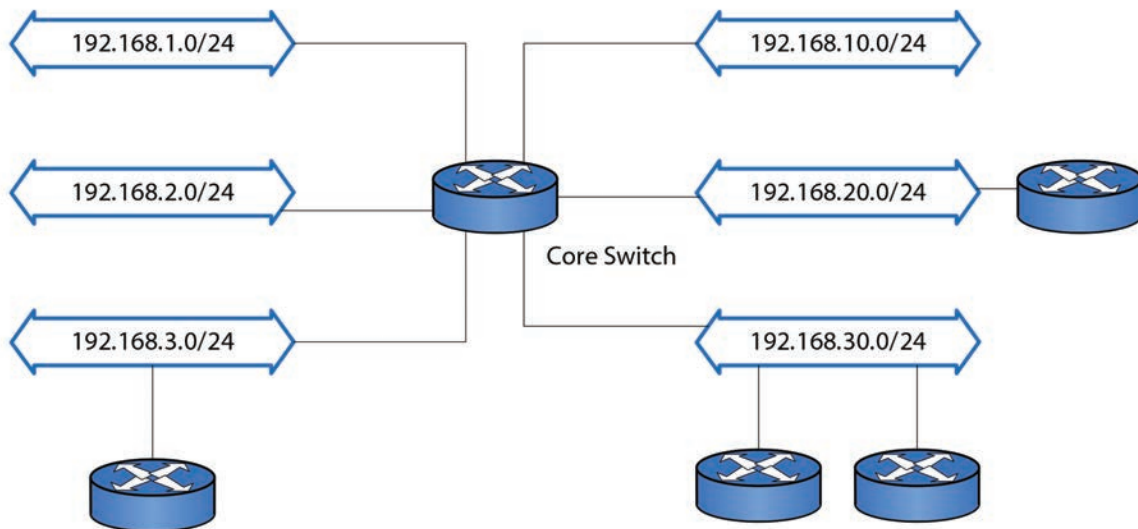


Figure 5 - Step 2 of mapping an existing network

Now you need to log into each of the next hop devices and do the same thing. Soon you'll start seeing devices you've already drawn. Once you've drawn a circle for all of the devices inside your network, you're done.

The next step is to rearrange the networks and the Layer 3 devices until the picture makes sense. I often like to arrange my network diagrams in what I call **security order**, with the public Internet at the top and the most privileged networks (probably the user segments)

at the bottom. Segments that are at a similar security level will be at the same approximate elevation between the top and bottom of the page.

Another common arrangement is to put clients on one side and servers on the other to depict the network in terms of what might be thought of as **traffic flow order**. In this type of arrangement, the segments that are in some sense at a parallel level in terms of traffic flow will be adjacent to one another. In more complicated networks, it's not at all uncommon to have several different views.

Don't worry about fitting your map all on one page. The important thing is that it's readable.

**DON'T WORRY  
ABOUT FITTING  
YOUR MAP ALL  
ON ONE PAGE.  
THE IMPORTANT  
THING IS THAT IT'S  
READABLE.**

In fact, for very complicated networks, I usually draw one very high-level network diagram that shows many distinct zones and how those zones are related to one another. Then I have a separate picture for each zone.

After I'm happy with my Layer 3 diagram or diagrams, I move on to Layer 2 and Layer 1. These diagrams are considerably more difficult to do just by logging onto the devices.

You get Layer 2 topology information from looking at your switch configurations and identifying things like trunk interfaces. But not all switch-to-switch links are configured as trunks, and not all trunks are connected to other switches. This is where mapping starts to get a little more complicated.

One useful technique for figuring out Layer 2 next hop devices is to look at the MAC address table. The switch uses the MAC address table to figure out how to forward Layer 2 Ethernet frames. It maps all of the Ethernet MAC addresses that the switch knows about to physical interfaces.

If there's only one MAC address associated with a particular interface, then it's likely (but not definite) that this interface is only connected to a single physical device. If there are many MAC addresses then it's likely (but not definite) that this interface connects to another switch.

But that's a lot of work. If I have to map a network by hand using the MAC address table, I often use the trick of importing the MAC address table into a spreadsheet and sorting it by interface to see which ones appear to connect to more than one MAC.

On Cisco switches, an easier technique is to use **Cisco Discovery Protocol (CDP)**. CDP is a Layer 2 protocol that allows Cisco devices to dynamically learn about one another. If you log into all of your switches and turn on CDP everywhere, then wait a few minutes for the devices to discover one another, you can get a much better first iteration of your Layer 2 map.

**MAPPING A NETWORK BY HAND CAN BE A TIME-CONSUMING AND TEDIOUS PROCESS.**

Once again, I pick one device that I think is in some sense central to my network and draw a little rectangle in the middle of the page to represent it. Then I draw similar boxes to represent all of the other switches that I can find by the methods described above.

Once again, this is an iterative process. Log into each of the next hop devices and repeat the procedure, again drawing all of their connections.

Phew! As you can see, mapping a network by hand can be a time-consuming and tedious process.

Even if you're using a **drawing program** like Visio, Gliffy, or LucidChart, you'll still need to manually collect all of the information you'll be plotting. Drawing software only helps you draw.

But here's where **automation tools** come in handy. There are several tools on the market, some free, some rather expensive, that can map your network for you. Each one has its strengths and weakness.

