Excerpts taken from:

Network Troubleshooting
By Othmar Kyas

An Agilent Technologies Publication

Section I

Basic Concepts

Chapter 1

Network Availability

1 Network Availability

"Waiting for an alarm is not the ideal form of network management."

BOB BUCHANAN, THE NETWORK JOURNAL

1.1 The Strategic Importance of Information Technology

Growing financial and competitive pressures in the business world mean that companies everywhere must continuously optimize their internal and external structures in order to survive. All business processes and routines must be reviewed regularly for effectiveness ("Are we doing the right things?") and for efficiency ("Are we doing things right?"). Most business processes today consist of physical activities, such as the manufacture of a metal part, combined with information flow: How many parts should be produced? When? In what sizes? Increasingly, key business processes—in insurance companies, travel agencies, banks, and airlines, for example—consist entirely of information flow. Today, of course, the flow of information is largely dependent on information technology—that is, computers, databases and networks. A high-performance, high-availability information technology (IT) system is becoming a prerequisite for successful execution of the business practices that are decisive in maintaining a leadership position in today's competitive markets. The role of the computer in business has changed radically over the past few years, from a tolerated plaything to a cornerstone of corporate infrastructures. This change has taken place so rapidly that in many companies IT still has not taken a central position in managerial circles, even though it has long since become indispensable for day-to-day business functions.

The reliance of enterprises on smoothly functioning IT infrastructures will continue to grow in the coming years. Areas of business that until recently had little to do with computer technology, such as marketing and customer service, are increasingly IT-based. This is largely due to the advent of customer interfaces that allow consumers to perform many transactions electronically, such as placing orders or making reservations. In fact, the proportion of people who work directly or indirectly with IT has grown in recent years to more than 50 percent (see Figure 1.1).

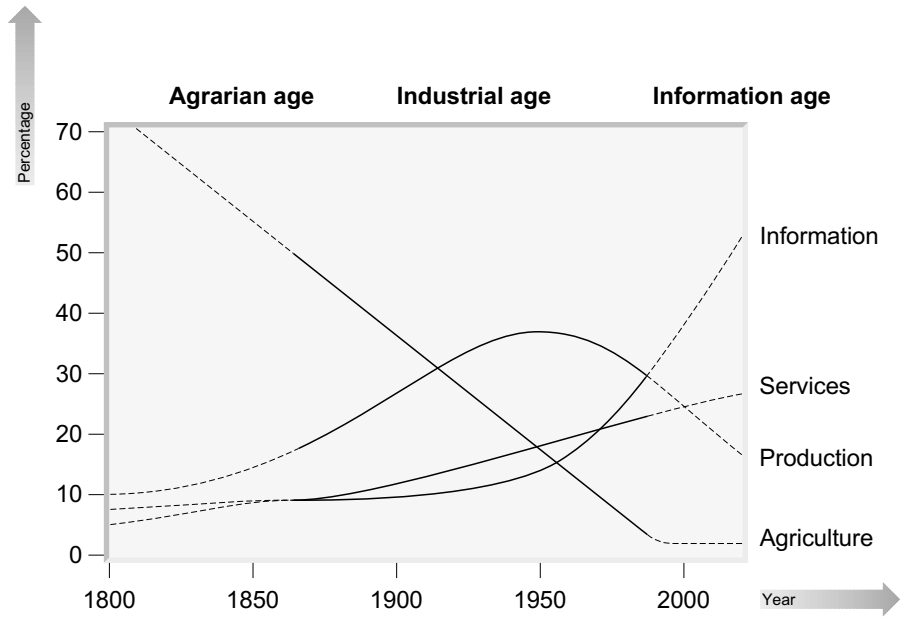

Source: Leo a. Nefiodow, "The fifth Kondratieff"

Figure 1.1 Changes in employment patterns in Western industrial countries since 1800

Industry	Business processes	Downtime costs per hour (US$)
Financial services	Stock trading	6,000,000
Financial services	Credit card/telecash transactions	2,400,000
Media	Pay-per-view	150,000
Retail	Home shopping (TV)	100,000
Retail	Mail order (catalog)	80,000
Travel/tourism	Airline reservation	82,500
Shipping	Parcel service	25,500

Source: AT&T/Gartner Group

Figure 1.2 Costs of network failure

In keeping with these developments, the professional operation and management of computer networks has long since ceased to be a necessary evil. On the contrary, it has become a decisive strategic necessity for the success of almost any enterprise. A network failure that lasts only a few hours can cost millions of dollars. According to a study carried out by AT&T, companies that deal in financial services, such as investment brokerages or credit card firms, can suffer losses of 2.5 to 5 million dollars from just 1 hour of network downtime (see Figure 1.2).

1.2 Intranets and the Internet: Revolutions in Network Technology

The difficulties involved in the professional operation of high-performance data networks have been further complicated by the Internet revolution, which has brought about radical changes in network technology and applications. Since the mid-1990s the Internet has not only developed into a universal communications medium, but has also become a global marketplace for the exchange of goods and services. As a result, growing numbers of business are faced with the necessity of providing their employees with Internet access. Special network infrastructures are now required in order to provide electronic access for increasing numbers of Internet-based consumers. Once it was sufficient to have just a few carefully controlled wide-area network (WAN) links in an otherwise homogenous local-area network (LAN). Today, however, a secure, high-performance LAN-WAN structure is indispensable.

