1

Carrier Networks

We have come a long way in a short time. Instant communication arrived relatively recently in the history of man, with the invention of the telegraph at the beginning of the nineteenth century. It took three quarters of a century before we saw the first major improvement in mass communication, with the arrival of the telephone in 1874, and another half century before a national network and transcontinental communication became common in the US. But what was a middle class convenience years ago is now a common necessity. Today, our worldwide communication Network is a model of egalitarian success, reaching into the enclaves of the wealthy and the street vendors in villages of the developing world with equal ease. Remarkably it remains largely in the hands of the private sector, and is held together and prospers through forces of cooperation and competition that have served society well.

The Network is made up of literally millions of nodes and billions of connections, yet when we choose to make a call across continents or browse a web site in cyberspace we just expect it to work. I use the proper noun Network when referring to the global highway of all interconnection carrier networks, such as Nippon Telephone and Telegraph (NTT), British Telecom (BT), China Telecom, AT&T, Verizon, Deutsche Telekom, Orange, Hurricane Electric, and many others, just as we use the proper noun Internet to refer to the global public Internet Protocol (IP) network. The Internet rides on the Network. If the Network were to fail, even within a city, that city would come to a halt. Instead of purchasing fuel at the pump with a credit card, drivers would line up at the register while the attendant tried to remember how to make change instead of waiting for the Network to verify a credit card. Large discount retail outlets would have their rows of registers stop and for all practical purposes the retailers would close their doors. Alarm systems and 911 emergency services would cease to function. Streets would become congested because traffic lights would no longer be synchronized. So what are the mechanisms that keep the Network functioning 24 hours a day with virtually no failures?

1.1 Operating Global Networks

The global nature of networks is a seismic change in the history of modern man. In the regulated world of the past franchise carriers completely dominated their national networks.

Global Networks: Engineering, Operations and Design, First Edition. G. Keith Cambron.

Interconnection among networks existed for decades, but carriers did not over build each other in franchise areas. That all changed in the latter decades of the twentieth century as regulation encouraged competition and data services emerged. International commerce and the rise of multinational companies created a demand for global networks operated by individual carriers. Multinational companies wanted a single operator to be held accountable for service worldwide. Many of them simply did not want to be in the global communications business and wanted a global carrier to sort through interconnection and operations issues inherent in far reaching networks.

In parallel with globalization was the move to the IP. The lower layers of the Open System Interconnection (OSI) protocol stack grew because of global scale, and upper layer complexity; the complexity increased with new services such as mobility, video, and the electronic market, largely spurred by Internet services and technology. Operators were forced to reexamine engineering and operating models to meet global growth and expanding service demand. Before deregulation reliability and predictability were achieved through international standards organizations, large operating forces, and highly structured and process centric management regimes. Deregulation, competition, global growth, and service expansion meant that model was no longer economic and could not respond to the rapid introduction of new services and dramatic growths in traffic.

Operating models changed by applying the very advances in technology which drove demand. Reliable networks were realized by reducing the number of failures, by shortening the time for repair, or both. In the old model central offices were staffed with technicians that could respond on short notice to failures, keeping restoral times low. In the new model networks are highly redundant, well instrumented, constantly monitored, and serviced by a mobile work force.

1.1.1 The Power of Redundancy

This section introduces the foundation of global network reliability, redundancy using a simple systems model.

1.1.1.1 Simplex Systems

In the model following, a subscriber at A sends information i_0 to a subscriber at B. The information arrives at B via a communications system S_0 as i_1 after a delay of t (see Figure 1.1).

Subscribers care about two things, the fidelity of the information transfer and transmission time. Fidelity means that the information received, i_1, should be as indistinguishable from the information sent, i_0, as possible. If we assume for simplicity that our communication depends on a single system, S_0, that fails on average once every year, and it takes 4 h to

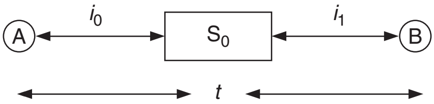

Figure 1.1 Simplex operation.

get a technician on site and restore service, the service will be down for 4 h each year on average, yielding a probability of failure of 4.6×10^{-5}, or an availability of 99.954%. That means we expect to fail about once for every 2000 attempts. For most communications services that is a satisfactory success rate.