## SolarWinds

The SolarWinds Network Topology Mapper generates maps that are like the first passes I do on paper; generally speaking, they're not all that meaningful. The software doesn't know anything about the functions of different network segments or systems. It also doesn't do a great job of distinguishing between the Layer 2 and Layer 3 functions of the same physical device.

It's also important to remember that the maps SolarWinds generates are essentially snapshots. They don't automatically evolve if the network topology changes. If you want to see how things have changed, you need to run the discovery tool again.

## Spiceworks

Spiceworks provides a free network mapping tool that many people find useful. The tool has fewer features than SolarWinds, but it's still quite good at gathering connectivity information for an undocumented network.



## netViz, net-viz, and netTerrain

One of the most advanced tools for network mapping used to be netViz, which was a very good package until it was abandoned by CA. Confusingly, there's another Python package called net-viz, which doesn't have a lot of features. Even more confusingly, Graphical Networks has developed a netViz work-alike called netTerrain. netTerrain is useful but requires significant amounts of user input to produce meaningful network maps.

## Auvik

Auvik is a particularly useful tool for network mapping. It has two important advantages. It does a much more sensible job of laying out your network topology than most other systems. And it automatically updates your network maps in response to topology changes as they happen.

Particularly for smaller networks, Auvik can actually produce useful and usable network diagrams right out of the box.

# Useful Topologies

There are several useful network topologies for LAN and WAN construction. Next we'll talk about these designs and where they're best applied.

## SINGLE SWITCH

The most basic type of small office network is a single switch. It's probably the fastest and most stable network you can build. But of course, it has the problem of being a massive single point of failure.

In a single-switch network, one switch supports everything: workstations, printers, and servers. The single switch performs the functions of access, aggregation, and network core all at once.




Figure 6 - Single switch small office network topology




Figure 6 Cont. - Small office network topology

For this single-switch network, I've shown separate Layer 1, 2, and 3 diagrams.

- The Layer 1 diagram shows Layer 1 information such as switch ports, media types, and link speeds.
- The Layer 2 diagram shows Layer 2 information, which in this case is just VLAN numbers.
- The Layer 3 picture shows how the VLANs correspond to subnets, and how the VLANs are interconnected.

The really important thing about these diagrams is to notice that the Layer 3 picture shows completely different types of information than the other two pictures. However, you'll probably have noticed that in a simple network like this one there's really no reason not to combine the Layer 1 and 2 information.

The only extra thing that might be in the Layer 1 diagram that doesn't belong in the Layer 2 picture are patch panels and physical layouts. These items may not be relevant in a very small network like this one though.

## A single switch topology is useful for

- Small networks with very few devices all located in close proximity to one another.

## Avoid when

- You have more devices than will fit on a single switch.
- There are geographical challenges, such as individual devices more than 100m away from the switch.

## Single switch with Internet connection

Figure 7 shows a slightly more complicated single-switch network. It's a little different from the previous network because of the inescapable fact that the Internet needs to be treated differently.

The Internet is the most hostile place on Earth. *Do not* put your external Internet connection on a VLAN on your internal switch.




Figure 7 - Small network with Internet connection topology (Visio version)

For interest's sake, here's how the same diagram would look if the Auvik tool had mapped it.