Internet technologies are also being introduced into company networks, leading to the development of “corporate intranets”. This has necessitated further restructuring so that broad areas of internal data processing can be adapted to the transport mechanisms, protocols and formats used in the Internet. All of these developments have caused the World Wide Web (WWW) to take on a position of global importance as a uniform user interface. At the same time, these changes have placed enormous demands on network managers. In many cases, the skills and tools available for managing computer systems and networks can barely keep up with the increasing complexity of data network structures. And to add to the difficulty of the task, the technology cycles in data communications—the intervals at which new and more powerful data communication technologies are introduced—are getting shorter all the time. Whereas the classic 10 Mbit/s Ethernet topologies shaped computer networking throughout the 1980s, the 1990s have seen the introduction of new technologies almost every year, including LAN switching, 100 Mbit/s Ethernet, Gigabit Ethernet, ATM (Asynchronous

Transfer Mode), IP (Internet Protocol) switching, Packet over Sonet (PoS) and ADSL (Asynchronous Digital Subscriber Mode), to name just a few. Product life cycles in the IT field are often measured in months now rather than years. This rapid pace of technological development puts manufacturers and users alike under tremendous pressure to keep abreast of constant innovation (see Figure 1.3).

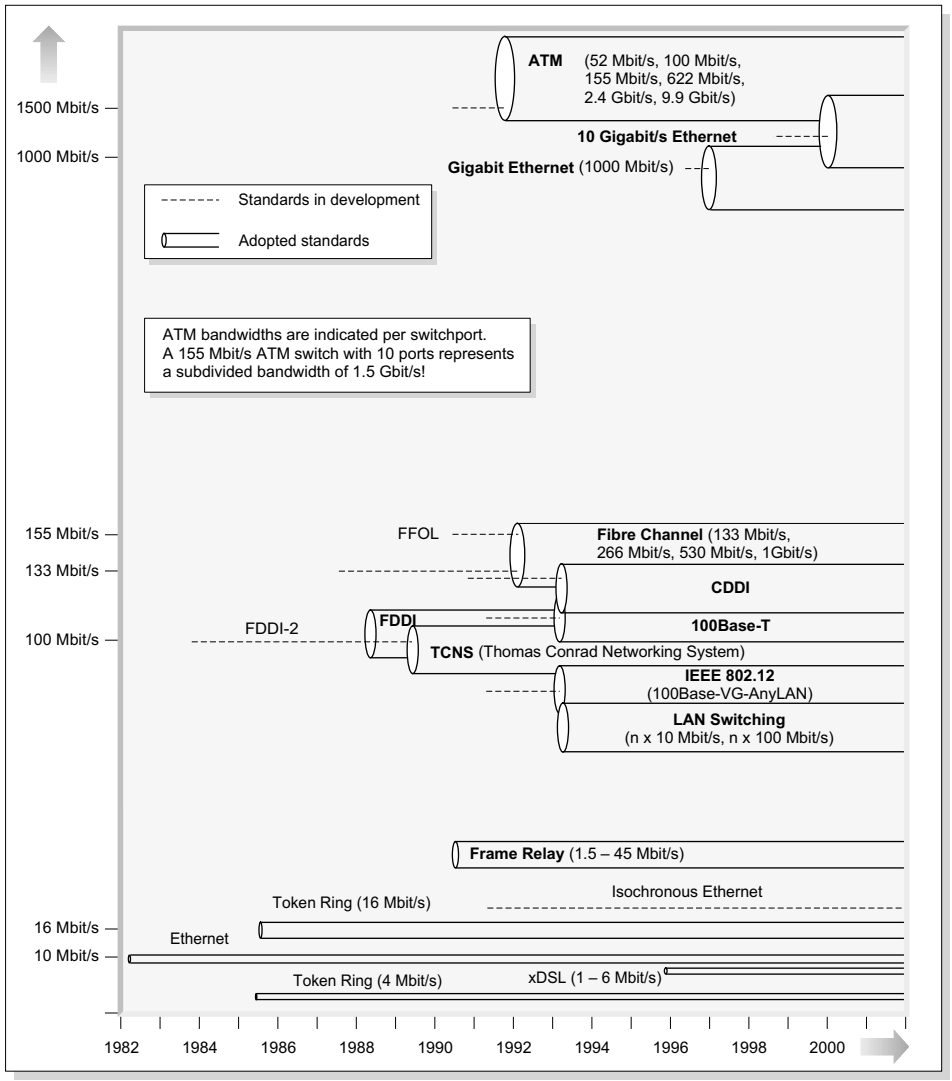

Figure 1.3 The development of data transmission technologies: 1980 to 2000

1.3 The Behavior of Complex Network Systems: Catastrophe Theory

The enormous technological complexity in combination with the large numbers of hardware and software components used in networks makes operation and management a difficult task, to say the least. Communication media, connectors, hubs, switches, repeaters, network interface cards, operating systems, data protocols, driver software, and application software must all function smoothly under widely varying conditions, including network load, number of nodes connected, and size of data packets transmitted. Even when a given system has attained a relatively stable operating state, its stability is constantly put to the test by dynamic variations as well as by operator errors, administrative errors, configuration changes, and hardware and software problems. In general, the more complex a system is and the greater the number of parameters that influence it, the more difficult it is to predict its behavior. Catastrophe theory (see René Thom, 1975) offers an excellent model for describing the behavior of systems as complex as computer networks. This theory can provide at least qualitative descriptions of system behavior, especially for non-linear operating states, such as those that often accompany a network breakdown. Catastrophe theory postulates seven elementary catastrophes, which behave in a given manner according to the *number* rather than the *type* of control parameters influencing the system. The behavior model for catastrophes determined by two parameters, for example, is called a cusp graph. The cusp is a three-dimensional surface whose upper side represents balanced states, while the lower surface represents unstable maxima. Catastrophe theory can be applied to Ethernet networks, for example, to show the effects of two control parameters, slot time and network load, on throughput. Slot time, which is defined as twice the time it takes a signal to travel between the two nodes that are farthest apart in an Ethernet segment, is influenced by the network components that cause signal transmission delay or latency, such as cables, repeaters or hubs. Figure 1.4 shows the behavior pattern for throughput when all other variables, such as network load, average packet size and number of network nodes, are constant.

An increase in traffic in a network with a given slot time a moves the operating state across the upper surface of the cusp. The rise along the x-axis indicates increasing throughput. Starting from the higher slot time b, however, the same increase in traffic drastically reduces network efficiency. All processes take place on the surface of the cusp and are thus linear. If the operating state is a when the network load increases, and subsequently the slot time increases from state c (Figure 1.5), an abrupt departure from the balanced state takes place at point d in order to arrive directly at point e. Point e represents a stable operating

state, but one in which throughput is minimal. The abrupt transition from *d* to *e* constitutes a catastrophe.

The model provided by catastrophe theory clearly illustrates how complex and unpredictable a network can be. The symptoms that indicate problems in a network are often caused by a series of errors. One event triggers another, and the resulting state yet another, and so on. Feedback may either amplify or reduce the effects of error events. When the error symptom is finally detected, it may be far removed from its original locus in a completely different form and appear to have been triggered by some trivial event.

