



On the Job with a Network Manager

This chapter presents a number of scenarios to give an impression of the types of activities that are performed by people who run networks for a living. We refer to them collectively as network managers, although they perform a wide variety of functions that have more specialized job titles. In fact, strangely enough, the term *network manager* is rarely used for the people involved in managing networks. Instead, terms such as *network operator*, *network administrator*, *network planner*, *craft technician*, and *help desk representative* are much more common. Each of those terms refers to a more special function that is only one aspect of network management.

The chapter also provides an overview of some of the tools network managers have at their disposal to help them do their jobs. The intention is to give you a taste of the kinds of tasks and challenges that network managers face and how network management tools support their work.

Ultimately, the network management technology introduced in this book exists in an operational context. Although this idea might seem self-evident, it must be understood and emphasized, particularly for people who are not themselves users but are providers of network management technology—application providers, equipment vendors, and systems integrators. Network management involves not just technology, but also a human dimension—how people use management tools and management technology to achieve a given purpose, and how people who perform management functions and who are ultimately responsible for the fact that networks and networking services are running smoothly can best be supported. In addition, the organizational dimension must be considered—how the tasks and workflows are organized, how people involved in managing a network work together, and what procedures they have in place and must follow to collectively get the job done.

Reading this chapter will help you understand the following:

- The types of tasks that people involved in the day-to-day operations of networks face
- How network management technology supports network operators in those tasks
- The different types of management tools that are available to help people running a network do their job

A Day in the Life of a Network Manager

Let us consider some typical scenarios people face as they run networks. No single scenario is representative by itself. Scenarios differ widely depending on a number of factors. One factor is the type of organization that runs the network. We refer to this organization as the *network provider*. The IT department of a small business, for example, runs its network quite differently than the IT department of a global enterprise or, for that matter, a global telecommunications service provider. Another factor is the particular function that the network manager plays within the organization. An administrator in an IT department, for example, has different responsibilities than a field technician or a customer-facing service representative. To cover the diversity of possible scenarios, this chapter examines the roles of several network managers.

The examples in this chapter are intended to be illustrative. Therefore, they are by no means comprehensive. The examples contain simplifications, and, in reality, the details described differ widely among network providers. Even people who have the same job description might perform their job functions in different ways. Ultimately, how they manage their networks differentiates network providers from one another, hence the presented scenarios should not be expected to be universally the same. Finally, don't worry if you are not familiar with all the networking details that are contained in the examples; they constitute merely the backdrop against which the storylines play out.

Pat: A Network Operator for a Global Service Provider

Meet Pat. Pat works as a network operator at the *Network Operations Center (NOC)* of a global service provider that we shall call GSP. She and her group are responsible for monitoring both the global backbone network and the access network, which, in essence, constitutes the customer on-ramp to GSP's network. This is a big responsibility. Several terabytes of data move over GSP's backbone daily, connecting several million end customers as well as a significant percentage of global Fortune 500 companies. Even with the recent crisis in the telecommunications industry, GSP is a multibillion-dollar business whose reputation rests in no small part on its capability to provide services on a large scale and global basis with 99.999% (often referred to as "five nines") service availability. Any disruption to this service could have huge economic implications, leading to revenue losses of millions of dollars, exposing GSP to penalties and liability claims, and putting jobs in jeopardy.

Pat works directly in command central in a large room with big maps of the world on screens in front, showing the main sites of the network. Figure 2-1 depicts such a command central.

Figure 2-1 *An Example of a Command Central Inside a NOC*



(Figure used with kind permission from ish GmbH&Co KG)

In addition to the big maps, several screens display various pieces of information. For example, they show statistics on network utilization, information about current delays and service levels experienced by the network's users, and the number of problems that have been reported in different geographic areas. This gives everybody in the room a good overall sense of what is currently going on—whether things are in crises mode or whether everything is running smoothly.

Normally, everything on the map appears green. This means that everything is operational and that utilization on the network is such that even if an outage in part of the network were to occur, network traffic could be rerouted instantly without anyone experiencing a service outage. The network is designed to withstand outages and disruptions in any one part of the network. However, Pat still remembers the anxiety that set in on a couple occasions when suddenly links or even entire nodes on the map turned yellow or red. Once, for example, a construction crew dug through one of the main fiber lines that connect two of GSP's main hubs. And who could forget 9/11, when suddenly millions of people wanted to call into New York at the same time, while at the same time seemingly every news organization in the world requested additional capacity for their video feeds?

On Pat's desk is an additional, smaller screen that shows a list of problems that have been reported about the network. Pat has been assigned to monitor a region of the southeastern United States for any problems and impending signs of trouble. Pat sees on her screen a list of so-called *trouble tickets*, which represent currently known problems in the network and are used to track their resolution.

Those trouble tickets have two sources: problems that customers have reported and problems in the network itself. Let's start with customer-reported problems.

For every call that is received from a customer about a network problem, one of the customer service representatives at the help desk in building 7 opens a trouble ticket. The rep provides what GSP refers to as "tier 1 support." Those service reps have their own procedures. The person who first answers the call records a description of the problem, according to the customer, and asks the customer a series of questions, depending on the type of problem reported. If the service rep cannot help the customer right away, the customer is transferred to someone who is more experienced in troubleshooting the problem. That person is part of the second support tier. If this more experienced rep cannot solve the problem, or if it takes him or her too long to do so, the ticket is assigned to the people in Pat's group and shows up on Pat's screen. Pat's group provides the third tier of support.

The tickets contain a description of the problem, who is affected, and contact information. At least, this is what they are supposed to contain; sometimes Pat's group gets tickets with little or no information. In those cases, someone from Pat's group must call the service rep who first entered the ticket and find out more, which is always painful for everyone involved. It can be embarrassing when, in the worst case, Pat's co-workers need to call the customer back and the customer realizes that GSP is only starting to follow up on a serious problem hours after it was reported.

