INTRODUCTION

VINT CERF AND BOB KAHN

I.1 INTRODUCTION

There can be little doubt in 2002 that Internet protocol (IP) networks, among which the public portion of the Internet is but one, have become essential elements in this century's communications repertoire. A broad definition of current-day network management might include order entry, provisioning, billing, fault detection, isolation and repair, traffic engineering, performance monitoring, and assurance. One might even include network security under the general rubric of network management. As these IP networks assume greater importance with their proliferation and increased use, a number of challenges in their management and operation emerge along with new areas of concern that have little relationship to current network-management problems.

I.2 SCALING

Chief among the existing challenges for the larger networks is the problem of scaling. As networks get larger, they incorporate increasing numbers of components. The information records describing the current configuration of the network also increase in size and the cascading effects of specific failures produce larger quantities of alarm indications. New techniques may be required to evaluate such alarms in real-time, as well as to filter and classify them as to severity, priority, and origin. Rule-based systems that apply artificial intelligence techniques to the assessment of alarms have proven to be useful tools, but this area is likely to remain people intensive for the indefinite future. Another side effect of

Managing IP Networks. Edited by Saleh Aidarous and Thomas Plevyak
ISBN 0-471-39299-5 © 2003 Institute of Electrical and Electronics Engineers

scale is the difficulty of assuring that the actual configuration of the network is matched in detail by the service provider's information records that are supposed to represent the same information. New tools and new companies are being formed to address this verification and validation challenge. Key to this will be the use of unique identifiers and other network metadata, as the number of entities increase and the use of more convenient terminology takes hold. A good example of this is in the use of familiar names to denote the burgeoning set of IPv6 addresses within a given organization.

A related complexity arises from the increasingly common practice of implementing IP networks atop virtual communications substrates. The dedicated, hard-wired, point-to-point circuits of yesterday are being replaced with virtual circuits derived from asynchronous transfer mode (ATM), frame-relay, synchronous optical (SONET), sychronous digital hierarchy (SDH), and multiprotocol label switching (MPLS) networks. These virtual resources have their own provisioning and configuration complexities, not the least of which is that virtual misconfiguration of a network resource is sometimes harder to detect than physical misconfiguration. From the network-management point of view, even when physical circuits are working properly, virtual circuits derived from them may be inoperable, making fault detection, isolation, and repair that much more complex. A related equipment problem may occur when embedded computing hardware is working, but certain software functions are rendered inoperable. Increasing scale exacerbates all of these problems.

Another measure of scale comes from the increasing numbers of networks that are interlinked in the Internet. Some of the network-management problems may be the result of faulty interactions between the networks (or due to equipment at the boundary of the networks). As the component networks increase in number, the number of potential interactions among the networks can increase more than linearly. Virtually all of the fault isolation, detection, and repair challenges of one network are multiplied as we attempt to resolve operational problems in the global Internet.

I.3 TRAFFIC ENGINEERING

Traffic engineering in IP networks is usually accomplished in several "layers." The primary management of traffic flow is a consequence of traffic routing at the IP layer and is animated by a variety of routing protocols including border gateway protocol (BGP), intelligent scheduling and information system (IS-IS) and open shortest path first (OSPF) to name three. These protocols are generally destination-based in the sense that for each possible IP destination, each router picks a particular "next-hop" router to which to send that traffic. In the future, the notion of layer may not accurately describe the situation accurately enough, particularly where dynamic interactions take place among various data structures to produce a given result. For example, agent-based systems will likely operate on this principle, and various operations may actually be composite and require cooperation at various "layers."

Network operators have found it useful to introduce an additional layer of traffic management mechanism in the form of virtual circuits below the IP layer. Traffic is typically categorized by destination and then distributed across alternative paths so as to make efficient use of the underlying transmission capacity. Virtual circuits (e.g., ATM or frame relay permanent virtual circuits, MPLS label switched paths) are used to created *adjacencies* between routers at layer 2, and the standard routing procedures are used to determine the next-hop router at layer 3. In effect, the topology of the links between routers and their capacity can be altered at layer 2 in accordance with the apparent traffic flow requirements between pairs of routers. At layer 3, it is possible that all routers might appear to be one hop from each other in a fully connected virtual network.