But real world connections are composed of a string of systems working in line, possibly in the hundreds, any one of which can fail and impede the transfer. For a linear connection of 100 such systems, our failure probability grows to 4.5×10^{-3} and availability drops to 95.5%. Approximately 1 in 20 attempts to use the system will fail.

1.1.1.2 Redundant Systems

The chances of success can be dramatically improved by using a redundant or duplex system design, shown in Figure 1.2.

In the design two identical systems, S_0 and S_1 are each capable of performing the transfer. One is active and the other is on standby. Since only one system affects the transfer, some communication is needed between the systems and a higher authority is needed to decide which path is taken.

In the duplex system design the probability of failure drops to 2.1×10^{-5} for 100 systems in line, an improvement of more than 100× for an investment of 2×. Availability rises to 99.998%. We expect to fail only once in each 50 000 attempts.

Implicit in the model are some key assumptions.

- Failures are random and non-correlated. That is the probability of a failure in S_1 is unrelated to any failure experienced by S_0. Since it's likely the designs of the two systems are identical, that assumption may be suspect.
- The intelligence needed to switch reliably and timely between the two systems is fail-safe.
- When S_0 fails, Operations will recognize it and take action to repair the system within our 4 h timeframe.

1.1.1.3 Redundant Networks

Redundancy works within network systems; their designs have two of everything essential to system health: power supplies, processors, memory, and network fabric. Adopting reliable network systems doesn't necessarily mean networks are reliable. Network systems have to be connected with each other over geographical expanses bridged by physical facilities to build serviceable networks. Physical facilities, optical fiber, telephone cable, and coaxial cable are exposed to the mischiefs of man and of nature. Dual geographically diverse routes

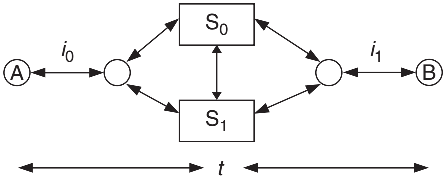

Figure 1.2 Duplex model.

to identical network systems preserve service if the end nodes recognize that one route has failed and the other is viable. Global networks rely on redundant systems within redundant networks. The combination is resilient and robust, providing any failure is recognized early and maintenance is timely and thorough.

The next sections explore this foundational model in more depth in an attempt to understand how it works, and how it can break down in real networks.

1.1.2 The Virtuous Cycle

In the 1956 film *Forbidden Planet*, an advanced civilization called the Krell invents a factory that maintains and repairs itself automatically. In the movie, although the Krell are long extinct, the factory lives on, perpetually restoring and repairing itself. Some academics and equipment suppliers promote this idea today using the moniker "self-healing network." An Internet search with that exact phrase yields 96 000 entries in the result; it is a popular idea indeed. Academic papers stress mathematics, graphs, and simulations in search of elegant proofs of the concept. Yet real networks that perform at the top of their class do so because of the way *people* design, operate, and manage the technology. It is the blend of systems, operations, and engineering that determine success or failure. *Systems and people* make the difference. Figure 1.3 illustrates the Virtuous Cycle of equipment failure, identification, and restoral.