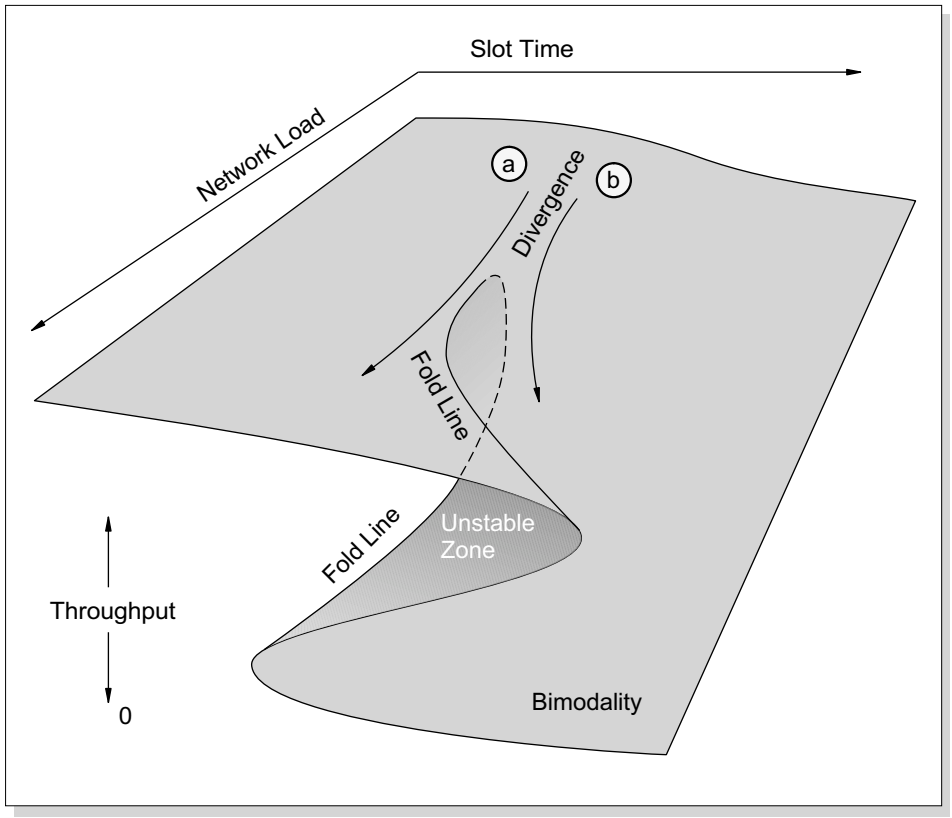

Figure 1.4 Catastrophe with two control parameters: the cusp

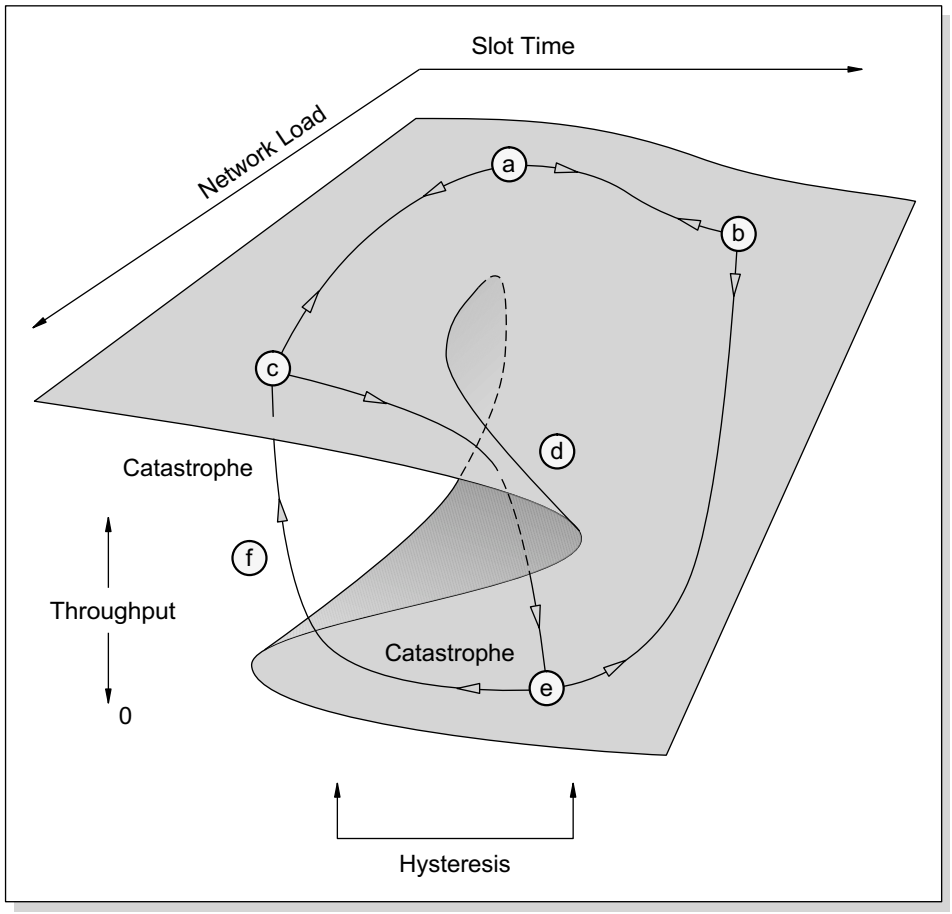

Figure 1.5 A non-linear operating state (network failure) in an Ethernet network

1.4 The Causes of Network Failure

There are five categories of errors that can lead to system failure:

- Operator error
- Mass storage problems
- Computer hardware problems
- Software problems
- Network problems

1.4.1 Operator Error

On the average, operator error is responsible for over 5 percent of all system failures—a large enough proportion to merit a closer look. Operator errors can be classified as intentional or unintentional mistakes, and as errors that do or do not cause consequential damage.

The term "intentional error" does not necessarily indicate that the error itself was the operator's intent, but rather that it resulted from some intentional action, such as trying to take a shortcut. The belief that a given process can be shortened, or that certain quality control or safety guidelines are superfluous, can lead to error situations with or without consequential damage. Less common are the truly intentional errors motivated, for example, by an employee's desire for "revenge" against a superior or the company, by the desire to cause trouble for a colleague (by making mistakes that the colleague will be blamed for), or out of destructiveness brought on by general frustration.

Unintentional errors usually result either from insufficient understanding of a given process or from poor concentration. Other common causes include software and hardware errors (the system does not behave as it should even though it is configured and operating correctly) or installation and configuration errors (errors occur when the system is operating correctly and the software or hardware is functioning according to specification). Sometimes a series of minor errors, which individually go undetected because no harmful effects are noted, are eventually compounded so that serious errors or even system failures result.

1.4.2 Mass Storage Problems