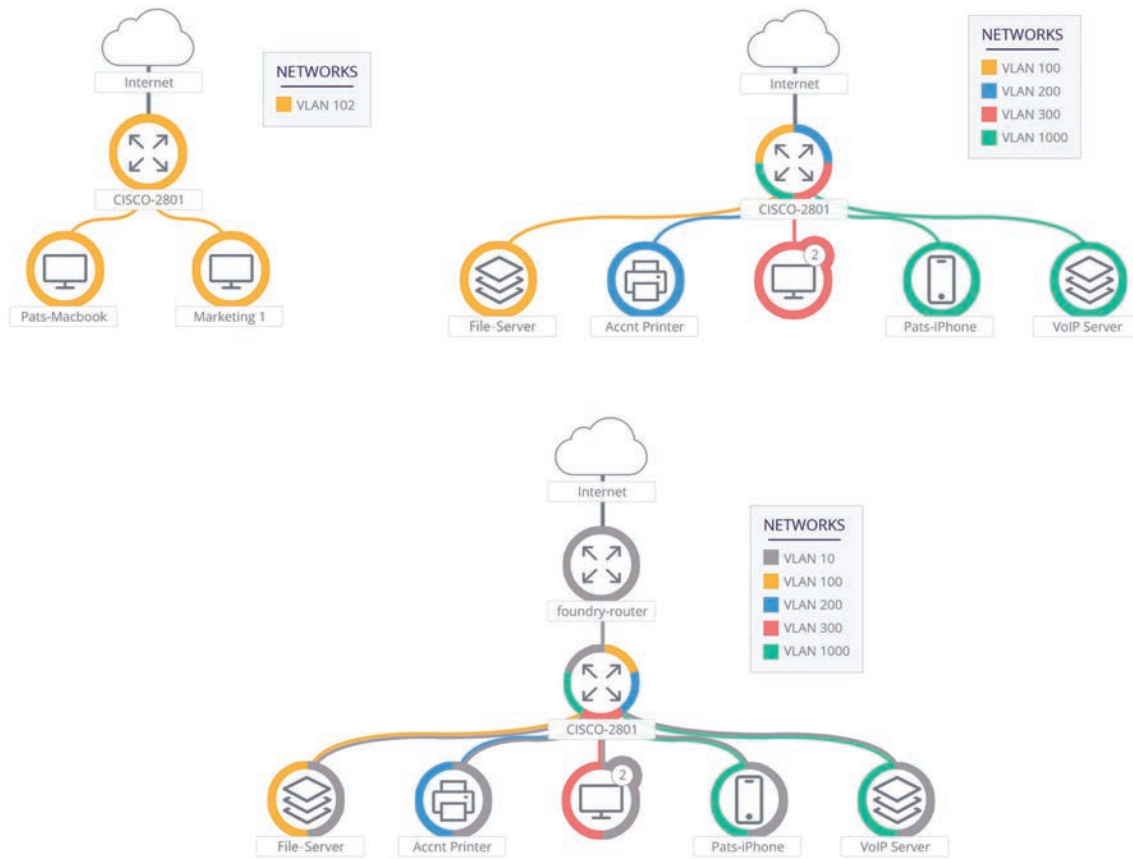


Figure 8 - Small network with Internet connection topology (Auvik version)

For this type of small network design, you really have three choices.

1. Put all the internal devices on one flat segment, together with the inside interface of the Internet firewall.
2. Split the internal network into multiple VLANs, and put every VLAN on a different interface on the firewall.
3. Interconnect the internal segments with the Layer 3 functions of the switch itself.

Each of these options has its strengths and weaknesses.

The first option is great if you have a very simple network. It's certainly how most home networks are set up. But if there's any requirement to separate devices, it's not useful.

You might separate the devices onto different segments for security reasons. For example, devices like IP phones and printers that run their own IP stacks but can't run anti-virus

software are vulnerable to certain types of attacks, so I like to separate them from sensitive servers and workstations.

The second option, putting each segment onto a different firewall interface, is one way to solve the security problem. It provides good separation of the logical segments. But it has two main drawbacks.

First, it doesn't scale well. If you need to expand the number of segments or even the number of devices, you'll quickly run out of resources on that firewall, both interfaces and throughput.

Second, it can be quite awkward to configure firewall rules that will allow arbitrary Windows server traffic. If you wind up configuring wide open firewall rules between two segments, you probably shouldn't be using a firewall.

So the third option is to use the Layer 3 routing features of the switch to separate segments. You can still put some segments on separate firewall interfaces if you really want the secure separation. The biggest disadvantage with option 3 is that adding Layer 3 switching features makes the switch more expensive.

## A single switch with Internet topology is useful for

- Small networks with Internet connections.

## Avoid when

- Your network is larger or geographically dispersed.
- You expect your network to grow quickly.
- You want to provide a web site to the external Internet. (In that case, you should really build a proper DMZ.)

## TOUCHDOWN SEGMENT

The touchdown segment is a technique for connecting a LAN to one or more WAN routers. The idea is very simple: Create a separate VLAN and put the other routers on it, then connect this segment to the internal network through either a router or a firewall.

It's a good technique for several reasons.

First, you might have security concerns about the external networks. In this case, the router or firewall that connects the internal network to the touchdown segment can be used as a security control point. It can have access control lists (ACL) or firewall rules to restrict what traffic can enter the internal network from the external routers.




Figure 9 - Touchdown segment topology

Second, it's often necessary to separate the routing protocols. You don't want external networks controlling your internal routing. At the same time, you often do need to learn about the external network's routing, and you similarly need that network to know about at least some of your internal routes. So it makes sense to carefully redistribute and filter the routes in both directions.

Third, related to this routing issue, reaching the external networks through a common point on your network means the router or firewall is an excellent place to summarize routing tables. This greatly simplifies the routing of both internal and external networks, which in turn contributes to better routing convergence times and more stable networks.

Of course, the touchdown segment is also obviously useful for simple reasons of organization. If you can put all of your WAN and external-facing routers in a common corner of your network, it's very easy to handle things like troubleshooting or expanding the network design.

## A touchdown segment topology is useful for

- Connecting WAN routers and external networks to an internal LAN

## Avoid when

- You're connecting to an extremely large number of external networks.
- The external WAN connections have a very high aggregate bandwidth (in which case the touchdown segment could present a bad congestion point).

## THREE-TIER

The classic three-tier network has a core tier, aggregation tier, and access tier, as shown in Figure 10. You'll often see this kind of topology in textbooks and in Cisco's reference architectures for large corporate networks.




Figure 10 - Three-tier large LAN topology

The main concept of the three-tier network is that the normal end devices, such as workstations and printers, should connect to access-tier switches. Access switches then connect to aggregation switches, and aggregation switches connect to core switches.

Three-tier design is really a Layer 1 concept. There's nothing in this basic design that dictates how VLANs or IP segments should be distributed or isolated. In principle, you could do just about anything you wanted. For example, you could have a particular group of users geographically dispersed among any of the access-tier switches.



The strength of the three-tier network architecture is its scalability. If the access switches each have 48 ports and the aggregation switches also have 48 ports, and there are two core switches also with 48 usable ports, then your network could readily support over 100,000 end devices. You could even increase this further by using larger switches at the core or aggregation tiers, or by expanding the core to include more physical switches.

In practice, if your network is large enough to need this many access ports for end devices, then you probably need to do a lot of segregation. For this reason, it would be unusual to freely distribute VLANs across the network.

Further, if you have this many end devices, you'll probably encounter serious congestion problems if all the traffic between network segments concentrates on the uplinks between the aggregation and core switches.