The second source of tickets is the network itself. These tickets are reported by systems that monitor alarm messages sent from equipment in the network. The problem with alarm messages is that they rarely indicate the root cause of the problem; in most cases, they merely reflect a symptom that could be caused by any number of things. Pat doesn't see every single alarm in the network—that would be far too many. For this reason, the alarm monitoring system tries to pre-correlate and group alarm messages that seem to point to the same underlying problem. For each unique problem that alarm messages seem to point to, the alarm monitoring system automatically opens a ticket and attaches the various alarm messages to it, along with an automated diagnosis and even a recommended repair action. Ideally, the underlying problem can be corrected and the ticket closed before customers notice service degradation and corresponding customer-reported trouble tickets are opened.

Seeing messages grouped in this way is much more practical than having to deal with every single alarm individually. The sheer volume of alarms would quickly overwhelm Pat and her group. Also, tickets that are system generated are typically issued against the particular piece of equipment in

the network that seems to be in distress. This makes system-generated tickets a little easier to deal with than customer-generated tickets, which often leave Pat's group feeling puzzled over where to start.

Pat remembers that tickets generated by alarm applications were problematic in the past. Often many more trouble tickets were generated than there were actual problems, so Pat sometimes saw 20 tickets that all related to the same problem. However, GSP has made significant progress in recent years—system-generated trouble tickets have become pretty accurate, with redundant tickets generated only in a small portion of cases. GSP's investment in developing better correlation rules for their systems paid off. Although Pat is an operator, not a developer, she knows that she was an important part of the development process because she provided much of the expertise that was encoded into those correlation rules. She still remembers being interviewed by a group of consultants for that purpose. During numerous sessions over the span of several months, they asked about how she determined whether problems that were reported separately were related.

Of course, despite all the progress made, many tickets still relate to the same underlying root cause. Many of those are tickets that were not automatically generated but instead were opened by customers. Perhaps a particular component in the access network through which customers were all connected to the network has failed, causing all of them to report a problem.

When clicking on a trouble ticket, Pat can see all the information associated with it. Pat must first acknowledge that she has read each ticket that comes in. If she does not acknowledge the ticket, it is automatically escalated to her supervisor. In busy times, this feels almost like a video game: Whenever a new ticket appears on the screen, she effectively "shoots it down" to stop it from flashing. Of course, acknowledging is only the first step. Next, Pat must analyze the ticket information. For the most part, her tasks are fairly routine. First she checks whether there are other tickets that might relate to the same problem. If there are, she attaches a note to the ticket that points to the other ticket(s) already being worked on. The system is intelligent enough to update the information in the other ticket to cross-reference the new one, thereby providing additional information that could prove useful in resolving it. This effectively leads to a hierarchy of tickets in which the original ticket constitutes a master ticket and the new ticket becomes a subordinate to the master. Pat then tables the resolution of the subordinate ticket until the master ticket that is already being worked on is resolved. At that point, she revisits the ticket to see whether the problem still exists or whether it can be closed also.

If she does not identify an existing ticket that might be related, she starts diagnosing the root cause of the problem. Let us assume that, in this case, the ticket was opened by a customer. Pat brings up the service inventory system to check which pieces of equipment were specifically configured to help provide service for that customer. With this knowledge, she brings up the monitoring application for the portion of the network that is affected to see for herself what is going on. This application offers her a view with the graphical representation of the device from which she can see the current state of the device, how its parameter settings have been configured, and the current communications activity at the device. She begins troubleshooting, starting with verifying the symptoms that are reported in the network.

In some cases, Pat eventually decides that a piece of equipment needs to be replaced, such as a card in a switch. In those cases, she brings up another tool, a work order system. She creates a new work order and specifies which card needs to be replaced. She enters the identifier of the trouble ticket as related information. This automatically populates the fields in the work order that identify the piece of network element, and also where it is located. Pat considers this to be a particularly nice feature. In the old days, she had to manually retype this information and also look up the precise location of the network element in the network inventory system. Now all those back-office systems are interconnected. She enters additional comments and submits the work order, and off it goes. This is all that she has to do for now.


It is not Pat's responsibility to dispatch a field technician or to check the inventory for spare parts; this is the job of her colleagues in the group that processes and follows up on equipment work orders. Actually, there are several groups, depending on where the equipment is located. Sometimes the equipment is in such a remote location that people have to physically get out there—"roll a truck," they call it. This is often the case for equipment in the access network. As mentioned earlier, the access network is the portion of the network that funnels network traffic from the customer sites to GSP's core network. In other cases—specifically, when the core network is affected—the equipment is at the NOC, in an adjacent building. Pat was once able to peek inside a room with all the equipment—many rows of rack-mounted equipment, similar to Figure 2-2.

Figure 2-2 *Rack-Mounted Network Equipment*



Pat's friends tell her that the NOC equipment is more compact than it is used to be, but Pat still finds it very impressive, especially the cables (cables are shown in Figure 2-3). Literally hundreds, if not thousands, of cables exist; taken together, they would surely stretch across many miles. You would never want to lose track of what each cable connects to. Although it all looks surprisingly neat, Pat can only imagine what a challenge it must be to move the NOC to a different location if that ever becomes necessary.

Figure 2-3 *Cabling and Equipment Backside*



(Figure used with kind permission from ish GmbH&Co KG)

Pat knows that the groups that do equipment work orders operate in similar fashion to her own group. The workflows are all predefined, and their work order system takes them through the necessary steps, autoescalates things when necessary, and generally makes sure that nothing can fall through the cracks—for example, it ensures that a work order does not sit unattended for days. It's impressive how integrated some of the procedures have become. For example, Pat has heard that when the technicians exchange a part, they scan it using a bar-code scanner that automatically updates the central inventory system. The system then warns them right away if they are scanning a different component than the one they are supposed to enter with the work order. In the past, occasional mismatches occurred between the equipment that was deployed and the equipment that was supposed to be there. This could lead to all kinds of problems—for example, equipment might be preconfigured in a certain way that would then no longer work as planned, or the installed equipment had different properties than expected. Those were rare but nasty scenarios to track and resolve.

Pat notes in the trouble ticket what she did and enters the identifier of the work order and when resolution is expected. For now, she is finished.


When the work order is fulfilled, Pat will find in her in-box a notification from the work order system identifying the trouble ticket that was linked to the work order and that should now be resolved. When she receives this notification, she does a quick sanity check to see if everything is up and running, and then closes the ticket for good.