Problems with hard disks are the most common cause of failures in data processing. More than 26 percent of all system failures can be traced to faults in mass storage media. Although high-performance mass storage can attain a mean time between failures (MTBF) of over 10^6 hours, this could still mean replacing hard disks almost every month if the system has a large number of disk drives. There is usually a wide gap between theoretical MTBF and the operational MTBF that can be achieved in practice. The probability that a hard disk drive with a theoretical MTBF of 10^6 hours (almost 114 years) will actually run that long without error is only 30 percent. To calculate the number of hard disks that will have to be replaced within a certain period of time in a given system, multiply the total number of hard disk drives in the system by the period of system service in hours, and then divide this number by the theoretical MTBF. For example, in a system that has 1,000 disk drives, each of which has a theoretical MTBF of 10^6 hours, the number of failures A in the first 5 years (43,800 hours) comes to 44 (see the following equation).

$$A = \frac{1{,}000 \cdot 43{,}800 \ \frac{hours}{disk}}{1{,}000{,}000 \ \frac{hours}{disk}} = 44$$

This is based on the assumption that all of the hard disk systems have the same MTBF and are operated under similar conditions. Tests have shown that mass storage units operating in warm ambient conditions tend to show a lower actual MTBF than those operated in well-cooled environments. Furthermore, frequent disk search operations and changes in location have both been shown to have negative effects on the service life of mass storage media. For this reason, some hard disk manufacturers use another value in addition to the theoretical and operational MTBF to indicate the probable period of error-free operation for their products. This value, called the cumulative distribution function (CDF), indicates the probability that a mass storage medium will fail within a specified time. For example, a CDF of 4 percent over 5 years means that there is a 4 percent chance that the medium in question will break down within the first 5 years of use.