In Figure 11, I show a subtle modification to the standard three-tier concept. Instead of each switch connecting to one upstream switch, I place two redundant connections at each point. Even the connection between the two core switches is redundant.




Figure 11 - Modified three-tier topology

The result is that I've treated the aggregation-tier switches in pairs, just like the core switches. You could even opt to have a separate link between the two paired aggregation switches. But doing this would require being careful about your spanning tree configuration, as it might just be another chance to make loops in your network.

In the real world, where scaling to hundreds of thousands of LAN interfaces is rarely required, the three-tier network architecture is used in very limited ways.

You might have a core tier with a few switches connected to an aggregation tier. You might decide to connect critical servers directly to some of the aggregation or even the core switches and a handful of access switches that are really only used for geographical reasons to get connectivity to physically remote groups of users.

That would still be considered a three-tier network, and would likely prove to be more useful.

### A three-tier topology is useful for

- Very large corporate networks with many workstations.

### Avoid when

- You have a smaller network that doesn't need many end devices.
- You have lots of distinct security zones that need firewalls between them.

## WEB-TIER

Figure 12 (on the next page) shows a classic web-tier network design, a standard DMZ.

A web-tier design doesn't even need to involve web servers. Any application that runs on servers and has a back-end could, in principle, use this model. However, web-based services are an extremely common way of delivering applications.

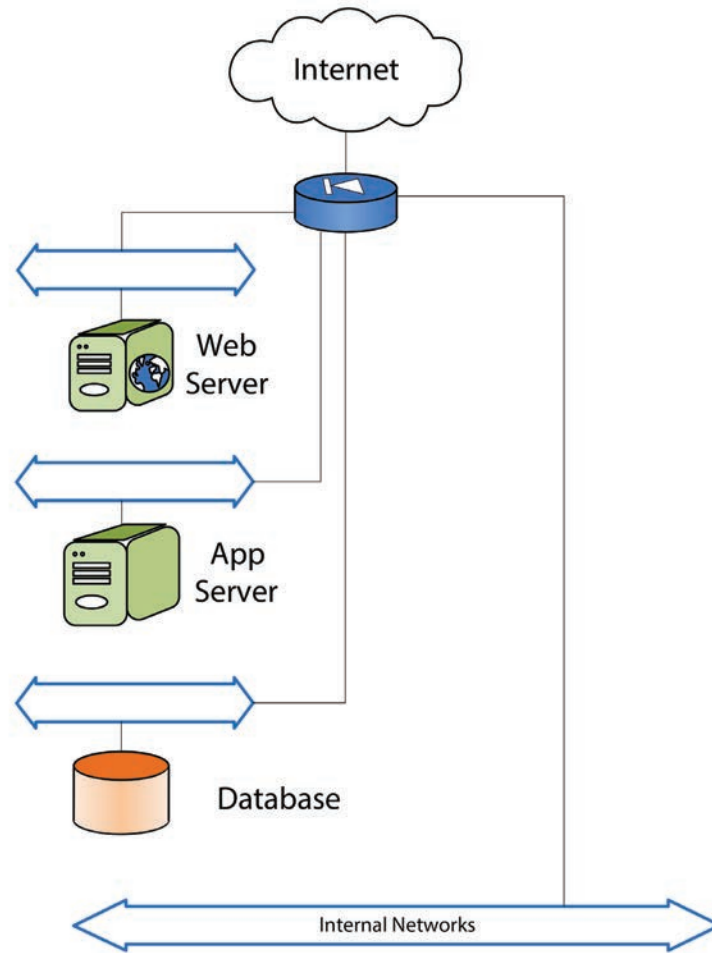


Figure 12 - Web-tier DMZ

In this topology, we put a firewall in front of everything. This is the Internet firewall. Its function is to protect the web servers from direct access from the Internet.

In most cases, the firewall will do two things. It will allow access to the web server and only the web server based on its destination IP address. And it will allow access only to the specific HTTP and HTTPS TCP ports (80 and 443) on this web server.

In many cases, it will also do Network Address Translation (NAT) to allow the web server to have a private address but still be accessed from the public Internet.

Inside the external Internet firewall is the network segment holding the web server. The external connections can only reach this web server. In turn, the web server is allowed to send packets through a firewall to an application server on a separate internal segment. And the application server can communicate to a database server on yet another segment, also behind a firewall.

In some DMZ designs the web and application servers might be combined, or the application and database servers might be combined. It's an extremely bad idea, though, to put your real data on the web server unless that data is completely public, like brochure material. Web servers get compromised so often that you have to operate on the assumption it will happen sooner or later.

There are actually a few ways to construct a web-tier network. The simplest is to use the same single firewall for all three server segments as well as the Internet access. This is a simple design, but it requires special attention to firewall rules. One of the critical features of this design is that the servers in each tier can only be accessed from the servers in the next higher tier.

A more complicated web-tier design uses different firewalls to separate the tiers. This requires either putting two firewalls on each LAN segment (Figure 13) or deploying the individual servers with multiple LAN interfaces (Figure 14).