When Pat first started her job, she was sometimes tempted to close the tickets right away without doing the check. Her department kept precise statistics on the number of tickets that she processed, the number of tickets that she had outstanding or was currently working on, the average duration of resolution for a ticket, and the number of tickets that had to be escalated. Of course, Pat wanted those numbers to look good because they were an indication of her productivity. Therefore, it was seemingly rewarding to take some shortcuts. It appeared that even in the unexpected case that a problem had not been resolved, someone would simply open a new ticket and no harm would be done. However, Pat soon learned that any such procedure violation would be taken extremely seriously. She now understands that procedures are essential for GSP to control quality of the services it provides. Doing things the proper way has therefore become second nature to her.

Chris: Network Administrator for a Medium-Size Business


Meet Chris. Together with a colleague who is currently on vacation, Chris is responsible for the computer and networking infrastructure of a retail chain, RC Stores, with a headquarters and 40 branch locations. RC Stores' network (see Figure 2-4) contains close to 100 routers: typically, an access router and a wireless router in the branch locations, and additional networking infrastructure in the headquarters and at the warehouse.

The company has turned to a managed service provider (MSP) to interconnect the various locations of its network. To this end, the MSP has set up a Virtual Private Network (VPN) with tunnels between the access routers at each site that connects all the branch locations and the headquarters. This means that the entire company's network can be managed as one network. Although the MSP worries about the interconnectivity among the branch offices, Chris and his colleagues are their points of contact. Also, the contract with the MSP does not cover how the network is being used within the company. This is the responsibility of Chris and his colleagues.

Figure 2-4 *RC Stores' Network*

Chris has a workstation at his desk that runs a management platform. This is a general-purpose management application used to monitor the network. At the core of the application is a graphical view of the network that displays the network topology. Each router is represented as an icon on the screen that is green, yellow, orange, or red, depending on its alarm state. This color coding allows Chris to see at first glance whether everything is up and running.

Even though the network is of only moderate size, displaying the entire topology at the same time would leave the screen pretty cluttered. Chris has therefore built a small topology map in which multiple routers are grouped into “clusters” that are represented by another icon. Each cluster encompasses several locations. In addition, there is a cluster each for the headquarters and the warehouse. This configuration enables Chris to display only the clusters and thereby view the whole network at once. Chris can also expand (“zoom into”) individual clusters when needed to see what each consists of. As with the icons of the routers, the icons for the clusters are colored corresponding to the most severe alarm state of what is contained within. This way, Chris does not miss a router problem, even though the router might be hidden deep inside a cluster on the map. As long as the cluster is green, Chris knows that everything within it is, too. Figure 2-5 shows an example of a typical screen for such a management application.

Figure 2-5 *A Typical Management Application Screen (Cisco Packet Telephony Center)*

Mike calls from upstairs. Someone new is starting a job in finance tomorrow and will need a phone. Chris notes this in his to-do list. He will take care of this later. First, he is trying to get to the bottom of another problem.

Chris received some complaints from the folks at the Richmond branch that the performance of their network is a little sluggish. They have been experiencing this problem for a while now; they first complained about it ten days ago when access to the servers was slow. At the time, Chris wondered whether this was really a problem with the network or with the server. As an end user, there was really no way to tell the difference. Eventually, the problem went away by itself and Chris thought it might have been just a glitch. Then three days ago, the same thing happened, and it did this morning again. This time Chris tried accessing the server himself with the Richmond people on the call but did not notice anything unusual.

Chris thinks that perhaps it really is a problem with the network. He wonders whether the MSP really gives them the network performance that they have promised. The MSP sold Chris's company a service with 2 Mbps bandwidth from the branch locations and "three nines" (99.9%) availability from 6 am until 10 pm during weekdays, 98% during off hours. The people from the MSP did not contact Chris to indicate that there was a problem on the MSP's side, but maybe they don't know—and besides, why would they worry if they didn't get caught? Chris wonders whether he should have signed up for MSP's optional service that would have allowed him to view the

current service statistics, as seen from the MSP's perspective, in near-real time over the web. Although Chris doesn't think the MSP can be entirely trusted, this would have provided an interesting additional data point.

From his management platform, Chris launches the device view for the router at the edge of the affected branch by clicking the icon of the topology map. The device view pops up in a window and contains a graphical representation of the device from which the current state, traffic statistics, and configuration parameter settings can be accessed. Currently, not much traffic appears to be going across the interface. From another window, Chris "pings" the router, checking the round-trip time of IP packets to the router. Everything looks fine.

Chris decides that this problem requires observation over a longer period of time, so he pulls up a tool that enables him to take periodic performance snapshots. He specifies that a snapshot should be taken every 5 minutes of the traffic statistics of the outgoing port. Chris also wants to periodically measure the network delay and jitter to the access router at company headquarters and to the main server. The tool logs the results into a file that he can import into a spreadsheet. Spreadsheets can be very useful because they can plot charts, which makes it easy to discover trends or aberrations in the plotted curves. (Of course, sometimes management applications support some statistical views as well, as shown in Figure 2-6.)

Figure 2-6 *Sample Screen of a Management Application with Performance Graphs (Cisco Works IP Performance Monitor)*



For now, that seems all that he can do. Chris takes a look at his to-do list and decides to take care of the request for the new phone. He doesn't know whether they have spare phones, so he goes to the storage room to check. One is left, good. He will have to remember to stock up and order a few more. He then peeks at the cheat sheet that he has printed and pinned in his cubicle, which has the instructions on what to do when connecting a new user. Most phones in RC Stores' branch locations are assigned not to individual users, but to a location, such as a cashier location, so changes do not need to be made very often.

RC Stores recently replaced its old analog private branch exchange (PBX) system with a new Voice over IP (VoIP) PBX. This enables the company to make internal phone calls over its data network. It also has a gateway at headquarters that enables employees to make calls to the outside world over a classical phone network, when needed. Chris remembers that, to make phone calls, the old PBX worked just fine, but programming the phone numbers could be a pain. Phone numbers were tied to the PBX ports, so he had to remember which port of the PBX the phone outlet was connected to so he could program the right phone number. Because RC Stores had never bothered documenting the cabling plan in the building, there were sometimes unwelcome surprises. Connecting one new user wasn't that bad, but Chris would never forget when they were moving to a new building and he and his colleague spent all weekend to get the PBX network set up to ensure that everyone could keep their extensions.