To achieve the benefits of this form of traffic engineering, it is useful to gather information about the flow of traffic across the network. Some vendors offer systems for gathering such data directly from routers; but under heavy traffic loads, such data gathering may potentially interfere with forwarding of traffic. An alternative is to gather the data directly from the physical circuits that connect routers (or the lower-level switches) together by copying the traffic to an analyzer that can filter source and destination Internet addresses and packet sizes from the traffic and produce a database of source/destination traffic requirements for the network. Such noninvasive measurement methods are attractive when the backbone circuit bandwidths and traffic loads are sufficiently high that self-monitoring by a router has a negative impact on its ability to forward traffic. Finally, it may be useful to send typed data through the network such that the network operators are aware of the nature of the digital information being routed in the network and can organize the network-management system to handle it accordingly.

I.4 SERVICE QUALITY

Among the key performance metrics in IP networks are delay, throughput, and any related measures of packet loss. Even in the absence of any service quality guarantees, it is vital for network operators to have concrete measures of network performance. Packet loss has powerful and negative effects on performance, especially for transmission control protocol (TCP) that interprets loss as congestion and responds by reducing its rate of transmission and retransmission of packets. While the discard of a small percentage of packets in transit by routers may be useful in avoiding localized congestion, packet loss, in general, will adversely affect throughput and delay, and can have a deleterious effect on real-time packet communications. Consequently, it is important for operators of IP networks to know whether and when a network is approaching limits to capacity. Once again, measurements are key and network-management engineering must take this into account.

There is increasing interest in the user population to find ways to assure the quality of service on IP networks. Requirements range from constraining end-to-end packet delay and loss to prescribed parameters or even to assure that the variance in packet interarrival times (“jitter”) can be constrained. The latter may prove very important for real-time applications, where information communicated in digital form to the user is converted at the user’s site into what is usually known as voice or video or some combination thereof. In addition to these performance characteristics, there is increasing interest by the network providers in being able to assure capacity for preferred customers or applications.

In times of crisis, the ability of the network to guarantee performance for a subset of critical applications can be a high priority. Finding ways to achieve this effect while circumscribing its abuse (e.g., by users who have *not* been subscribed to such preferred service) is an important design challenge. For example, the ability to mark packets as having priority needs to be balanced by the ability to confirm that the originator of the packets has the authority to make such markings. Alternatively the markings might be applied by an edge router only after verifying that packets originating on a given access circuit should be given priority (or a subset of them based on the class of traffic/packet type).

Even this becomes a challenge if traffic is flowing *between* networks operated by different service providers. In this case, either the providers must trust one another to properly mark the traffic or all such markings might have to be ignored. Even if there is trust, there may be a question about a commercial agreement between the service providers to give priority to such traffic, perhaps at a premium charge associated with a net settlement commercial arrangement. Few if any service providers are prepared to offer such internetwork services, but it seems likely that this will become an important requirement in the not-to-distant future. Of course, this determination may be further governed by terms and conditions in associated metadata that can be communicated with the data themselves.

I.5 ORDER ENTRY, PROVISIONING, BILLING

Network management depends in large measure on the use of repositories of information that describe existing network topology, the connectivity among various monitored devices, and the way in which customers at the edge are connected and serviced. Without reliable and accurate information, billing may be difficult to impossible; fault detection, isolation, and repair an impossible challenge; and traffic engineering a distant hope. In designing network-management systems, then, it is vital that the design include considerations for obtaining customer orders, correctly calculating the provisioning required to service the order, and properly capturing all of these data to enable effective fault detection, billing, and traffic engineering. Actual data capture to meet these requirements can also be subject to considerable engineering debate. The use of alarms versus polling for status is one such area for debate. The larger the network, the longer it may take to capture critical information through polling. But, if information is sent automatically, under some circumstances the data gathering can be overwhelmed by an excess of alarm data. The design of the network-management system must balance these competing alternative methods, using hierarchies of polling, alarm, and filtering to cope with the scaling problems.

As discussed earlier, interactions between networks can produce unanticipated problems. For example, lack of sufficient buffering in the hand-off of traffic from one network to another can result in loss of packets during specific high traffic intervals. It is conceivable that neither network may be aware of the problem, or know how to fix the anomaly without collaborating. Indeed, two such providers may not even know they have a problem without the specific input of others. More subtle problems, such as changes in the variance of key performance parameters, may only be determined on a “black-box” type of network characterization at the boundary of the network, or at the boundaries of a collection of networks.