Figure 13 - Web-tier DMZ with multiple firewalls

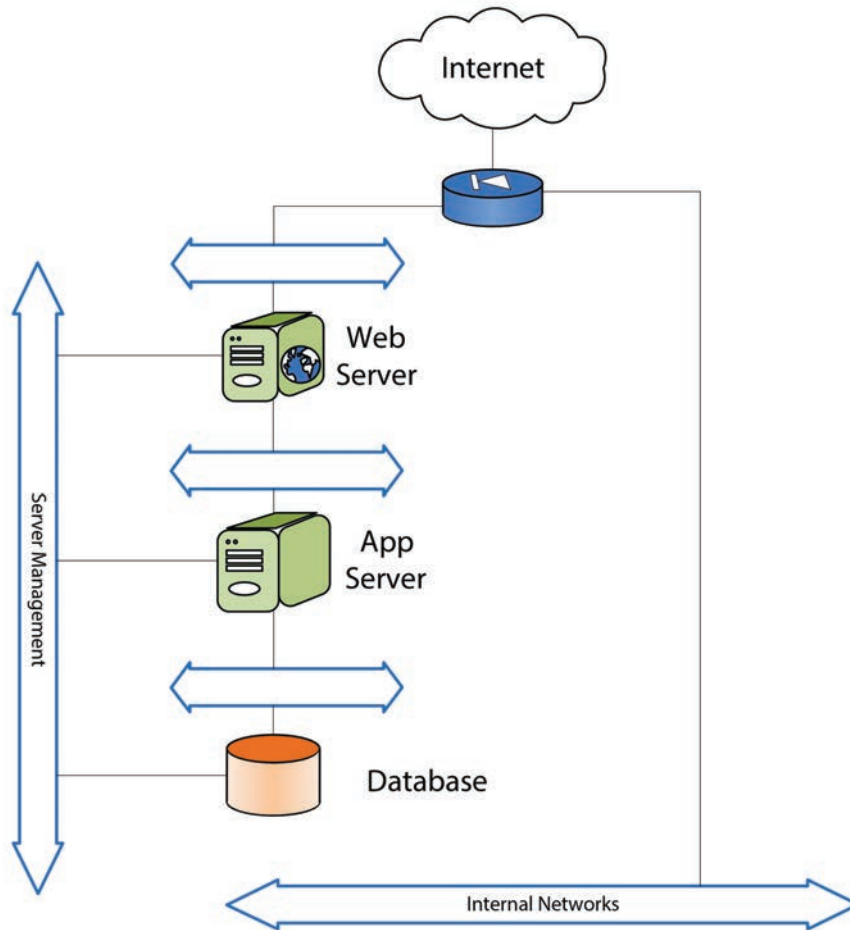


Figure 14 - Web-tier DMZ with multiple server LAN interfaces

I personally believe using multiple firewalls is marginally more secure because a single firewall configuration can be compromised if an attacker is able to somehow subvert the single firewall. Then, because the single firewall connects to all segments, the attacker could gain direct access to the potentially sensitive database.

Conversely, in a multiple firewall configuration, the attacker must compromise the first tier before moving on to the second and third tiers.

The disadvantage to the multiple firewall topology is that it forces greater complexity on the servers. They must now have different routes for traffic going outward to the client than for traffic going inward to the next tier.

Single versus multiple firewalls is definitely a trade-off. You'll have to choose the balance that's right for your environment.

Another important modification to the multiple-tier web model involves adding a load balancer, as shown in Figure 15. I always put the load balancer inside the firewall.

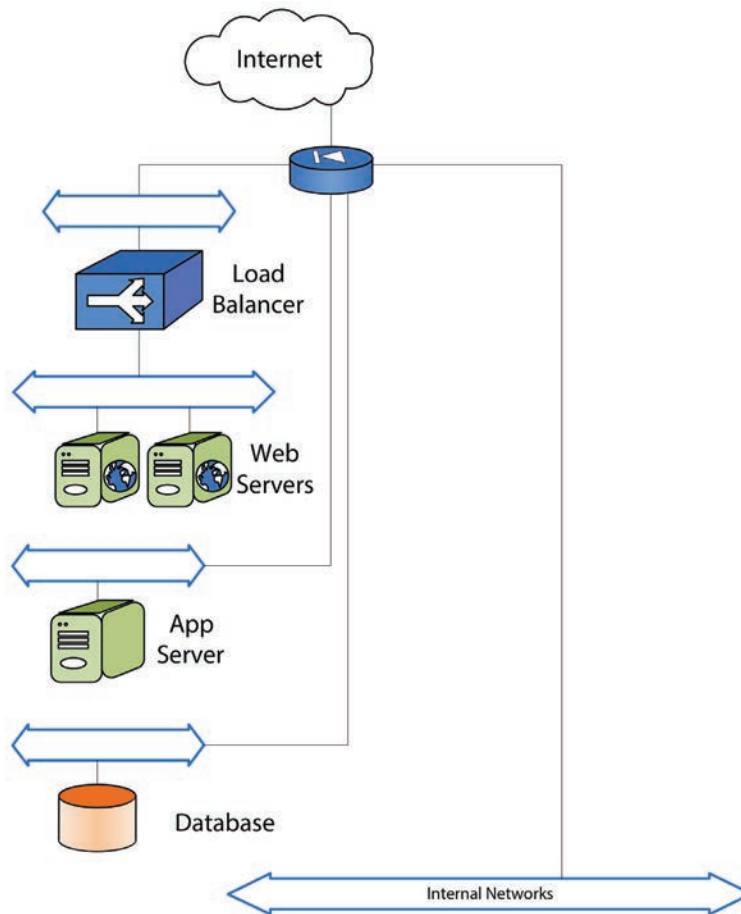


Figure 15 - Web-tier DMZ with load balancer

I also like to terminate SSL sessions for HTTPS on the load balancer because it simplifies how certificates are handled for the end user. And I like to do the NAT on the load balancer because load balancers are fundamentally NAT engines. It's better and simpler to do NAT only once rather than first on the external firewall and again on the load balancer.