Now it is a simpler. Chris jots down the MAC address from a little sticker on the back of the IP phone and brings up the IP PBX device manager application. He also gets his sheet on which he notes the phone numbers that are in use. His method to assign phone numbers is nothing fancy. He has printed a table with all the available extensions. Jotted on the table in pencil is the information on whether a phone number is in use. Chris selects a number that is free, crosses it out, and notes the name of the new person who is assigned the number, along with the MAC address of the phone.

Chris then goes into the IP PBX device manager screen to add a new user. The menu walks him through what he needs to do: He enters the MAC address and the phone extension, along with the privileges for the phone. In this case, the user is allowed to place calls to the outside. Now all that remains to be done is to add voice mail for the user. He starts another program, the configuration tool for the user's voice-mail server. RC Stores decided to go with a different vendor for voice mail than for the IP PBX. Chris often moans over that decision. Although having different vendors resulted in an attractive price and a few additional features, he now has to administer two separate systems. Not only does he need to retype some of the same information that he just entered, such as username and phone number, but he also needs to worry about things such as making separate system backups. Chris leaves the capacity of the voice mail box at 20 minutes, as the application suggested for the default; it is the company's policy that everyone gets 20 minutes capacity except department heads and secretaries, who get an hour.

The phone extension is now tied to the phone itself, regardless of where on the network it is physically plugged in. Chris walks over to the Human Resources (HR) person upstairs and asks where the new employee will sit. He carries over the phone right away, plugs it into the outlet, and makes sure that it works. He must remember to send a note to HR to let them know the number he

assigned so they can update the company directory. Chris has been intending for some time to write a script that provisions new phones and automatically updates the company directory at the same time. Unfortunately, he has not gotten around to it yet. Maybe tomorrow.

Chris goes back to his desk and checks on the performance data that is still being collected. Things look okay; he will just let it run until the problem occurs again so that he has the data when it is needed. In addition, he decides that he wants to be notified right away when sluggish network performance is experienced. He goes again into his management platform and launches a function that lets him set up an alert that is sent when the measured response time between any two given points in the network exceeds a certain amount of time. He configures it to automatically check response time once per minute and to send him an alert to his pager when the response time exceeds 5 seconds. He hopes that this will give him a chance to look at things while the problem is actually occurring, not after the fact.

Chris realizes that the response time is needed for two purposes—once for the statistics collection function, once for the alerting function. Currently, there is no way to tie the two functions together. Therefore, the response times will simply be measured twice. Although this is not the most efficient method, there is no reason for Chris to worry about it.

Thinking about it, Chris suspects that the problem is related to someone initiating large file transfers. Perhaps an employee is using the company's network to download movies from the Internet. If this is the case, it would be a clear violation of company policy. Not only does it represent an abuse of company resources, but, more important, it also introduces security risks. For example, someone could download a program containing a Trojan horse from the outside and then let it run on the company network. Of course, Chris has set up the infrastructure to regularly push updates of the company's security protection software to the servers, but this alone does not protect against all possible scenarios. All the efforts to secure the network against attacks from the outside do not help if someone potentially compromises network security from the inside. Chris thinks that this hypothesis makes sense. The gateway that connects the company to the Internet is located at headquarters, and from the remote branch someone would have to go first via the company's VPN to that gateway to go outside. The additional traffic on the link between the remote branch and headquarters might be enough to negatively affect other connected applications. So maybe the problem resides with RC Stores after all, not with the MSP.

In any event, Chris knows that when the symptom occurs again, he will be able to find out what is going on by using his traffic analyzer, another management tool. He will be able to pull up the traffic analyzer from his management station to check what type of data traffic is currently flowing over a particular router—the gateway to the Internet, in that case—and where it originates.

Before Chris leaves in the evening, he forwards his phone extension to his mobile, in case something comes up. Also, he brings up the function in the alarm management portion of his management platform application and programs it to send him a page if an alarm of critical severity occurs, such as the failure of an access router that causes a loss in connectivity between a branch and headquarters. Chris has remote access to the VPN from home and can log into his management application remotely, if required.

Sandy: Administrator and Planner in an Internet Data Center

Meet Sandy. Sandy works in the Internet Data Center for a global Fortune 500 company, F500, Inc. The data center is at the center of the company's intranet, extranet, and Internet presence: It hosts the company's external website, which provides company and product information and connects customers to the online ordering system. More important, it is host to all the company's crucial business data: its product documents and specifications, its customer data, and its supplier data. In addition, the data center hosts the company's internal website through which most of this data can be accessed, given the proper access privileges.

F500, Inc.'s core business is not related to networking or high technology; it is a global consumer goods company. However, F500, Inc. decided that the functions provided by the Internet Data Center are so crucial to its business that it should not be outsourced. In the end, F500, Inc. differentiates itself from other companies not just through its products, but by the way the company organizes and manages its processes and value supply chains—functions for which the Internet Data Center is an essential component.

Sandy has been tasked with developing a plan for how to accommodate a new partner supplier. This will involve setting up the server and storage infrastructure for storing and sharing data that is critical for the business relationship. Also, an extranet over which the shared data can be accessed must be carved out. The extranet constitutes essentially its own Virtual Private Network that will be set up specifically for that purpose.

Sandy has a list of the databases that need to be shared; storage and network capacity must be assessed. Her plan is to set up a global directory structure for the file system in such a way that all data that pertains to the extranet is stored in a single directory subtree—perhaps a few, at most. She certainly does not want the data scattered across the board. Having it more consolidated will make many tasks easier. For example, she will need to define a strategy for automatic data backup and restoration. Of course, Sandy does not conduct backups manually; the software does that. Nevertheless, the backups need to be planned: where to back up to, when to back up, and how to redirect requests to access data to a redundant storage system while the backup is in progress.


Sandy's main concern, however, is with security. Having data conceptually reside in a common directory subtree makes it much easier to build a security cocoon around it. Security is a big consideration—after all, F500, Inc. has several partners, and none of them should see each other's data. A major part of the plan involves updating security policies—clearly defining who should be able to access what data. Those policies must be translated into configurations at several levels that involve the databases and hosts for the data, as well as the network components through which clients connect.

