Network Management

A primer for consumers, advocates and policy makers

Network management – the process used to maintain the integrity and operations of a communications network – is a beneficial and necessary function of modern networks that has nonetheless become embroiled in the controversies related to network neutrality and traffic congestion on the Internet.

This three-part series examines network management – what it is, how it works, what it means to consumers, and what policies would most benefit consumers and network operators in their efforts to maintain effective communications. Part I defines network management.

The Internet was initially designed with the simple goal of assuring an ability to connect among private, government and academic networks. The quality and reliability of the connection was not a key focus.

The original concept was that the network would remain simple and "dumb" in order to maximize the ability to establish connections, and that management of the communications would be the responsibility of the software on the devices that used the connection.

But the theoretical concept of a "dumb" network quickly gave way to the necessity for management within the network itself -- to keep data flowing smoothly by avoiding congestion points caused by variable loads and equipment failures. The original IP protocol was augmented with a "Transmission Control Protocol" (TCP) to manage problems of data loss, leading to the TCP/IP protocol suite that is the foundation of the "best efforts" Internet today.

Other network management enhancements followed -- real-world needs for firewalls, network address translators, web content caches and the like been necessary to allow for the more sophisticated applications and critical data that flow over the commercial Internet

Network Management

1. **All networks require management to carry out their mission of effectively connecting two or more points on the network.**
2. **The purpose of network management is to enhance the reliability and viability of the network.**
3. **Network management practices differ for each layer of the network.**
4. **Practices at each level must interconnect with and support one another.**
5. **Network management practices should optimize the experience of all users.**
6. **No application or user should be permitted to contravene, interfere with or compromise these management practices, since to do so would reduce the reliability and viability of the network.**

To put it in simple terms, every network must have some system to ensure the flow of data. Even a network consisting of two tin cans and a string must have a system to ensure that the string remains connected and taut; otherwise, communications will fail. In the case of today's communications networks, that system is called "network management."

How The Internet Works

Some parties continue to portray "The Internet" as a single global structure of "dumb pipes" that connect computers, users and applications the world over. In fact, the Internet consists of tens of thousands of individual networks that are separately owned and administered.

Each of these networks has its own mission, its own requirements, its own priorities and its own network management practices. Equally important, the Internet operates on underlying physical networks that also support a multitude of other uses and users, both IP-based and non-IP-based capabilities. How, then, is the Internet able to operate?

The answer lies not only in a common protocol, TCP/IP, but in the structured and rational manner by which the underlying physical networks operate.

A Model Of The Internet

In order to understand this, it is helpful to understand the layers of functionality that work together to make communications possible.

The commonly-accepted model for the Internet has five layers.[1]

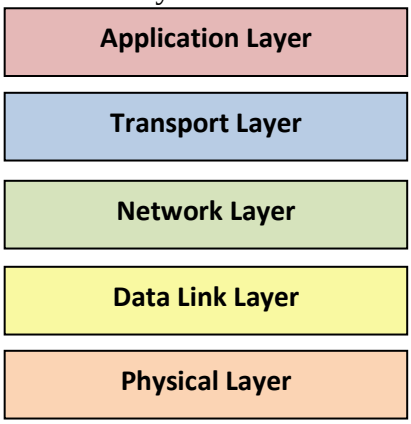

The Five-Layer Internet Model

Under this model, the Physical Layer represents the building blocks for a connection – the physical cables, wireless frequencies, the satellite signals, the microwaves and light beams over which communication is carried. It also includes the physical interfaces between the various elements and the specification for how "bits" are transmitted.

The Data Link Layer represents the ability to create an end-to-end connection between two fixed locations (or physical addresses). In effect, it is the roadway between two locations.

The Network Layer is where software and computers control the flow of communications over the networks described above. This layer identifies the association between physical addresses (fixed locations on the network) and an element that is not necessarily stationary – such as a laptop or cell phone.

---

[1] The "Internet model" is actually derived from the seven-layer "Open Systems Integration model." The Internet model is not official, but is commonly used as an easy way to think of how the Internet works. It combines the OSI model's Application, Presentation and Session layers into a single layer labeled "Application". For our purposes, this consolidation is not an issue.