Finally, there are a lot of cases where you won't have a database at all. Your application might not need or use one.

The important concept of the web-tier design is that sensitive systems should not be accessible from the external network. You need to go through multiple layers to reach them, and you ideally want to build your applications so there's inspection and validation of the data at each stage.

This ensures that an attacker can't just craft a request at the web layer and pass malicious code all the way to the sensitive internal systems.

## A web-tier topology is useful for

- Internet DMZs where you need to provide access to users on an external untrusted network like the Internet.

## Avoid when

- Your applications touch only trusted networks.
- Your applications can't be readily separated into tiers.

## HUB-AND-SPOKE

Hub-and-spoke networks are similar to the three-tier LANs we discussed earlier. Or, more precisely, three-tier networks are specific examples of a hub-and-spoke topology. Hub-and-spoke networks are also a common WAN topology.

In a basic hub-and-spoke network topology, there's a single central device that connects to several remote devices. Usually these will be Layer 3 connections in a WAN, but they could be Layer 2 trunk connections in a LAN.

Layer 3 hub-and-spoke networks scale very well because you can increase the size of the hub by creating a ring of hub routers as the routing core.

The trouble with this model is that the hub device is a single point of failure for the entire network. So a common variant uses dual redundant hubs. Both single hub and dual hub options are shown in Figure 16.




Figure 16 - Single and dual hub-and-spoke topologies

A more typical modern WAN topology is the Multiprotocol Label Switching (MPLS) cloud. It looks somewhat like a hub-and-spoke network design, but in fact most MPLS WANs allow any-to-any communication. This means any of the spoke routers can communicate with any other spoke router, as well as the hub.

The key defining feature of a hub-and-spoke topology is that the spoke nodes are not interconnected to one another. If I can connect the network in a flatter topology, it will generally provide greater reliability overall.

To see why this is, consider the difference between a triangle and two separate remote nodes, each individually connected to a central node. It's clear the triangle topology will be more resilient against different failure modes.



## A hub-and-spoke topology is useful for

- Connecting geographically diverse network zones.
- Campus LANs or tall buildings where you need to interconnect many remote areas to a common central data center.

## Avoid when

- There's a lot of spoke-to-spoke communication. (For example, this kind of topology wouldn't work well if you have VoIP users calling one another between spokes.)
- The main data center is not located at the hub.

## RING

A ring topology features three or more interconnected switches. The nice thing about a ring topology is that it provides redundancy with a minimum number of links. Any one device or link can fail without disrupting connectivity for any of the others.

I usually make rings in a small part of the network, then connect other switches or routers to the ring using redundant links to two different ring members. I show this design in Figure 17.




Figure 17 - Ring topology

There are two particularly useful places to deploy ring structure in networks. One is a ring of access switches and the other is a ring of core switches.

In principle, you could also have a ring of aggregation switches in a three-tier network model. But it's hard to imagine a network large enough to require such a structure yet small enough to not suffer from spanning tree hop count problems.

The old spanning tree rule of thumb is that a network becomes unstable if you have more than 16 hops between any two devices. You'll typically get a high number of hops when the ring contains a dozen or more switches, or in smaller rings when each switch in the ring has many downstream switches.

I sometimes like to use rings of access switches in cases where, usually for reasons of physical geography, I can't run all of the individual access switches all the way back to the core but I still need the redundancy of multiple connections.

For example, I might connect the switch on the first floor to the switch on the second floor, which connects to the switch on the third floor, which in turn connects back to the first floor. Then I separately connect two of these switches to two different switches in my network core for full redundancy without requiring a full mesh and using as few expensive cable runs as possible between the floors.

I also like to use rings of core switches in situations where I need to have more than two core switches or routers.

For example, this strategy was useful in a data center network I designed. In that case, the client had a relatively large number of access switches connected directly to the core switches, but it was more cost-effective to use fixed form-factor core switches than to go to a higher density chassis switch at the core.

In the end, we connected several of the core switches together in a ring. The access switches were then connected to the core switches using redundant uplinks to two of the core switches. A similar design is shown Figure 18.

Here, the core switches are interconnected using the highest speed links available. That might be 10Gb or 40Gb Ethernet if the switches support it, or it could be a LACP bundle of links.

The ring in this case is broken up into two groups. Switches A and B are used for server connections in the data center and switches C and D are used for downstream connections to the user access switches. Each downstream access switch has redundant uplinks to the core switches for fault tolerance.

Note that the higher speed links are indicated with thicker lines, a drawing technique we discussed earlier.