Several layers of security must be configured: Sandy needs to set up a new separate virtual LAN (VLAN) that will be dedicated to this extranet. A VLAN shares the same networking infrastructure as the rest of the data center network but defines a set of dedicated interfaces that will be used only by the VLAN; it allows the effective separation of traffic on the extranet from other network traffic. This way, extranet traffic cannot intentionally or unintentionally spill over to portions of the data center network that it is not intended for. The servers hosting the common directory subtree with

the shared data will be connected to that VLAN. Sandy checks the network topology and identifies the network equipment that will be configured accordingly.

Figure 2-7 shows a typical screen from which networks can be configured. This particular screen allows the user to enter configuration parameters for a particular type of networking port.

Figure 2-7 Sample Screen of a Management Application That Allows the Configuration of Ports (Cisco WAN Manager 15.1)



In addition, access control lists (ACLs) on the routers need to be set up and updated to reflect the new security policy that should be in effect for this particular extranet. ACLs define rules that specify which type of network traffic is allowed between which locations, and which traffic should be blocked; in effect, they are used to build firewalls around the data. This creates the second layer of security.

Finally, authentication, authorization, and accounting (AAA) servers need to be configured. AAA servers contain the privileges of individual users; when a client has connectivity to the server, access privileges are still enforced at the user and application levels. Any access to the data is logged. This way, it is possible to trace who accessed what information, in case it is ever required, such as for suspected security break-ins.

However, before she can proceed with any of that, Sandy needs to assess where the data will be hosted and any impact that could have on the internal data center topology. After all, without

knowing what servers should be connected, it is premature to configure anything else. When the partner comes online, demand for the affected data is sure to increase.

Sandy pulls up the performance-analysis application. She is not interested in the current status of the Internet Data Center because operations personnel are looking after that. She is looking for the historical trends in performance and load. Sandy worries about the potential for bottlenecks, given that additional demand for data traffic and new traffic patterns can be expected. She takes a look at the performance statistics for the past month of the servers that are currently hosting the data. It seems they are fairly well utilized already. Also, disk space usage has been continuously increasing. At the current pace, disk space will run out in only a few more months. Of course, some of the data that is hosted on the servers is of no relevance to the partnership; in effect, it must be migrated and rehosted elsewhere. This should provide some relief. Still, it seems that, at a minimum, additional disks will be needed. Given the current system load, it might be necessary to bring a new server with additional capacity online and integrate it into the overall directory structure. Sandy might as well do this now. This way, she will not need to schedule an additional maintenance window later and can thus avoid a scheduled disruption of services in the data center.

Of course, the fact that data is kept redundantly in multiple places will be transparent (that is, invisible) to applications. All data is to be addressed using a common uniform resource identifier (URI). The data center uses a set of content switches that inspect the URI in a request for data and determine which particular server to route the request to. The content switch can serve as a load balancer in case the same data and same URI are hosted redundantly on multiple servers. The content switch is another component that must be configured so it knows about the new servers that are coming online and the data they contain. Sandy makes a mental note that she will need to incorporate this aspect into her plan.

Observations

This should suffice for now as an impression of the professional lives of Pat, Chris, Sandy, and many other people involved in running networks. At this point, a few observations are key:

- Pat, Chris, and Sandy handle their jobs in different ways. For example, in Pat's case, there are many specialized groups, each dealing with one specific task that represents just a small portion of running the network. On the other hand, Chris more or less needs to do it all. Sandy is less involved in the actual operations but more involved in the planning and setup of the infrastructure. This work includes not just network equipment, but computing infrastructure as well. There is no "one size fits all" in the way that networks are run.
- Pat, Chris, and Sandy all have different tools at their disposal to carry out their management tasks. We take a look at some of the management tools in the next section. Not all tools that they use are management systems; in Chris's case, we saw how a spreadsheet and a piece of paper can be effective management tools.

- A major aspect of Pat's job is determined by guidelines, procedures, and the way the work is organized. Systems that manage operational procedure and workflows are as much part of network management as systems that communicate with the equipment and services that are being managed. Their importance increases with the size and complexity of the network (and network infrastructure) that needs to be managed.
- Some tasks are carried out manually; some are automated. There is no one ideal method of network management, but there are alternative ways of doing things. Of course, some are more efficient than others.
- Management tasks involve different levels of abstraction and, in many cases, must be broken down into lower-level tasks. Chris and Sandy both were at one level concerned with a service (a voice service in one case, an extranet in the other case), yet they had to translate that concern into what it meant for individual network elements. Sandy had to worry about how security policies at the business level, that state which parties are allowed to share which data, could be transformed into a working network configuration that involved a multitude of components.
- Many functions are involved in running a network—monitoring current network operations, diagnosing failures, configuring the network to provide a service, analyzing historical data, planning for future use of the network, setting up security mechanisms, managing the operations workforce, and much more.
- Integration between tools affects operator productivity. In the examples, we saw how Pat's productivity increased when she was supported by integrated applications, which, in that case, included a trouble ticket, a work order, and network monitoring systems. Chris, on the other hand, had to struggle with some steps that were not as integrated, such as needing to keep track of phone numbers in four different places (company directory, number inventory, and IP PBX and voice-mail configuration).

Later chapters will pick up on many of the themes that were encountered here, after discussing the technical underpinnings of the systems that enable Pat, Chris, and Sandy do their jobs. Before we conclude, however, let us take a look at some of the tools that help network providers manage networks.

The Network Operator's Arsenal: Management Tools

We conclude this chapter by taking a look at some of the tools that assist people who manage networks for a living—people like Pat, Chris, and Sandy. Ultimately, it is the goal of network management technology to provide tools that make people efficient. Having an impression of what such tools can do provides a helpful context for material covered in later chapters.

We start with simple and relatively basic tools and move progressively toward tools of greater complexity, concluding with tools that are typically found only in large-scale network operations.