Finally, in a multinetwork environment, the costs for various services may be the result of costs incurred by several parties (both network operators and applications providers). While end users may not need to know the detailed breakdown of costs, the parties involved in providing the services may eventually need to know the detailed breakdown of charges, in order to keep end-user costs under control.

I.6 NETWORK SECURITY

Because network management includes the provisioning of network resources, it is vital that this capability be enabled only for authorized personnel and systems. In the wrong hands, a network-management system can simply disable the entire network. Consequently, the design of the network-management system must incorporate highly reliable authentication of individuals empowered to operate the system and also authentication of various network control subsystems, so that the controlled components can confirm the authenticity and authority of these subsystems to issue commands that affect the configuration of all managed network components. In the wake of an increasing incidence of various forms of network attacks and potential for *inside* abuse by disgruntled or compromised employees, it is no longer sufficient to rely on systems such as firewalls to defend the network from various forms of attack. Security methods including the incorporation of strong authentication technology are called for. Plainly these techniques are only as effective as the observation of procedures for their use. Installing locks does no good if people fail to lock doors or if they leave the keys lying about.

Key management may be an increasingly important part of network management. Already the number of services that require passwords and other forms of authentication is beyond the ability of many to treat them separately. Keeping large "key rings" is one possibility, but more likely is the use of a few services that provide identity-management services. These will also be useful for managing internal network operations, and could also be offered to users of the network. Authenticity of users may require a combination of network-management techniques combined with identity management.

I.7 FUTURE CONSIDERATIONS

As we peer into a somewhat uncertain future, it is possible to discern some interesting possibilities. In addition to finding billions of people and Internet-enabled devices on the network, information itself may become a kind of first class citizen requiring its own management infrastructure. The idea that *digital objects* have an identity and a configuration analogous to the physical devices of the Internet leads one to imagine management requirements for constellations of active digital objects. The elevation of information to a place in the network once reserved only for physical devices suggests that network management will be a lively and challenging area in which to work in the coming decade.

One can imagine accessing digital objects whose location (whether inside the net or provided as an external service) is transparent to a given user. If such objects cannot be accessed, should this be treated as an applications failure and not the responsibility of the network? What if the network makes the selection of where on the network to access the digital object? How can the operator determine the overall health of the system in this context, independent of whether a given user has lodged a complaint or not?

A major concern within today's intellectual property community is the unauthorized access to literary and/or musical works structured in various digital forms. The provision of moving pictures experts group (MPEG) audio layer-3 (MP3) song files over the Internet being easily achieved technically has brought this issue to the forefront. Higher-speed network access may soon make access to audiovisual works in digital form (e.g., "movies") economical over IP networks as well, further compounding the matter, if acceptable solutions are not found.

Today's network infrastructure does not provide required service levels based on a determination of content (i.e., based on the nature of a digital form of expression), but rather attempts to do so based on stated "quality of service" (QOS). If agreement can be reached with the owners of intellectual property on how to determine if a given object is allowed to be accessed, as a prerequisite for performing expressly stated operations on it, then the use of QOS may suffice, even if the underlying information is unknown or even encrypted. This issue is also likely to be controversial, as network operators have historically had no obligation to consider such matters. In the wake of recent terrorist actions, this issue may have ramifications for other types of communication as well. Issues of successful binding of content (in whatever form) to authorizations for communication will have network-management implications as well as free-speech overtones. This will be especially true when the need to track unauthorized access and other communication services is raised for national security purposes.

I.8 SUMMARY

Network management has become more complex as networks have evolved in capacity, scale, and capability. The Internet is no exception. Its implementation has made increasing use of virtual resources that must themselves be managed in addition to managing the system as seen from an IP perspective. As the Internet becomes increasingly ubiquitous, and as traditional applications such as telephony, radio, and video are converted to various digital forms for purposes of access and other communication services, the network-management challenges for Internet operators will continue to evolve and make new engineering demands. These challenges will also appear with regard to other applications, such as interactive games, peer-to-peer services, mobile services, and Internet-enabled appliance management. They will have high priority as service quality *between* Internet service providers (ISPs) becomes as important as service quality *within* an ISP. As with many infrastructure problems, the perceived value of the new communications services is evolving along with the importance of discovering solutions to the associated network-management challenges associated with them.