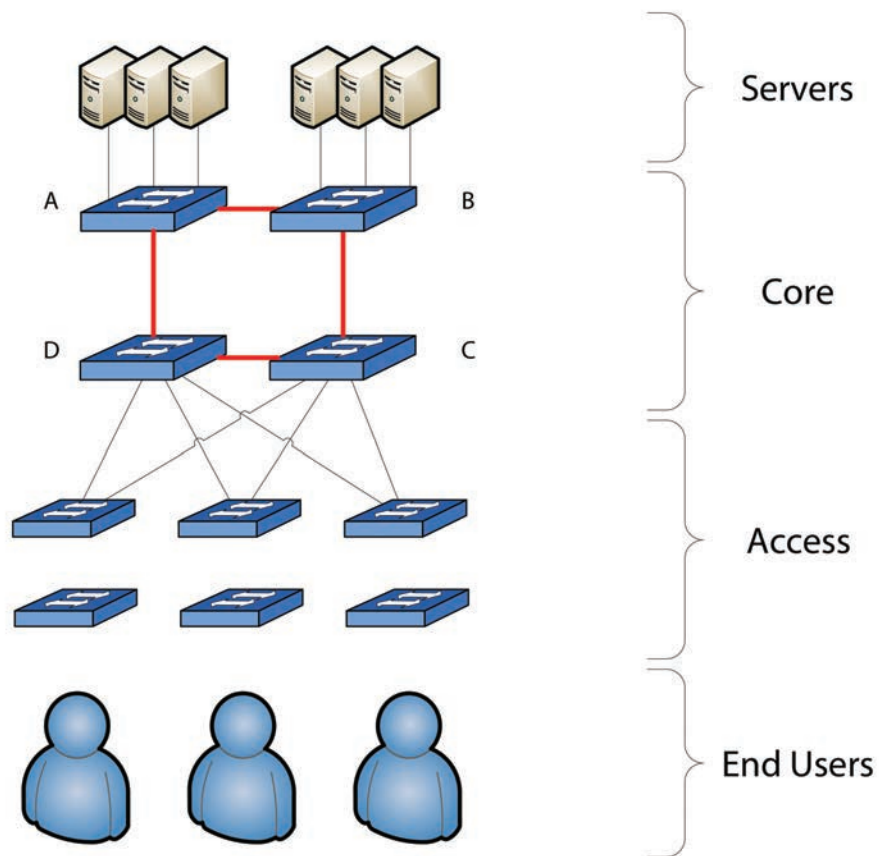


Figure 18 - Backbone rings

## A ring topology is useful for

- Interconnecting a group of switches in a redundant topology that uses as few links as possible.
- Situations where you need to save cabling costs due to geographical problems.
- Conserving switch uplink ports when bandwidth is not an issue.

## Avoid when

- The number of switching hops is high.
- Your network has high traffic, particularly if there's real-time traffic such as voice or video. The ring topology aggregates data flows onto a small number of backbone links. This can cause congestion, which results in latency and jitter problems.

## INTERNAL FIREWALLS

Security is becoming an increasing vital consideration in network design. I have a lot of clients who deploy firewalls to separate internal security zones within their networks, not just at the perimeter.

There are a lot of reasons to do this. You might have a set of servers containing extremely sensitive data, such as credit card information, private customer information, financial records, human resources data, or really anything that could be sensitive to your business.

It might be legally sensitive, or it might just be the secret sauce that gives your business its competitive advantage. Whatever the reason, access to that data must be controlled, and firewalls are often a good mechanism for controlling access.

Figure 19 shows an example of an internal security zone. In this picture, I show three parallel internal security zones for different applications.




Figure 19 - Internal firewall topology

I like to deploy internal firewalls in a way that's a little different from perimeter firewalls. For perimeter firewalls, I like to have VPN and IDS/IPS inspection features. These features are important when facing the hostile public Internet.

On the other hand, the important thing for an internal firewall is simply controlling access to the applications and servers. So internal firewalls don't need the same web-tier infrastructure design as an Internet DMZ.

Instead, for internal firewall segregated security zones, I like to build the zones in parallel on the same set of switches. If the server environment is virtualized, then I often virtualize the firewalls too. If the firewalls are virtualized, then I give every separate parallel application zone its own separate virtual firewall.

If the server environments are physical, then I generally prefer to use a single common firewall with one interface on the corporate side and a separate interface in each internal zone.

I tend to avoid using a single common physical firewall, whether virtualized or physical, when I need a lot of throughput into and out of the application zones.

For example, if the amount of network traffic reaching the application inside such a zone is so high that it would require an extremely expensive firewall, then an internal firewall design won't work well. In that case, it might make sense to re-engineer the application so it has a web front-end and put the heavy traffic on an application back-end behind the web server. Then the internal security zone model works very well again.

Internal firewall topology is also not useful when there's a lot of traffic among the different secure application zones. If your applications exchange huge amounts of data, they probably don't want to be separated. In that, case it might make sense to put them all in a single secure zone.

## An internal firewall topology is useful for

- Protecting sensitive data or applications that are internal to your network, intended for internal consumption, but where the data is so sensitive that it requires special access controls. (The "inside" of the firewall is the secure zone, and the "outside" is the corporate network.)
- Separating software development and test environments from your production data. (The "inside" or secure part of the network is the corporate network and the "outside" is the development or test zone.)

## Avoid when

- You need more throughput into and out of the application zones.
- There's a lot of traffic among the different secure application zones.

# Helpful Resources

---

In this short guide, we've covered the basics of network design and documentation. If you want to dig deeper, here are some recommended resources.