The Transport Layer is a mechanism that validates that content reaches its destination. It is at this level that error detection and device-to-device flow control come into play. When TCP is employed, it adjusts the flow rate between two specific connected devices in response to congestion.

Finally, the Application Layer is where bits are broken down and reassembled to enable users to interact with web browsers, email systems, and other applications that are reached through the public Internet.

Network Management

Just as each of these five layers performs a different function in making communications possible, each uses different kinds of network management:

- At the Physical Layer, network management focuses on monitoring the performance of physical equipment and facilities in order to detect problems such as a cut cable, a faulty circuit board and the like. When an issue is detected, network management procedures generally call for a switch to a back-up component or a switch to another path to prevent or minimize disruption to the flow of information. After the traffic flow is shifted, the defective element can be repaired.

- The Data Link Layer is where the performance of the connection between two points in a larger physical network are monitored and managed. Management practices at this level focus on the integrity of the flow of data and the capacity available to handle it. It is not concerned with either the content in the packet or the destination.

- The Network Layer is where a continually updated "picture" of network configurations is stored. This allows multiple uses and users to share the same underlying physical capacity – one of such uses being the public Internet. Network management reduces the cost for all users by optimizing the use of available physical capacity. Critically, the available capacity is not constant due to the continually changing load as well as equipment and facility failures.

  For the Internet, this is where a path is selected for sending packets between users and content sources. A fixed end-to-end path does not exist. Rather, a path with available capacity is defined between points (or routers) and packets are sent between them one hop at a time.

  Each hop relies on software logic to continually move the packets closer to their ultimate destination. This path is subject to continual updating to reflect changing network conditions.

  A substantial amount of network status information must be updated continuously in order for the Internet to efficiently move billions of packets per second. Whenever a physical route is overloaded, an alternate route must be identified and used until the capacity constraint is resolved.

  When this happens, packets may be slowed in their journey (increased latency), may arrive at a non-uniform rate (increased jitter) or, in the extreme, not arrive at all (packet drop). Such events can happen at any time within

any individual physical network used in the Internet communication chain.

- The Transport Layer addresses the computer-to-computer packet flow. This is where the "TCP" part of the Internet plays its important role.[2] TCP operates to monitor for increased latency, jitter or packet loss, any one of which could indicate network congestion and could affect application performance.

- Finally, there is the Application Layer. The application layer, as it relates to network management, insulates the user from visibly experiencing minor degrees of variability in connection performance. Primarily, this involves "buffering" – the temporary storage of incoming data so that it can be delivered at a more consistent speed and interval. Application layer functions may also include screening or filtering to prevent unwanted incoming data, such as viruses or other harmful transmissions, from being processed.

Internet users want their communications to be reliable, even though the original design of the Internet favored simple connections over reliability.[3] Relatively few people understand that the Internet is not managed by a single service provider. Rather, each network that makes up the Internet is managed separately by that network's operator.

In the real world, cables get cut (by accident or otherwise), electronics and software periodically fail. While engineers carefully plan for all type of uses of the physical capacity, they are not perfect in their predictive powers. As a result, load spikes, whether due to legitimate or malicious traffic, are a fact of life. When these events occur, engineers must take steps to assure all competing uses of capacity, including the Internet, have equitable access to the network's capacity.

Network management practices, at all levels, are designed to assure sensible sharing of a physical network and to permit a variety of uses to co-exist and operate in an efficient manner. This saves money for everyone.

Because the Internet is a network of thousands of individual private networks -- each with their own design, policies and practices -- effective network management is a challenging task.

There is no central organization or body to impose a single set of rules on a configuration that extends to almost every nation and billions of individuals around the world. There is no simple "cook book" for network management, nor could one be created.

Because of the complexity and scope of the Internet, network engineers and network managers must continually reassess and adapt the approaches they use. Network managers need a wide array of tools that can be quickly applied and adapted to the situation at hand.