The list is by no means complete but covers many of the most important tool categories. It illustrates the kaleidoscope of different functionality that is available to network providers. Perhaps it also explains why it is not uncommon to find literally hundreds of different management applications at large service providers. Don't worry, though. Many environments use far fewer applications, as with enterprise IT departments of medium-size businesses like the one encountered in the example with Chris. In addition, although the breadth of tools and functions might seem overwhelming at first, in later chapters we discuss how to bring order to all of this. For example, in Chapters 4, "The Dimensions of Management," and 5, "Management Functions and Reference Models: Getting Organized," we discuss systemic ways of categorizing and organizing management functionality; Chapter 10, "Management Integration: Putting the Pieces Together," picks up on the challenge of how to integrate different tools into one operational support environment.


Device Managers and Craft Terminals

Craft terminals, sometimes also referred to as device managers (not to be confused with element managers, discussed shortly), provide a user-friendly way for humans to interact with individual network equipment. Craft terminals are used to log into equipment one device at a time, view its current status, view and possibly change its configuration settings, and trigger the equipment to execute certain actions, such as performing diagnostic self-tests and downloading new software images. Frequently, craft terminals provide a graphical view of the equipment that shows the physical configuration of the equipment with its different cards and ports, viewed from both the front and the back sides. Figure 2-8 shows an example of such a view. The view might even be animated to show which LEDs will be currently lit or blinking, depending on the device's status.

Contrary to most other management tools, craft terminals generally do not retain any information about the managed equipment in a database, nor do they offer electronic interfaces to other management applications. All they provide is a remote real-time view of the equipment you want to look at, one at a time. In some cases, managed equipment might already provide a "built-in" craft interface, for example, by way of a mini-web server that renders a device view. In this case, separate craft terminal software is not needed because all that a user needs to do is point a web browser at the device.

Craft terminals are often used by field technicians, who might have craft terminal software loaded onto their notebook computers with which they connect to the device that needs to be managed through a universal serial bus (USB) or serial interface, much as you find on most PCs. In general, craft terminal functionality can also be launched from other management applications, such as from management platforms (see the section, "Management Platforms"), to provide a remote graphical view of the managed device. This was also the case in the earlier scenario, when Chris was using the function of a craft terminal that he launched from a management platform to take a look at the router at the edge of the branch that was having a performance problem.

Figure 2-8 *Sample Screen of a Device Manager Offering a Graphical Device (Chassis) View (CiscoView for Catalyst 6500)*



Network Analyzers

Network analyzers come under many different names, including packet sniffers, packet analyzers, and traffic analyzers. They are used to view and analyze current traffic on a network, generally to understand the way in which the network is behaving and to diagnose and troubleshoot particular problems. Network analyzers capture or “sniff” packets that flow over a port of a network device, such as a router or switch, and present them in a human-readable format that an experienced network operator can interpret. In the earlier example, Chris was planning to use a network analyzer to analyze the type of traffic that occurred during times when the network performance problem was observed.

Element Managers

Element managers are systems that are used to manage equipment in a network. Typically, element managers are designed for equipment of a specific type and of a particular vendor; in fact, they are often provided by an equipment vendor. Element managers are similar to craft terminals, in that they allow operators to access devices to view their status and configuration, and possibly modify their parameter settings. The functions of element managers, however, far exceed the functions of craft terminals. For example, element managers typically include a database in which they retain information about all the various devices (at least, for those that are supported) in the network.

This enables users to view how devices are configured without the devices themselves needing to be repeatedly queried. More important, it enables users to back up and archive how devices are configured, to restore device configurations if that ever is required, and to manage the distribution of software images to the devices. In addition, element managers can receive event messages from the devices, which enables users to monitor the various pieces of equipment across the entire network, not just one device at a time. Element managers might also be able to automatically discover equipment that is deployed on the network. The tool that Chris used in the earlier example to manage the IP PBX was an element manager.

Element managers also often offer an electronic interface to other applications. This allows other applications to manage the equipment through the element manager instead of having to interface to the equipment directly. This can have important advantages:

- Less possibility exists that data about the network will run out of synch between different applications. The element manager not only serves as an authoritative data store about the device, but also coordinates management requests that applications might issue concurrently.
- The interface that the element manager offers might be easier to use and, hence, build to than the interfaces offered by the devices themselves. The element manager can also shield applications from minor variations in device interfaces.
- The management load on the managed equipment is reduced. Not only can the element manager coordinate requests that are received from other management applications, but, in many cases, it can respond to requests by providing information about the device from its own database instead of needing to talk to the device.

Management Platforms

Management platforms are general-purpose management applications that are used to manage networks. The functionality of management platforms is generally comparable to that of element managers. However, management platforms are typically designed to be vendor independent, offering device support for equipment of multiple vendors. Typically, the primary task of a management platform is to monitor the network to make sure it is functioning properly. Therefore, it was also the main tool that Chris used in the earlier example. Management platforms are often accompanied by development toolkits. Those toolkits enable users, systems integrators, and third-party management application developers to adapt and extend the management platform. Its functionality can be customized and adapted to different environments, it can be extended with new capabilities, and it can be integrated with additional management applications whose functionality is made accessible through the management platform.

These capabilities can make management platforms resemble a sort of “operating system” for management applications. Indeed, in some ways, analogies between a management platform and a PC operating system can be drawn.

For example, the PC operating system includes basic functionality such as a file explorer and Internet browser, and might come bundled with a basic word-processing program and spreadsheet, with an abundance of additional applications available that run on top of it and make the PC operating system more useful. Those applications leverage certain operating system infrastructure, such as the file system. The management platform, on the other hand, provides out-of-the-box support for basic management needs such as network monitoring and discovery, with additional add-on applications available to cater to more advanced needs. Those applications use management platform infrastructure. An example are functions that allow applications to communicate with network devices, as well as functions that keep an inventory of the equipment in the network and that cache their configuration in an internal database. Also, where a PC operating system offers plug-in support for additional device drivers, management platforms need to support similar capabilities to support additional networking equipment.

Collectors and Probes

Collectors and probes are auxiliary systems that offload applications from simple functions.

Collectors are used to gather and store different types of data from the network. An example is Netflow collectors, which collect data about traffic that traverses a router. Such data can be generated by routers in high volumes and is commonly represented in a format known as Netflow. Another example is loggers, which collect so-called syslog messages from network equipment that provides a trail of the processing and activities that occur at a router.

Probes are similar to collectors but are “active,” in the sense that they trigger certain activities in the network and collect the responses—for example, they perform periodic tests. In the earlier scenario, Chris used a probe to take periodic measurements of the network response time over a certain link.