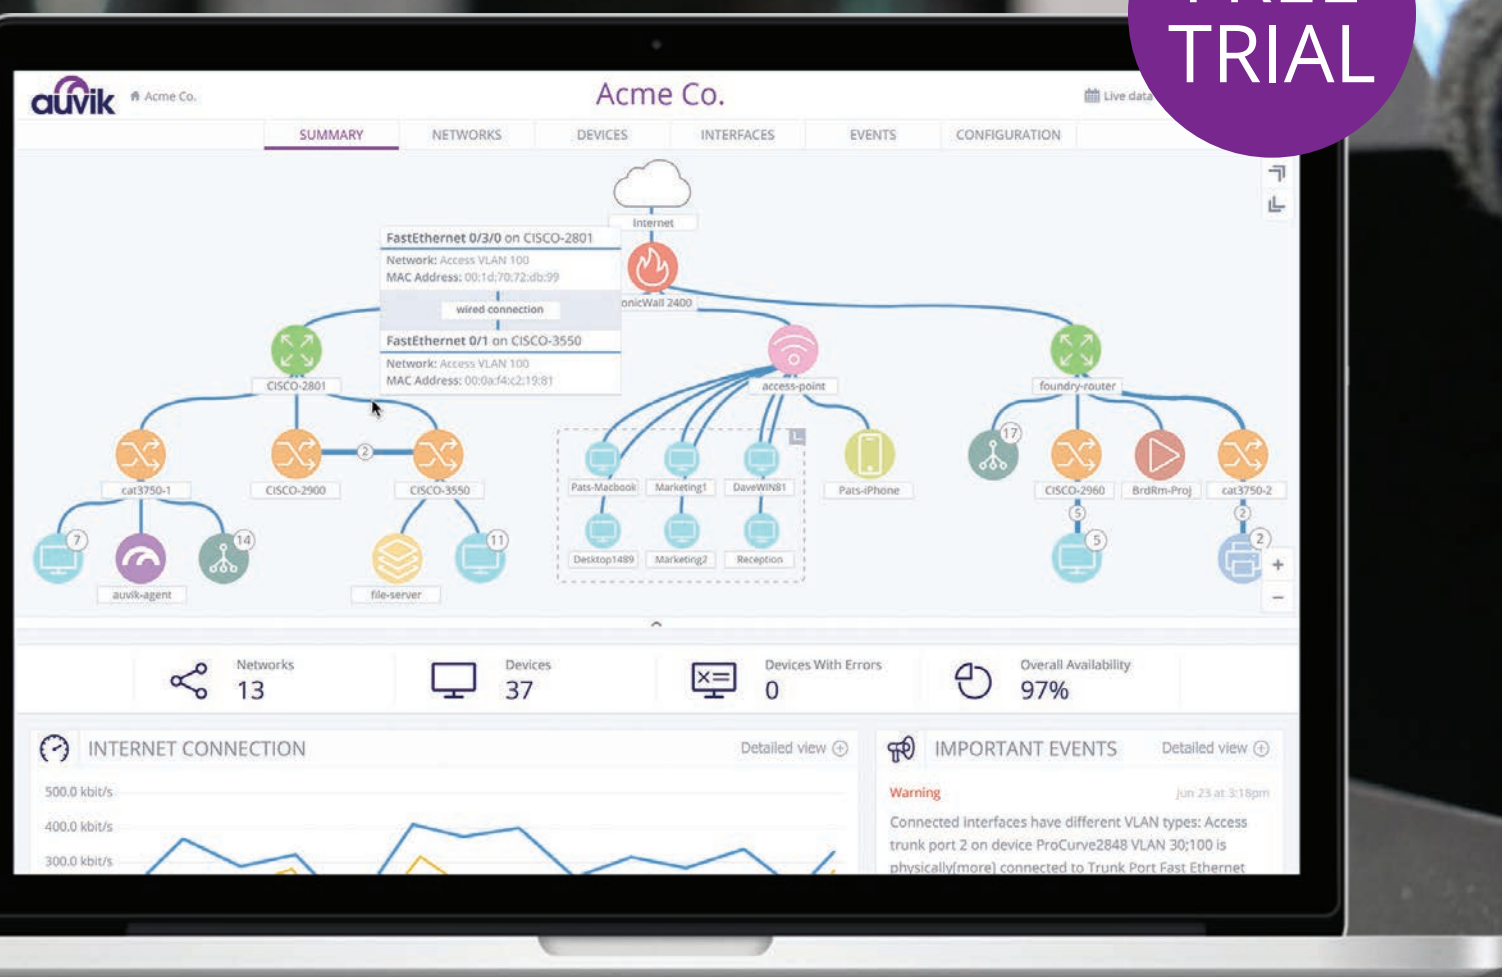
- An excellent discussion of the theory of TCP/IP can be found in the first volume of *TCP/IP Illustrated (2nd edition)* by W. Richard Stevens. The book was published in 1993 so it's a little dated, but it does a great job of explaining how the protocols work.
- For network topology and network design, there are two books that I like to recommend. The first is *Top-Down Network Design (3rd edition)* by Priscilla Oppenheimer. The second is my own book, *Designing Large-Scale LANs*.
- For larger and more complicated networks, you'll want to dig into the theory and implementation of routing protocols. For this, I really like *Routing TCP/IP (2nd edition)* by Jeff Doyle and Jennifer Carroll. The two-volume set is a tad Cisco-centric but I've never encountered a better book for explaining routing protocols.

# SEE



- Your network topology in minutes
- Every config you're running
- Real-time performance stats

FREE TRIAL



<https://www.auvik.com>

Auvik is a cloud-based software-as-a-service (SaaS) application that makes it dramatically easier for small and mid-sized businesses to build and manage their IT networks.

Follow us