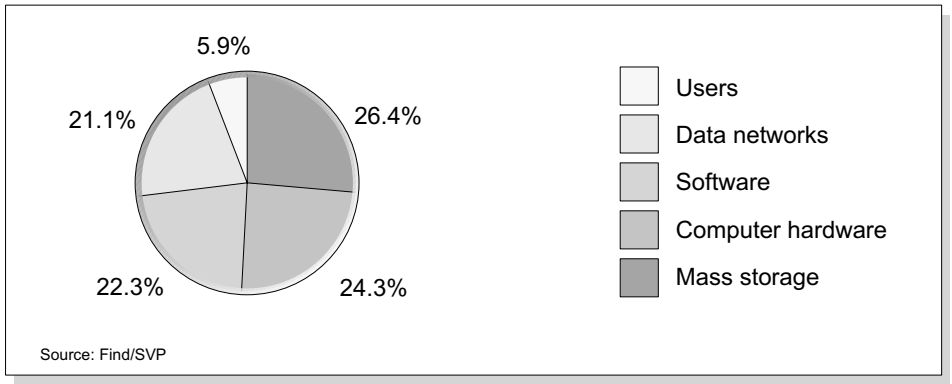

Figure 1.6 Causes of system failures in data processing

1.4.3 Computer Hardware Problems

Roughly one-quarter of all system failures are caused by computer hardware problems. By definition, this includes problems with any computer hardware component, including monitor, keyboard, mouse, CPU, RAM, hard disks and floppy disk drives. The average error-free service life of a system is calculated from the sum of the MTBF values of its components divided by the number of components. The following are some average MTBF values for various computer system components:

- RAM chips: 8,000,000 hours
- Floppy disk drives, mice, CD-ROM drives: 2,000,000 hours
- 10Base-T interface cards: 5,000,000 hours
- FDDI, ATM interface cards: 400,000 hours
- CPUs: 100,000 hours

The MTBF values calculated for today's computer systems average between 10,000 and 50,000 hours. In general, the more complex a system is, the lower the average MTBF. A system with multiple processors and multiple network links, for example, is more error-prone than a comparatively simple server with only one processor.

The Annual Failure Rate (AFR) is a better indicator of reliability than the MTBF. The AFR is the MTBF divided by the number of hours per year that the system is in operation. When a server system with a MTBF of 25,000 hours is in constant operation, the AFR amounts to 25,000/8,760 = 2.8, or about 3 failures every year. Another important parameter for the availability of computer systems is the mean time to repair (MTTR), which indicates the average length of time it takes to repair the system after a failure. The MTTR is the total repair time divided by the number of system failures. Typical MTTR values lie between 2 and 3 hours when the repair time used in the calculation is the amount of work time actually spent repairing the system.

1.4.4 Software Problems

Software problems cause almost as many failures as hardware problems do. The widespread use of client-server architectures and distributed platforms in enterprise networks have led to such complex combinations of software that it is almost impossible to monitor system behavior under all network loads and in all operating states. In the age of corporate intranets and the Internet, the update schedules for software applications are becoming shorter all the time, so that sufficient time is not allowed for detailed testing before software is released. Automatic testing tools, such as LoadRunner (from Mercury Interactive—*www.loadrunner.com*) or AutoTester (from AutoTester Inc—*www.autotester.com*), which attempt to simulate various extreme operating situations, provide only limited assistance. Problems with new software that can lead to system failure arise not only at the application level, but also as a result of unstable software drivers, faulty installation or backup procedures, or operating system errors.

1.4.5 Network Problems

The fifth major category of IT problems encompasses errors that occur within the network itself. When the software and hardware problems that are directly related to network operation are included in this category—such as problems with network interface cards or with certain components of application software, protocols and card drivers—this group accounts for more than one-third of all IT failures. These network errors can be classified by OSI layer. As shown in Figure 1.7, 30 percent of all LAN errors occur on OSI layers 1 and 2. Typical causes are defective cables, connectors, or interface cards; defective modules in hubs, bridges or routers; collisions (in Ethernet networks); beacon processes