In each case, the data that is collected is made available to other applications, such as a management platform.

Intrusion Detection Systems

Intrusion detection systems (IDSs) help network providers to detect suspicious communication patterns on the network that might be indicative of an ongoing attack. Attacks include attempted break-ins into routers or, much more common, into servers, and denial-of-service (DoS) attacks that could be caused by Internet worms designed to overload and, hence, effectively shut down a service. IDSs use a wide variety of techniques, including analyzing traffic on the network,

listening to alarms, inspecting activity logs, and observing load patterns. IDSs help operators quickly recognize such threats and mitigate their effects—for example, by shutting off network ports through which attacks occur.

Performance Analysis Systems

Performance analysis systems enable users to analyze traffic and performance data, with the goal of recognizing trends and patterns in that traffic. They have to deal with massive amounts of data that has been collected over long periods of time; hence, they frequently involve data mining (techniques to recognize common patterns in large amounts of data), as well as advanced visualization techniques to display data in the form of graphical patterns that make sense to a user. Users such as Sandy from our earlier scenario use this information for a variety of activities, such as for network planning. Sandy can use information she gathers from a performance analysis system to anticipate where additional capacity will be needed in the near future and to tune data center performance based on an analysis of bottlenecks. Information gathered from performance analysis might even be helpful for tasks such as the development of pricing structures that will encourage communication behavior that helps “even out” the communications load on the network. Recognizing which services lead to a disproportionate load and frequently cause congestion in portions of the network might cause a service provider to charge extra for them.

Alarm Management Systems

Alarm management systems are specialized in collecting and monitoring alarms from the network. They help users to quickly sift through and make sense of the volumes of event and alarm messages that are received from the network. Often alarm management systems have additional capabilities to group (“correlate”) alarms that are likely to belong together, to offer initial diagnoses for the root cause of an alarm, or to provide impact analysis to forecast the fallout that an alarm might have. Sometimes, based on their analysis, alarm management systems generate additional synthetic event messages that aggregate and interpret the findings from a set of raw alarms. In many cases, alarm management systems also serve as preprocessors for other management applications, such as trouble ticket systems, like the one that we encountered in the earlier scenario with Pat.

Of course, other tools, such as management platforms and element managers, already include a certain degree of alarm management functionality. Dedicated alarm management applications, however, generally offer functionality that is more sophisticated and goes above and beyond what more general-purpose applications are offering.

Figure 2-9 shows a view of a screen of a typical alarm management application, displaying a list of alarm messages that can be expanded or searched and filtered for various purposes. For each alarm message, the screen shows a brief summary of what the message is about, along with information on which device it originated from, what category of alarm it belongs to, when it occurred, and (through its color coding) how severe the condition is.

Figure 2-9 Sample Screen of an Alarm Management Application (Cisco Info Center)

Node	Alert Group	Summary	Last Occurrence
loureed	Probe	A Probe process running on loureed has disconnected.	02/26/03 14:07:55
link4	Link	Link Down on port	03/28/03 11:05:24 1
wombat	Systems	Machine has gone offline	03/28/03 11:05:20
orac	Systems	Machine has gone offline	03/28/03 11:05:10
muppet	Systems	Machine has gone offline	03/28/03 11:05:05
link6	Link	Link Down on port	03/28/03 11:05:00
moose	Systems	Machine has gone offline	03/28/03 11:04:49
vixen	Stats	Diskspace alert	03/28/03 10:52:59
hal	Stats	Diskspace alert	03/28/03 10:42:18
vixen	Stats	Diskspace alert	03/28/03 10:53:23 1
hal	Stats	Diskspace alert	03/28/03 10:45:23 1
wombat	Systems	Machine has gone online	03/28/03 11:05:19
orac	Systems	Machine has gone online	03/28/03 11:05:18
angel	Link	Port failure : port reset	03/28/03 11:05:16
moose	Systems	Machine has gone online	03/28/03 11:05:06
muppet	Systems	Machine has gone online	03/28/03 11:05:04
dewey	Link	Port failure : port reset	03/28/03 11:05:03
link1	Link	Link Down on port	03/28/03 11:05:25

11 0 6 2 8 1 All Events

No rows selected. 03/28/03 11:06:10 root NCOMS [PRI]

Trouble Ticket Systems

Trouble ticket systems are used to track how problems in a network (such as those that are indicated by alarms) are being resolved. Note that this is different from managing the alarms themselves. Trouble ticket systems are used to capture information about problems that were observed in the network and to track the resolution of those problems. In many cases, trouble tickets are generated by users of the network who experience a problem, although they might also be created proactively by an application that monitors the network and detects a problem.

A trouble ticket system supports the resolution of problems in many ways. For example, the trouble ticket system can automatically assign trouble tickets to a ticket owner who has to take responsibility, or it can automatically escalate tickets that take too long to resolve. The trouble ticket system can also report statistics about the resolution process and generally ensures that problems are followed up on. Of course, the scenario that featured Pat was centered heavily on the use of a trouble ticket system.

Work Order Systems

Work order systems are used to assign and track individual maintenance jobs in a network. They also help organize and manage the workforce that carries them out. For each job, a work order is assigned whose resolution is then tracked. Similar to trouble ticket systems, work order systems offer a myriad of functions to capture information about jobs, to manage the assignment of jobs to a work force, to make sure those jobs are properly taken care of, and, in general, to track what the

work force that is maintaining the networking infrastructure is doing. We encountered a work order system in conjunction with the scenario of Pat when someone needed to be dispatched to replace a piece of faulty equipment.

Workflow Management Systems and Workflow Engines

A workflow management system helps manage the execution of workflows. A workflow is basically a predefined process or procedure that consists of multiple steps that can involve different owners and organizations. Workflow management systems pertain to business processes in general and are not specific to network management. However, they can be applied to network management when the processes and workflows in question involve the running of a network.

A workflow management system helps keep track of the steps in a workflow and ensures that predefined procedures are followed and policies are enforced. Workflows are usually defined using a concept called *finite state machines*. Each step along the way constitutes a state, and transitions between states occur according to well-defined interfaces and when well-defined events occur. The individual tasks are then pushed through these finite state machines as applicable, managed through the core of the workflow management system, the so-called *workflow engine*.