The cycle begins at the top, or 12 o'clock, where the Network is operating in duplex that is full redundancy with primary and alternate paths and processes. Moving in a clockwise direction, a failure occurs signified by the X, and the Network moves from duplex to

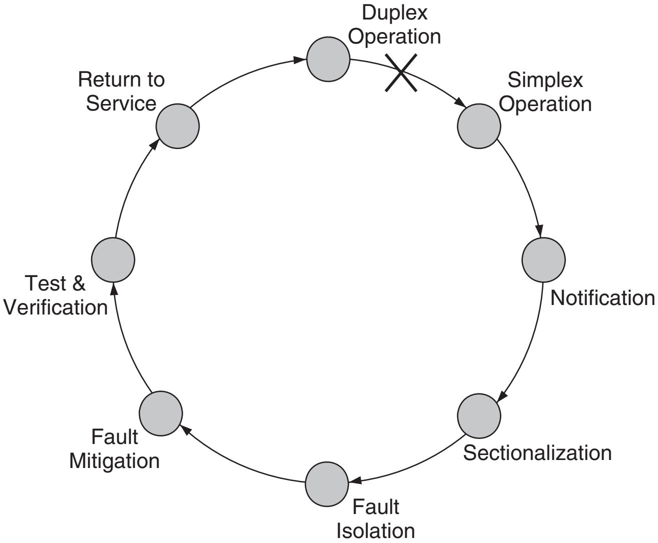

Figure 1.3 Virtuous Cycle.

simplex operation, although no traffic is affected. While the Network is operating in simplex it is vulnerable to a second failure. But before operators can fix the problem they need to recognize it. Notification is the process whereby network elements send alarm notifications to surveillance systems that alert network operators to the situation. Notifications seldom contain sufficient information to resolve the problem, and in many situations multiple notifications are generated from a single fault. Operators must sort out the relevant notifications and sectionalize the fault to a specific network element in a specific location. The failed element can then be put in a test status, enabling operators to run diagnostics and find the root cause of the failure. Hardware faults are mitigated by replacing failed circuit packs. Software faults may require a change of configuration or parameters, or restarting processes. Systems can then be tested and operation verified before the system is restored to service, and the network returns to duplex operation. Later chapters explore these steps in detail.

1.1.3 Measurement and Accountability

The Virtuous Cycle enables highly trained *people* to work with *network systems* to restore complex networks quickly and reliably when they fail. But it does nothing to insure the network has sufficient capacity to handle demands placed upon it. It does not by itself give us any assurance the service the customer is receiving meets a reasonable standard. We can’t even be sure network technicians are following the Virtuous Cycle diligently and restoring networks promptly. To meet these goals a broader system of measurements and accountability are needed. Carrier networks are only as good as the measurement systems and the direct line of measurements to accountable individuals. This is not true in smaller networks; when I speak with Information Technology (IT) and network engineers in smaller organizations they view carriers as having unwarranted overhead, rules, and systems. In small networks a few individuals have many roles and are in contact with the network daily. They see relationships and causality among systems quickly; they recognize bottlenecks and errors because of their daily touches on the network systems. Such a model does not scale. Carrier networks have hundreds of types of systems and tens of thousands of network elements. Carrier networks are more akin to Henry Ford’s production lines than they are to Orville’s and Wilbur’s bicycle shop. Quality and reliability are achieved by scaling measurement and accountability in the following ways.

- **End service objectives** – identify measurable properties of each service; commit to service standards, communicate them, and put them into practice.
- **Network systems measurement** – using service objectives analyze each network and network element and set measurable objectives that are consistent with the end to end service standard.
- **Assign work group responsibility** – identify which work group is responsible for meeting each of the objectives and work with them to understand how they are organized, what skills they have and what groups they depend upon and communicate with regularly.
- **Design engineering and management systems** – systems should support people, not the other way round. Find out what systems the teams already use and build on those if at all possible. Don’t grow YAMS (yet another management system).

1.2 Engineering Global Networks

Changes in operations as dramatic as they have been are greater yet for the design and engineering of global networks. Carriers in the US prior to 1982 were part of a vertically integrated company, AT&T® or as it was commonly known, the Bell System. AT&T General Departments operated the complete supply and operations chain for the US telecommunications industry. Wholly owned subsidiaries planned, designed, manufactured, installed, and operated the network. AT&T's integrated business included the operations support, billing, and business systems as well. Carriers (the operating companies) had no responsibility for equipment design or selection, and limited responsibility for network design. Today carriers have full responsibility for planning, designing, installing, and operating their networks. They also have a direct hand in determining the functionality and high level design of network systems, and operations and business systems. The sections that follow summarize responsibilities of carrier engineering departments.