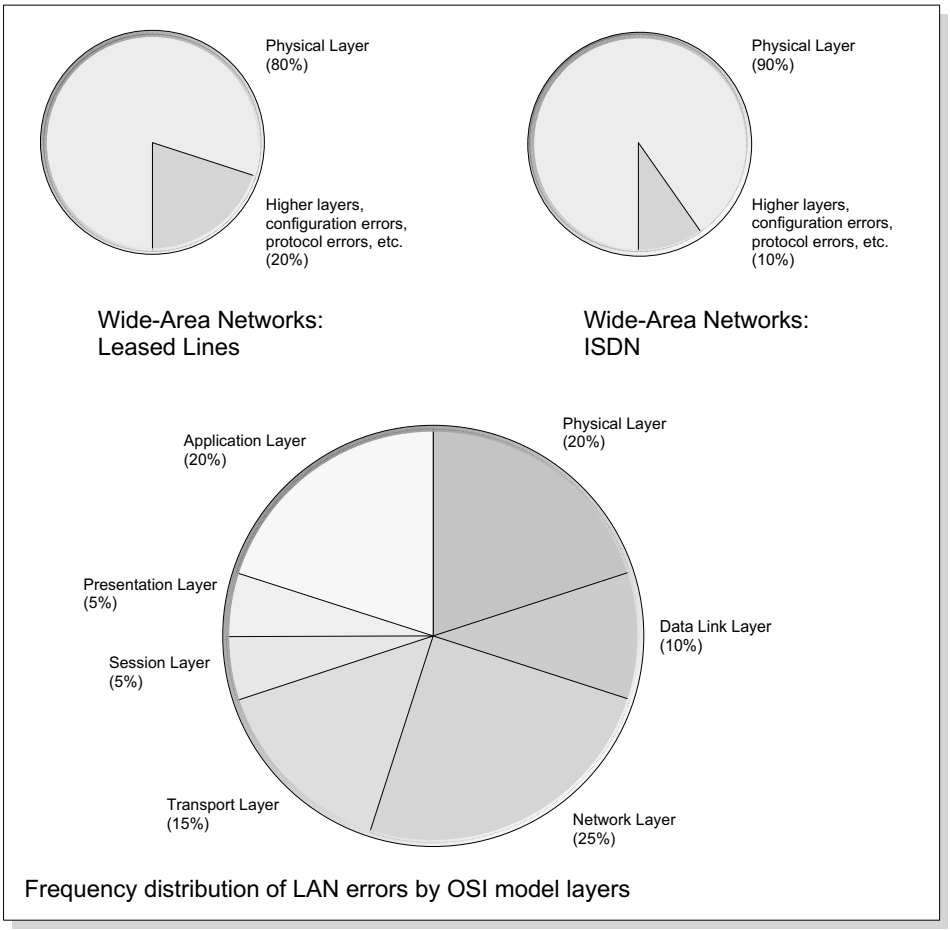

Figure 1.7 Distribution of data network problems in local- and wide-area networks

(in Token-Ring networks); checksum errors, and incorrect packet sizes. The development and implementation of more reliable hardware components, coupled with the continuous improvement of cabling systems, have meant a decrease in the absolute numbers of these types of errors, but developments in software have had similar effects on the higher OSI layers as well. More stable network operating systems and applications as well as mature protocol stacks have also reduced the number of failures per network segment. As a result, the distribution of error sources over the seven OSI layers remains roughly the same over the past years.

In wide-area networks (WANs), the proportion of errors occurring on the physical layer is even higher. Where permanent WAN links (leased lines) are employed, 80 percent of all errors—in the case of ISDN (Integrated Services Digital Network), as many as 90 percent of all errors—can be traced to component failure, defective modems, or cable and connector faults (see Figure 1.7).

1.5 Calculation and Estimation of Costs Incurred Due to Network Failures

It is becoming increasingly important to estimate the costs that are incurred in the event of network failure. These can be difficult to quantify, however. Nonetheless, a fairly clear idea of the financial impact of system failure is essential in order to determine the optimum infrastructure dimensions from the perspective of network management and maintenance. Knowing the costs of system failure enables the enterprise to make informed decisions regarding the level of investment in redundant components or network management and troubleshooting systems. All too often the costs of system failure are grossly underestimated. It may be true that the exorbitant losses of $100,000 per minute and more reported in some superficial studies apply only to a few special cases, such as when system failure affects production control systems, financial services offered by credit card companies, or investment brokerages. Nonetheless, the consequences of a network breakdown even in smaller- or medium-sized companies should not be underrated.

The average availability of a data network today is between 98 and 99 percent. A system that is in operation 10 hours a day, 5 days a week can expect network downtime totaling between 52 and 104 hours per year. If an average of 100 employees are affected by a network failure, this means a maximum loss of productivity between 5,200 and 104,000 hours. This type of oversimplified calculation, however, quickly leads to inflated figures that do not necessarily reflect real situations.

The first step toward a more realistic analysis of network downtime costs is to distinguish between immediate costs incurred within the first 24 hours following the failure and consequential costs that arise after the first 24 hours. Costs in each of these categories are further divided into direct and indirect costs. Direct costs include all expenditures that are directly involved in correcting the network problem, while indirect costs include such factors as lost employee productivity and delayed project completion.