Both trouble ticket and work order systems can, in fact, be considered specialized examples of workflows. However, a workflow management system is more general in nature and highly customizable, to allow for the incorporation of any type of workflow.

Inventory Systems

Inventory systems are used to track the assets of a network provider. They come in two flavors:

- Network inventory systems track physical inventory in a network, mainly the equipment that is deployed, but sometimes also spare parts. Inventory information includes the type of equipment, the software version that is installed on it, cards within the equipment, the location of the equipment, and so forth. We encountered a network inventory system in the scenario involving Pat when the network technicians replaced a part in the network and the work order system automatically updated the network inventory accordingly.
- Service inventory systems track the instances of services that have been deployed over the network and that can be traced to individual users and end customers. For example, this could be DSL and phone services for residential customers of a telecommunications service provider. They might also include information on which network equipment and which ports are used to physically realize the service. Knowing this information makes it easier to assess who will be affected in case maintenance operations need to be performed or in case of a network failure.

In addition to inventory systems, facility management systems are used to document and keep track of the physical cable and ducts in buildings that are used to interconnect networking equipment.

Service Provisioning Systems

Service provisioning systems facilitate the deployment of services over a network, such as Digital Subscriber Line (DSL) or telephone service for residential customers of large service providers. Service provisioning systems translate requests to turn on or to remove a service into a series of steps and configurations that are then driven into the network.

Service provisioning systems are typically very complex applications that can be found only in operational support environments of large service providers; we did not encounter any in our earlier scenarios. They allow service providers to roll out services on a very large scale, often at a rate of tens of thousands of service requests per day. In many cases, service provisioning systems do not even interact with human operators, except possibly in case of exceptions that require human intervention. For this reason, perhaps surprisingly, they often offer no graphical user interface (GUI) or only a very rudimentary one. Instead, requests are issued from another system, for example from a service order management system via an electronic application programming interface (API). Such an interface allows another system to automatically interact with the system without user involvement, for example to request a piece of information or to hand off a request. For example, a service order management system (which we encounter in the section that follows) might use the API to automatically dispatch a request for provisioning a service to the service provisioning system when the order for the services becomes due.

Service Order–Management Systems

Service order–management systems are used to manage orders for services by customers of large service providers. (As with service provisioning systems, such systems are generally not encountered in enterprise environments.) They are part of a larger category of systems that deal with customer relationship management (CRM), which, for example, also includes help-desk functions.

Managing service orders involves a set of specialized workflows, similar to managing work orders or trouble tickets. Service order management systems help service providers track and fulfill orders for services and automate many, if not most, steps along the way. This includes identifying needed equipment, locating required ports, performing customer credit checks, scheduling the fulfillment of service orders, and eventually dispatching requests to turn up services to a service provisioning system.

Note the distinction between service order management systems and service provisioning systems. The former help manage workflows and processes of an organization. The latter are applications that interact with a network to configure it in a certain way. Compare this to the earlier distinction between trouble ticket systems and alarm management systems.

Billing Systems

Last in our list, but not least, are billing systems. We did not discuss billing systems in any of our earlier scenarios, but we should not lose sight of the reason many network providers (service providers, in particular, not enterprise IT departments) are in the business of running networks in the first place: to make money. Billing systems are essential to the realization of revenues. They analyze accounting and usage data to identify which communication services were provided to whom at what time. Subsequently, a tariffing scheme that defines how services need to be charged for is applied to that data to generate a bill.

Many other functions than billing systems themselves are associated with billing. For example, fraud detection systems help detect suspicious patterns in activity that could indicate that services are being stolen. Billing systems might also need to interface with other systems that are used for customer relationship management so that, for example, customer databases can be updated with information on which customers are past due.

Chapter Summary

In this chapter, we took a look at a few scenarios that illustrate how networks are being managed in practice and the variety of tasks that are involved. We followed three fictitious network operators and administrators: Pat in the Network Operations Center of a large service provider, Chris in the IT department of a medium-size business, and Sandy in the Internet Data Center of a large enterprise. The three scenarios represented operational support environments that differ greatly, as do the daily routines of the persons involved.

The service provider scenario emphasized workflows, processes, and interactions. In fact, in service provider environments, a significant part of the management infrastructure is dedicated to managing those organizational aspects, not just the technologies deployed in the networks themselves. The medium-size enterprise scenario was characterized by a great variety of tasks that had to be performed by the individual and a greater reliance on the individual expertise and intuition of the operator. The Internet Data Center scenario, finally, was geared at a different part of the network's life cycle, the planning phase. Also, it showed how the boundary between managing a network and managing the devices, servers, and applications that are connected to the network can become blurry.

The scenarios are representative of some of the environments in which management technology is ultimately applied. The scenarios also illustrate that network management is not just a topic of management technology; there are other significant factors in the equation, such as organizational aspects and human factors.

In each case, the personnel were supported by a variety of tools. In the end, management technology is tasked with building such tools, which are supposed to facilitate to the greatest extent possible the task of running a network. A wide variety of different tools exist for a great variety of purposes, so it comes as no surprise that running the largest, most complex networks can involve literally hundreds of management systems and applications. Of course, many scenarios are much simpler; it all depends on the particular context.

Chapter Review

1. Is running a network only a matter of network management technology, or are there other considerations?
2. What does Pat's employer use to track the resolution of problems in the network?
3. How does the integration of the work order system with the trouble ticket system make Pat's job easier?
4. Which network provider do you think will be more vulnerable to human failures by operations personnel, Pat's or Chris's?
5. Which of the following can be used as management tools? A. alarm management system, B. spreadsheet, C. pencil and piece of paper, D. all of them.
6. In how many different places does Chris need to maintain the same phone number, and why could this be an issue?
7. When Chris is worried about compromised security of his company's network, does the threat come from outside attackers or from within the network?
8. Connectivity between different company sites is provided by an outside MSP. Why is Chris nevertheless concerned with monitoring traffic statistics across these outside connections?
9. When Sandy wants to implement a security policy for the Internet Data Center, at what different levels does she take security into account?
10. Why is Sandy interested in "old" performance data and traffic statistics, even though she is not monitoring actual network operations?