1.2.1 Architecture

High level technology choices are the responsibility of Engineering. Engineering architects analyze competing technologies, topologies, and functional delegation to determine the merits and high level cost of each. Standards organizations such as ITU, IETF, and IEEE are forums serving to advance ideas and alternatives. Suppliers naturally promote new ideas and initiatives as well, but from their point of view. Long range plans generally describe the evolution of networks but may not address practical and necessary design and operational transition issues.

1.2.2 Systems Engineering

A wide range of responsibilities rest with systems engineers. They begin with high level architectural plans and translate them into detailed specifications for networks and for the individual network elements. Equipment recommendations, testing, certification, and integration are all performed by these engineers. Operational support, IT integration, and network design are performed by systems engineers as well.

1.2.3 Capacity Management

There are four general ways in which network capacity is expanded. Each is described in the following.

1.2.3.1 Infrastructure Projects

Periodically major network augmentation is undertaken for a variety of reasons.

- Expansion into a new geography is a common trigger. A country adopts competitive rules that enable over building the incumbent.

- Technology obsolescence, such as the shift from Frame Relay to IP networks leads to a phased introduction of new technology. The new networks often must interwork with the legacy technology making the transition more challenging.
- Carrier mergers or acquisitions are followed by network rationalization and integration. The numbers and types of network elements are winnowed out to ease operational demands.
- New lines of business, such as Internet Protocol Television (IPTV) or content distribution, place new demands on the network requiring specialized technology design and deployment.

1.2.3.2 Customer Wins

Major customer contract wins significantly increase demand at large customer locations, rendering the existing capacity inadequate. Sometimes outsourcing of a Fortune 500 company network can be accompanied by an agreement to transfer their entire network, employees, and systems to the winning carrier. If they are of sufficient scope, the accompanying network augmentations are treated as separate projects with dedicated engineering, operations, and finance teams.

1.2.3.3 Capacity Augmentation

By far the most common reason for adding equipment and facilities to a network is the continuous growth in demand of existing services and transport. For decades voice traffic grew at a rate of about 4% each year. Data traffic and IP traffic specifically, have grown at an annual rate of 30–50% for over three decades. With tens of thousands of network systems and millions of facilities, automating demand tracking and capacity management is one of the most resource intensive jobs in engineering.

1.2.3.4 Network Redesign

This is the most neglected, and often the most valuable tool available to network engineers. The demand mechanisms cited above are all triggered by events. Capacity augmentation, the most common engineering activity, is triggered when a facility or network element falls below a performance threshold, such as packet discards or blocked calls. Network engineers generally look at those links exceeding the accepted levels and order augmentation or resizing. If a node nears exhaust, either because of port exhaust or throughput limits, engineers order a new node and rearrange the traffic between that node and adjacent ones. In effect they limit the problem and the solution space to a very narrow area, the particular link or node that exceeded a threshold.

Network redesign broadens the scope to an entire region or community. It is performed by systems engineers, not network engineers. It begins with A-Z (A to Z) traffic demand and uses existing topology, link, and element traffic loads as an informational starting point, not as a constraint. In voice networks Call Detail Records (CDRs) are a starting point since they have the calling party (A) and the called party (Z). In IP networks netflow data, coupled with routing information yield the necessary A-Z matrices. Redesigns are performed far too infrequently and the results often reveal dramatic changes in traffic patterns that no one recognized. Express routes, bypassing overloaded network elements, elimination of elements, and rehoming often result in dramatic savings and performance improvements.