---

[2] TCP is also used and useful for IP-based networks other than the public Internet. In addition, other transport flow mechanisms (e.g., UDP) can be used in conjunction with IP and may also operate on the Internet. TCP is discussed as a non-exhaustive example of Transport Layer functionality.

[3] Internet mythology says that the Internet was designed simply and with little or no reliability so that it could survive a nuclear attack. While this may have been a policy consideration, the over-riding engineering consideration was to find a way to connect literally hundreds of different types of networks in a way that would let them "talk" to one another.

Network Management In Practice

Network management practices can be summarized in six general statements:

1. **All networks require management to effectively support communications between two or more points.** In the real world, networks that are not managed have substantial performance variability. But, stable and predictable performance is critical for advanced applications such as online banking, video streaming, Voice over IP telephony and eHealth.

2. **The purpose of network management is to enhance the reliability and affordability of the network.** Network management, at its heart, is concerned with three things. First, it protects the network – and, by extension, every computer and user on the network – from security threats. Second, it works to ensure the performance of the applications using the network. Third, it seeks to maximize the availability and use of capacity. By succeeding in this area, costs are held down for all users.

3. **Network management practices differ for each layer of the network.** Each layer addresses different functionality, so in turn the management techniques must focus on the purpose served by a particular layer. Regulating one "layer" jeopardizes the overall effectiveness of all layers and likely requires active regulation of all layers.

4. **Practices at each level must work with and support one another.** The Internet, and any other IP-based network, is designed in a way that each layer supports those above and below. If there is a failure in the Physical Layer, for example, a particular path may not be available for use. But the same is also true of other layers – if a hacker attacks the Application Layer (i.e., the end user) it may generate such a load that the Transport Layer cannot adapt or the Network Layer cannot find enough capacity to deliver legitimate traffic.

   Therefore, the management practices at each layer must not only be compatible with one another, but must work in tandem to be effective. Any action taken to remove or weaken a practice at one level will affect every other level, potentially jeopardizing the reliability of the entire network.

5. **Network management practices should optimize the experience of all users.** In today's world, a myriad of applications, communications and users must be accommodated at the same time on each physical network used. Network management provides the means to handle day-to-day variations in capacity requirements implicit in a dynamic environment such as the Internet, and does so in ways that benefits all users.

6. **Networks users and application are expected to respect the integrity of the network and act in accord with management practices.** Unfortunately, some Internet users act in unacceptable ways. Hackers assault destinations on the Internet in an effort to steal information, inject viruses, disrupt traffic flow, or damage targeted organizations.[4] Others use the network to perpetrate frauds and some users send large volumes of spam. Users of the Internet

[4] The Pentagon, headquarters for the military of the United States, claims that it suffers from three million attempted hacker attacks each day.

have differing need for bandwidth, and those needs may vary by the hour. Some applications aggressively consume capacity to the detriment of other users. Network management can deter or address behavior that is detrimental to the networks and users of the Internet.

These basics of networking are straightforward and simple. But there are occasions when network management is not so simple – as is the case when an application is designed to unfairly attack or circumvent the network management practices.

In Part II of this series, we will consider the impact of network management practices on consumers.

In Part III, we will discuss the means to handle misunderstandings and conflicts that have emerged or may in the future emerge with respect to network management on the Internet, and specifically the calls for new or more expansive restrictions on network management practices.

Author: David P. McClure
Date: August 8, 2008
Published by: US Internet Industry Assn.
1800 Diagonal Road
Suite 600
Alexandria, va 22314
(703) 647-7440 Voice
(703) 647-6009 Fax
(703) 851-4784 Mobile
InfoUSIIA@usiia.org
http://www.usiia.org

Formed in 1994, the US Internet Industry Association is the primary trade association for companies engaged in Internet commerce, content and connectivity. USIIA serves its members through legislative advocacy and professional services. The association is headquartered in Alexandria, VA.

David P. McClure is President and Chief Executive Officer of the US Internet Industry Association. A technologist by education and experience, McClure has held positions in the Internet, computing, aerospace and environmental services industries. He is widely published on technical and business topics, and is the author of more than 40 white papers related to Internet and Broadband policy, governance and economics.