Realtime publishers
"Leading the Conversation"

The Definitive Guide™ To

Converged Network Management

Ken Camp

Copyright Statement

Chapter 6: Impact Analysis, Root Cause, and Event Correlation

To conduct full impact and root cause analysis, event correlation engines are often used to provide data about what happened. This chapter begins by building an understanding of protocols involved, examining their strengths and weaknesses. Key protocols include Simple Network Management Protocol (SNMP), the Internet Control Message Protocol (ICMP) tools, and even Network Time Protocol (NTP) for effective event correlation.

Beyond protocols, the chapter will explore *syslog*. Syslog servers provide a part of the picture, but they really provide data collection mechanisms, not analysis engines. Correlating events across an enterprise network of disparate systems presents a difficult challenge.

In *Business @ the Speed of Thought* (ISBN: 0446525685), Bill Gates described what he called the “digital nervous system.” He said, “The most meaningful way to differentiate your company from your competition...is to do an outstanding job with information. How you gather, manage, and use information will determine whether you win or lose.” When deploying converged networks over IP, you’re integrating voice technology with the critical data infrastructure. Building monitoring and management processes into daily network operations provides the information, or knowledge base, about the corporate nervous system that lets you manage a complex, almost organic, business operating environment. Your management and monitoring tools become a key part of your enterprise business intelligence.

SNMP

SNMP is a widely used protocol for monitoring the health and well-being of a network. It’s a simple, text-based protocol that uses a database called a management information base (MIB) to describe network device management data. Almost all network elements are SMNP-enabled. Most equipment comes from the manufacturer with the community strings of *public* and *private* enabled by default. Typically the *public* string provides read-only access. The *private* community string often provides write access also, and is often used for managing devices remotely and “pushing” updated configurations to routers and switches across the network.

SNMP was designed to ease monitoring and remote management of network elements. These include servers, routers, switches, and even workstations. It can provide monitoring for performance, utilization, and state information about the device. SNMP uses what are called “traps” to capture this information, which is then often passed on to a centralized management station in a network control center. These stations typically provide network maps, with icons representing each node being monitored. In many systems, a simple green-yellow-red icon allows easy monitoring of network element status from healthy and operational (green) to potential problems (yellow) to out of service (red) conditions.

What Is a MIB?

The MIB is a type of database, comprising a set of objects used to manage individual network elements. MIBs are structured based on the OSI/ISO network management model. In the public switched telephone network (PSTN), Abstract Notation One (ASN.1) has been used for years as a mechanism for describing the object data structure of that network's elements. The PSTN elements include things like Class-5 central office switches, carrier trunking technologies, and the SS7 signaling network elements. ASN.1 was jointly developed by the ISO and the ITU-T in 1984. Today's network MIBs are developed as a subset of this larger standard. This subset is defined in IETF RFC 2578.

✎ IETF RFCs for MIBS

RFC 1156 - Management Information Base Network

RFC 1157 - A Simple Network Management Protocol

RFC 1441 - Introduction to SNMP v2

RFC 2579 - Textual Conventions for SNMP v2

RFC 2580 - Conformance Statements for SNMP v2

RFC 2578 - Structure of Management Information for SNMP v2

RFC 3416 - Protocol Operations for SNMP v2

RFC 3417 - Transport Mappings for SNMP v2

RFC 3418 - Management Information Base for SNMP v2

RFC 3410 - Introduction and Applicability Statements for Internet Standard Management Framework

RFC 3411 - Architecture for Describing SNMP Frameworks

RFC 3412 - Message Processing and Dispatching for the SNMP

RFC 3413 - SNMP Applications

RFC 3414 - User-based Security Model (USM) for SNMP v3

RFC 3415 - View-based Access Control Model for the SNMP

RFC 3584 - Coexistence between SNMP v1, v2 and v3

A MIB Object is one of any number of specific characteristics of a managed device. Examples of MIB objects include:

- Output queue length, which has the name ifOutQLen
- Address translation table (like ARP tables) called atTable

The Architecture of SNMP

There are three components needed for managing a network with SNMP:

- Network management systems—The network management system (NMS) runs the applications that control and monitor the network devices. The NMS is an active system that sends and receives SNOM queries. It accomplishes this by "setting" SNMP traps, then "getting" the results.
- Agents—An agent is simply the SNMP component of the software running on the device being monitored. This might be an integral part of the OS software or it might be another process or *daemon* that is executed when the device boots. The agent has information about the local operating characteristics of the network device. This information makes up the MIB for that specific device. The agent translates that information and provides the communication with the NMS.
- Managed devices—These are the network elements that are monitored by the NMS. These devices are typically routers, switches, servers, printers, and other service delivery elements of the network. In a unified communications design, these also often include the gateways, session border controllers, signaling and voice servers, and voicemail systems. Although workstation OSs—such as Windows—include SNMP capability, they are generally not monitored except when required for some very defined requirement. These devices often collect and store some form of management information locally in either event logs or *syslog* files.