1.3 Network Taxonomy

To better understand network operations and engineering some grounding in networks and systems is needed. Networks are best described as communications pathways that have both horizontal and vertical dimensions. The horizontal dimension encompasses the different types of networks which, when operated in collaboration deliver end to end services. The vertical dimension is two tiered. Network elements, which carry user information and critical signaling information, are loosely organized around the OSI seven-layer model, one of the most successful design models in the last 50 years. As a word of warning, I use the terms *network system* and *network element* interchangeably. *Network system* was in wide use when I joined Bell Telephone Laboratories in the 1970s. *Network element* evolved in the 1990s and is institutionalized in the 1996 Telecommunications Act.

Above the network tier is a set of management systems that look after the health, performance, and engineering of the network tier.

The distinction between network and management systems is almost universally a clear line, as shown in Figure 1.4. Tests used to distinguish between the two systems types are based on how transactions are carried.

1.3.1 Voice Systems

In the first half of the twentieth century transactions meant one thing, a wireline phone call. A wireline call has six distinct stages.

1. The first stage is the service request. For a manual station set, that simply means taking the receiver off the switch hook and listening for dial tone.
2. The second stage is the signaling stage in which the originator dials the called party's number.

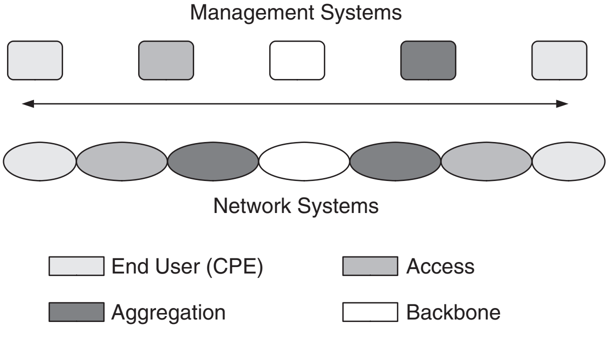

Figure 1.4 Network and management systems.

3. In the third stage, call setup, a signaling connection is established between the originator and the called party and a talking path is reserved.
4. Alerting or ringing the two parties takes place in the fourth stage. Audible ringing, sometimes called ringback, is applied to inform the calling party that the call has progressed. Power ringing causes the called station to emit an audible ringing sound alerting the called party.
5. For successful calls, where the called party answers, the fifth stage is the completion of the final leg of a stable two way talking path.
6. In the sixth and last stage, the two parties conclude the call and hang up, after which the end to end path is freed for other users.

These six stages are the same whether a call is originated or terminated by a human or a machine. A wide range of technologies has been used over the years in each stage, but the stages are more or less constant.

For voice services we can then distinguish among systems by applying the following tests:

- Does the system perform a critical function in one of the six stages of call processing?
- If the system is out of service, can existing subscribers continue to place and receive calls?

Network systems when tested yield a yes to the first test and a no to the second. The time frame for applying the tests is important; a reasonable boundary for applying these tests is an hour. A local switching system is the one that gives dial tone and rings wireline phones. If it fails, the effects are immediate. At a minimum no new originations or completions occur because dial tone and ringing are not provided. In severe cases, calls in progress are cut off.

A provisioning system is a counter example. That system is responsible for adding new customers, removing customers, and making changes to existing customers' services. It does not perform a critical function in any of the six stages of call processing. If the provisioning system fails, we simply can't modify a customer's service attributes until the provisioning system returns to service. Existing calls, new originations, and terminations are not affected, so the provisioning system is a management system, not a network system. A second example is billing systems. If a billing system fails on a voice switching system, calls are completed without charge. Unfortunately no one sounds a siren or sends a tweet to let users know the billing system has failed and you can make free calls. The design choice to let calls go free during billing system failure is a calculated economic decision. It is cheaper to let them go free than it is to design billing systems to network system standards. Occasionally losing some revenue is cheaper than building redundant fault tolerant recording everywhere.