1.5.1 Immediate Costs

(Costs arising within the first 24 hours)

Direct Costs

- Replacement parts (network cards, cable, repeaters, hubs, etc.)
- New components (bridges, routers, servers, etc.)
- Rental or purchase of diagnostic equipment (network analyzers, cable testers, etc.)
- Consulting fees charged by network specialists
- Consulting fees charged by software/hardware manufacturers
- Overtime compensation for network support staff

Indirect Costs

- Loss of employee productivity at computer workstations
- Loss of productivity on production lines, in shipping and receiving departments, or in warehouse management; downtime of automated warehousing systems, etc.
- Loss of consumer or customer orders and confidence

Easiest to calculate are the direct immediate costs, such as the purchase of replacement components or consultants' fees, because these are automatically documented by invoices. In mid-sized networks (around 500 nodes) with an availability of 99 percent, the average downtime of 52 hours results from an average of 10 to 20 failures of the network or parts of it, lasting between 1 and 5 hours each. If the direct immediate costs of solving the problem average $1,250 per case, then the direct cost of restoring operation after 10 failures comes to $12,500.

Quantifying the indirect immediate costs is more difficult. In general, only the cumulative loss of employee productivity is calculated. The extent of this loss, however, depends mainly on the degree to which employee productivity is dependent on network availability. Often a number of employee activities can be postponed until the next day, or at least for a few hours, without significant loss

of productivity. In mid-sized office environments, therefore, loss of employee productivity is usually estimated at roughly 25 percent of the total network downtime. For example, if an average of 100 employees are affected by the network breakdown, at 52 hours of downtime per year the loss of productivity amounts to $52 \cdot 0.25 \cdot 100 = 1{,}300$ hours. If the average gross salary costs come to $40 per hour, this puts the immediate indirect costs at $52,000.

1.5.2 Consequential Costs

(Costs arising after the first 24 hours)

Direct Costs

- New or adjusted hardware configuration in the network (restructuring of servers, bridges, etc.)
- Testing of other network segments for errors similar to those that caused the failure
- Documentation of the system failure

Indirect Costs

- Delayed project completion (product development, production, etc.)
- Delayed services (tenders, invoices, entering transactions in accounts, etc.)
- Loss of customer loyalty and satisfaction

Consequential indirect costs resulting from network failure are the most difficult to calculate. These costs are also referred to as "company losses" because they cannot be attributed to any one department or cost center. The amount of such costs is proportional to the degree to which the company depends on network-supported processes. Tenders may have to be printed and sent a day later than planned, for example. Incoming orders and payments may be similarly delayed. Incoming deliveries may be blocked if receiving slips cannot be printed or automated warehousing equipment cannot be operated. Urgent shipments sent by special courier result in higher shipping rates. Late charges may be incurred for bills that cannot be processed. Sales may be lost due to unavailable Web-based ordering systems. Customers may grow dissatisfied if they cannot reach a support hotline, which means a loss of future orders. These are just a few examples of company losses as consequential costs of network failures. At a fairly low estimate of $1,000 in consequential indirect costs and $250 in consequential direct costs per failure, the total loss per year in this category, based on the conditions described previously, is $77,000. This means each hour of downtime costs the company $1,480. Or, to look at the case from another perspective, an improvement of a mere 0.1 percent in network availability saves the company $7,700.

Network availability	99%
Annual downtime (hours)	52
Number of employees affected per failure	100
Dependency of employee productivity on network availability	25%
Average annual failures	10
Average direct, immediate costs per year / per failure (replacement parts, etc.)	$ 12,000 ($1,200 per failure)
Average indirect, immediate costs (loss of employee productivity)	100 · 0.25 · 52 = 1300 h · $40 = $52,000
Average direct, consequential costs per year / per failure (planning, failure documentation, reconfiguration)	$2,500 ($250)
Average indirect, consequential immediate costs per year / per failure (company losses)	$10,000 ($1,000)
Annual network failure costs	**$77,000**
Hourly network failure costs	**$1,481**
Network failure costs per 0.1% downtime	**$7,700**

Figure 1.8 Calculating the costs of network downtime

1.6 High Availability and Fault Tolerance in Networks

High-availability data processing infrastructures have become a basic requirement for smooth business processes in commercial data processing. Barring special measures taken to maximize network availability, the average availability of today's IT systems is between 98 and 99 percent, which corresponds to a total annual downtime of 50 to 100 hours. For a growing number of companies, however, even this is too much downtime. Special systems can be added to boost network availability to between 99.9 and 99.9999 percent (99.999 percent uptime is equivalent to 6.8 minutes downtime in one year). In this way the average downtime-per-year can be reduced to a few hours or even, in the extreme case, a few minutes.

The costs of availability, however, increase almost exponentially with each additional decimal place. Before planning a high-availability system, it is impor-

Architecture	Availability	Typical failure duration	Annual downtime
Uninterruptible operation	100%	None	None
Fault tolerance	99.9999%	Ticks	0.5 minutes
Fail-over by cluster	99.999%	Ticks to seconds	Up to 5 minutes
Fault resilience (fail-over)	99.99%	Seconds to minutes	Up to 50 minutes
High availability	99.9%	Minutes	Up to 8 hours
Standard system	99%	Hours	Several days

Source: AT&T/Gartner/TPPC

Figure 1.9 Availability levels and downtime

tant to specify exactly what service levels are required. This determines the degree of availability that must be guaranteed. Availability is expressed as a percentage, calculated from the total operating time and the downtime:

$$Availability = \frac{total\ operating\ time - downtime}{total\ operating\ time}$$

Another important factor is the average downtime resulting from system failure, which is called the mean time to repair or MTTR. In most cases, a large number of short service interruptions, lasting only seconds or minutes, is acceptable, while just a few failures that last for several hours each have serious consequences.