SNMP provides a standards-based protocol and mechanism for remote monitoring and management of the unified communications network on a large scale. SNMP currently exists in versions 1, 2, and 3 in the real world. SNMPv2 was not widely adopted due to disagreements over the security framework, but many networks are evolving to use SNMPv3.

Version 3 includes some important new features. The most notable is encryption of the data in transit. Earlier versions send data in plaintext, which is easily read, making SNMP a prime tool for a malicious intruder to learn about the network. Encryption ensures that only the NMS and authorized personnel can read and evaluate this information. Different SNMP versions can interoperate to a limited degree. Interoperability between versions is explained in IETF RFC 3584.

Remote Monitoring MIBs

Remote Monitoring (RMON) is another technical specification that provides for a different variety of network monitors and console systems. RMON is designed to support network probes and monitors (often called sniffers). It allows the integration of diagnostic tools from multiple vendors, which may be used for very specific diagnostics or analysis.

RMON was initially developed when LAN switching became popular. It allows for managing switched LAN segments from a central monitoring facility or Network Operations Center (NOC). RMON is simply another extension of the standards already described as part of the SNMP MIB.

Unlike SNMP, RMON uses only two components. The probe contains the agent and is inserted into the network. One example would be a sniffer inserted into a specific network segment or VLAN for troubleshooting purposes. The other component is a management station. This workstation is frequently a network engineer's workstation, used interactively in troubleshooting and problem diagnosis. Like SNMP, RMON information uses the MIB found locally on the device, but the RMON agent is most commonly embedded in the OS. RMON agents don't monitor the entire system; only the traffic flowing through the RMON device. An RMON sniffer placed in listening mode on a LAN segment can only report on traffic on that LAN segment.

There are several variations in RMON MIBs. The Token Ring RMON MIB, for example, provides specific objects for managing a Token Ring network. The SMON MIB extends RMON and provides support for RMON analysis of a switched network.

SNMP is a very simple application protocol. Because it doesn't require a full three-way handshake or guaranteed communications, it's encapsulated in User Datagram Protocol (UDP). All three versions of SNMP contain the same message components:

- Version—The SNMP version number. It's key that the agent software running on the network device and the NMS use the same version of SNMP. Messages that arrive tagged as being a different version are typically discarded by the NMS.
- Community—The community name, or string, is used to authenticate the management system and grant either read or write access to the agent. The most common default strings, described earlier, are *public* and *private*. Many vendors' products come with other SNMP community strings enabled by default.
- PDU (Protocol Data Unit)—The PDU types and formats are different for SNMPv1, v2, and v3. A PDU is a descriptor of how information is packaged for a given technology. For example, the PDU for LAN technologies such as Ethernet and Token Ring is called a "frame." Ethernet and Token Ring frames differ in format, but the PDU they transport is a frame. IP packets are the PDU that IP transmits and differ slightly in format from other packets.

Using SNMP

SNMP uses IPv4 but also supports IPv6 for the future. The following list highlights the capabilities provided by SNMP:

- Data gathering—Collect data from a device that is SNMP capable. Single requests can be submitted using the *snmpget* and *snmpgetnext* requests. Multiple requests can be stacked using the *snmpwalk*, *snmptable*, or *snmpdelta* commands.
- Configuration modification—The *snmpset* command provides for altering the configuration information.
- Status checking—Commands such as *snmpdf*, *snmpnetstat*, and *snmpstatus* allow for retrieval of status and other information.
- Translation—Converting MIB information, content, and structure from text and numeric forms to other formats for use in a wide variety of analysis systems is accomplished using *snmptranslate*.

Many current SNMP tools provide a graphical MIB browser. Most organizations use graphical tools that provide some underlying, automated mechanism of implementing the *snmptrapd* command to automate receiving of SNMP notifications. A GUI can provide a human-friendly view that makes changes in the environment quickly observable. These notifications can also be logged to a syslog server or an event log or exported to a plain-text file. They can also easily be forwarded to other SNMP management systems and passed to external applications for event correlation and further analysis.

snmpwalk

In UNIX systems, *snmpwalk* is a widely available application. An administrator can run a very simple snmpwalk command

```
snmpwalk -c [good community string] [target host]
```

and learn a great deal of information about a device. Windows users can download a variety of SNMP exploration tools from the Internet. These tools generally eliminate arcane command-line interfaces, making basic exploration of networks and devices a simple point-and-click operation.