But what about the power systems in buildings where communications systems are located? In general network systems operate directly off of DC batteries which are in turn charged by a combination of AC systems and rectifiers. These hybrid power systems are engineered to survive 4–8 hours when commercial AC power is lost. Most central offices have back up diesel generators as well, enabling continuous operation indefinitely, assuming the fuel supply is replenished. Cooling systems fall into the same category. These are systems that do not affect the six stages of voice network systems if they remain failed for an hour. So here is a class of systems that if failed, don't affect calls within our hour time frame, but can affect them after a few hours or possibly days, depending on the season. These systems are in a third category, common systems. This is an eclectic group, covering power, cooling, humidity, door alarms, and other systems that if failed, can imperil the network systems within hours under the wrong circumstances.

1.3.2 Data Systems

The original tiered distinction and design for network and management systems came from the wireline voice network, but it applies to data and mobile networks as well. Consider two common data services upon which we can form our definitions, Internet browsing and mobile texting, or Short Messaging Service (SMS). Browsing is generally performed by a subscriber accessing content on the Internet by sending requests to a set of servers at a web site. The subscriber unconsciously judges the service by the response time, measured from the time the return key is stroked until the screen paints with the response. In the case of SMS, the subscriber has no clear way of knowing when the message is delivered, or if it is delivered at all. However, if a dialog between two SMS subscribers is underway, a slow response or dropped message is likely to be noticed.

For mobile subscribers, many of the network systems that carry Internet service and SMS are common. Our criterion for distinguishing between network and management systems is set by the most demanding data service. Before the introduction of 4G mobile services under the banner of LTE, Long Term Evolution, Internet access was the most demanding service. But LTE, unlike prior mobile technologies, uses Voice over Internet Protocol (VoIP) for voice service. With LTE data (VoIP data) delay tolerances become more unforgiving.

For data systems we can use our voice tests to distinguish among systems by applying the same tests, with minor modifications:

- Does the system perform a critical function in the timely delivery of subscriber data?
- If the system is out of service, can existing subscribers continue to send and receive data?

The modifier timely was added to the first test. While it was not included in the comparable test for voice service, it was implied. Recalling the six steps of call processing, untimely delivery of any of the functions is tantamount to failure. If you pick up a wireline receiver and have to wait over 10 seconds for dial tone, it's likely one of two things will occur. If you're listening for a dial tone you may grow impatient and just hang up and try again. If you don't listen and just begin dialing, believing the network is ready, you'll either be routed to a recording telling you the call has failed, or you'll get a wrong number. Consider the case of not listening for dial tone before dialing your friend whose number is 679−1148. You could be in for a surprise. Suppose you fail to listen for dial tone and begin dialing. If dial tone is delivered after the 7, the first three digits the switching system records are 911. Now you will have an unplanned conversation with the Public Service Answering Point (PSAP) dispatcher. When Trimline®[1] phones were first introduced by AT&T they caused a rise in these wrong number events. Trimline was among the first station sets to place the dial in the handset and people did not immediately become accustomed to putting the

[1] Trimline is a registered trademark of AT&T.

phone to their ear, listening to dial tone, and then pulling the phone down to dial. Many just picked up the phone and began to dial. Eventually they learned. Users can be trained.

1.3.3 Networks

Our communication infrastructure is actually a network of networks. I mean that in two senses. Networks are different in *kind*, and different by *serving area*.

To say networks differ in kind is an economic and technical distinction. Networks evolve to perform specific roles and the economics, topology, and service demands determine the technologies used and the associated designs. So, local distribution networks that deliver broadband and phone service to homes look far different than backbone networks carrying Internet traffic between major peering points.

Residential distribution networks, whether they are designed by cable providers or telcos tend to be asymmetrical, delivering more bandwidth toward the home, and they are very sensitive economically. If you have to deliver that service to 30 million homes, a $10 saving for each home matters.