The main prerequisite for a high-availability infrastructure is the use of high-quality components. Even without any special equipment or configuration for ultra-high availability, the quality of components is an important factor in the reliability of hardware and software. Component quality is also decisive for the performance of diagnostic tools and system and network management applications, as well as for the level of maintenance and support that can be attained. If no concessions are made in these areas, the availability of the data processing structure is bound to be significantly above average. Availability can only be improved beyond this level by the addition of components and services. These can include:

- Redundant components
- Software and hardware switching

- Detailed planning of every scheduled downtime
- Reduction of system administration tasks
- Development of automatic error reaction systems
- Thorough acceptance testing prior to installation of new hardware or software components
- Specifications and practice drills for operator response to system failure
- Replicated databases and application software
- Clustering

Redundant components can reduce the number of single points of failure in the network. When a given network component fails, its redundant counterpart is activated automatically. If the installation of fallback components is combined with software and hardware switching technologies, the redundant components can take over for malfunctioning components within seconds or even fractions of seconds. Reducing the level of interaction between the network and adminis-

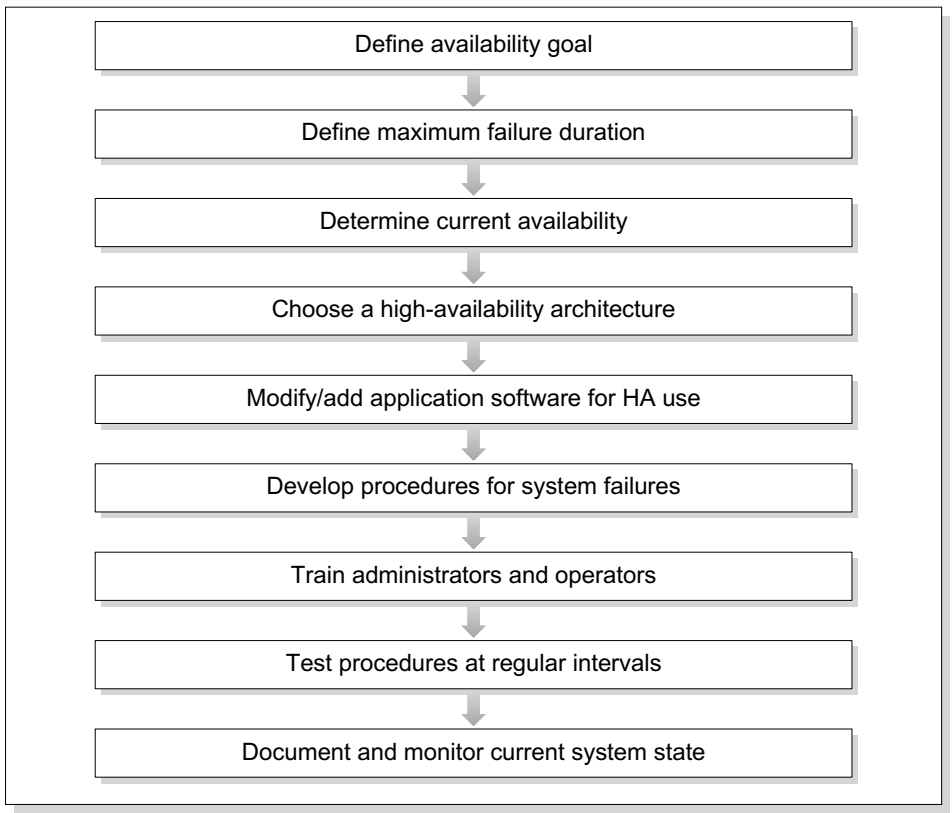

Figure 1.10 The introduction of a high-availability (HA) system

trator is a useful step in establishing deterministic reactions to different error scenarios—ideally, a given error should consistently trigger a single, defined process. The individual steps involved in introducing a high-availability system are shown in Figure 1.10.

1.7 Summary

Mission-critical systems in today's enterprise networks are growing more dependent every day on smoothly functioning data processing systems. Network managers are thus faced with the enormous challenge of increasing the availability of their data processing infrastructures while these infrastructures grow in both size and complexity. Network management is further complicated by the fact that corporate intranets are increasingly accessible through remote or public networks, such as the Internet, telecommunication service providers, customers' networks, telecommuters' systems and so on. It is no longer possible to have complete, end-to-end control over a company network. This makes it even more important to plan network operation and maintenance systematically, to implement appropriate procedures, and to have experienced network support staff equipped with advanced diagnostic and management tools.

For additional excerpts from this chapter and other Network Troubleshooting book sections, be sure to regularly visit our web site at:

www.FreeTroubleshootingBook.com

New chapters will be posted every 2 to 3 weeks.
Be sure to visit our web site and vote for the chapters you would like to see posted!

Agilent Technologies