Core IP networks carrying petabytes of traffic are at the other end of the technology and economic spectra. They are symmetrical networks that are fully redundant and possess sophisticated mechanisms for rerouting traffic in less than a second in the event of failure. A failure in the core affects all customers and has severe economic impact. Spending more in the core makes sense.

While the first distinction is according to kind, the second is by provider serving area. Each service provider designs, builds, and operates their network, generally under a franchise of the local communications regulator. The goal of universal service was established as U.S. Policy in the Communications Act of 1934 [1]. Two cornerstones of the act were that service should be extended to everyone, and that competing carriers should interconnect, enabling a national network of independent carriers. Prior to regulation in the twentieth century competing carriers often refused to interconnect. After 1934 interconnection and cooperation became common practice in the industry. It naturally extended to the Internet, although the U.S. Communications Act of 1934 does not directly apply to that network.

International cooperation and carrier interconnection are remarkable and beneficial practices that emerged from our twentieth century industrial society. Railroads in that era by comparison are a different story. Different gauges continued well into the twentieth century in Europe [2], inhibiting travel and commerce. When travelers reached an international border they often disembarked from one train and loaded aboard a different train because of the differences in railroad gauges. We take for granted our ability to place a call anywhere in the world, access any Internet site, and send e-mail and text messages to anyone anywhere. Interconnection only becomes news when it is taken away, as some Middle Eastern countries experienced in the Arab Spring uprisings. We'll explore network interconnection in depth in a later chapter.

1.3.4 Network Systems

Network systems support the Virtuous Cycle in the following ways.

- **Duplex operation** – achieving non-stop processing in the face of internal and external failures is founded upon redundant operation, and it requires the following functionality.

  - **Field replaceable units (FRUs)** – failures occur on components with active electronics and components with software. Active electronic components are housed in circuit packs on assemblies that can be replaced in the field by technicians.
  - **Hot swappable cards** – this takes FRUs one step further. Critical cards, such as system controllers and network fabric cards must be able to be removed and inserted while the system is in service, without affecting system performance. This is far more difficult to design than one might think. We will explore the necessary design steps.
  - **Fault detection and spare switching** – systems must detect faults and re-route traffic to functioning cards, once the card is ready to process the load. Craft should also be able to return traffic to the primary card once it has been verified as restored.
- **Notification and sectionalization** – effective implementation is contingent upon trouble identification, correlation, and communication. Identification is achieved through hardware and software monitoring of internal hardware and software systems, and external circuits. Off-normal conditions must be resolved in a hierarchical structure that reports the most likely cause to the upper layers. Sectionalization resolves off-normal conditions as local or far-end related.
- **Fault isolation and mitigation** – fault detection and re-routing must happen automatically in milliseconds under the control of the system. However isolation and mitigation are performed by craft and rely on the software management of boards and ports. Object management is closely tied to the notification process and the implementation of administrative and operational state management.
- **Test, verification, and return to service** – object management again plays a key role, coupled with on board diagnostics.

The following chapters describe how network hardware, software, management systems, and work group standards and practices make this cycle successful.

1.4 Summary

Networks expanded across the entire globe from humble beginnings 100 years ago. Today's networks interconnect billions of people with highly reliable voice, Internet, and video services around the clock. Redundancy, fault tolerant systems, and operators and management systems working together in the Virtuous Cycle detect and resolve failures before they affect customer service. Network engineers monitor network elements and links, adding capacity, and augmenting the networks on a daily basis. Systems engineers extend networks with new technology and certify technology for seamless introduction into live networks. Network systems can be upgraded and repaired in the field without service interruptions. Hardware and software designs are defensive, meant to continue to operate in the face of failures and traffic overload. As we'll see in the next chapter, network systems are designed to very different standards than support systems.

References

1. United States Federal Law Enacted as Public Law Number 416, Act of June 19 (1993).
2. Siddall, W.R. (1969) Railroad gauges and spatial interaction. *Geographical Review*, **59** (1), 29–57.