Figure 6.1 shows a simple snmpwalk of a print server on the author's network using a GUI tool from Solar Winds. In an enterprise network, routers, switches, servers, and VoIP service delivery systems can yield routing information, user account information, performance information, and details about TCP and UDP services running from this output information. When enabled, SNMP can provide an administrator with extensive information about an enterprise network very quickly.

> SNMP is a reconnaissance tool. If SNMP must be enabled, it is absolutely critical that default community strings be replaced. Just as a network administrator can use SNMP to perform quick network reconnaissance and learn information about the network that must be kept private, so too can an attacker. As a tool, SNMP is a double-edged sword, providing value while potentially exposing vital information.

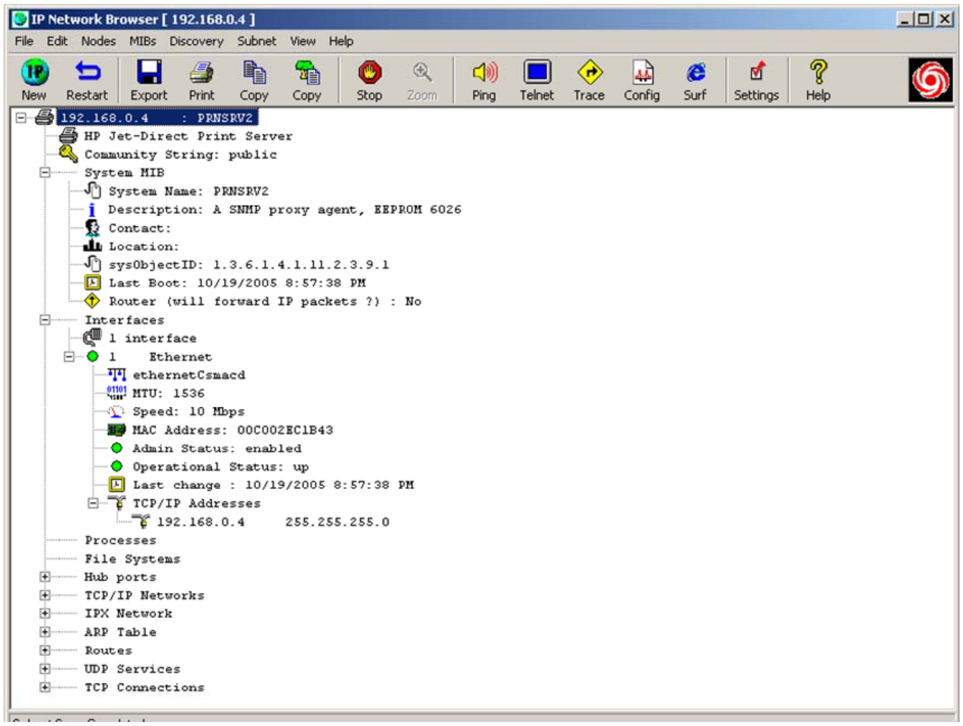
IP Network Browser [192.168.0.4]
File Edit Nodes MIBs Discovery Subnet View Help
New Restart Export Print Copy Copy Stop Zoom Ping Telnet Trace Config Surf Settings Help

```
192.168.0.4      : PRNSRV2
  HP Jet-Direct Print Server
  Community String: public
  System MIB
    System Name: PRNSRV2
    Description: A SNMP proxy agent, EEPROM 6026
    Contact:
    Location:
    sysObjectID: 1.3.6.1.4.1.11.2.3.9.1
    Last Boot: 10/19/2005 8:57:38 PM
    Router (will forward IP packets ?) : No
  Interfaces
    1 interface
    1     Ethernet
      ethernetCsmacd
      MTU: 1536
      Speed: 10 Mbps
      MAC Address: 00C002EC1B43
      Admin Status: enabled
      Operational Status: up
      Last change : 10/19/2005 8:57:38 PM
      TCP/IP Addresses
        192.168.0.4      255.255.255.0
  Processes
  File Systems
  Hub ports
  TCP/IP Networks
  IPX Network
  ARP Table
  Routes
  UDP Services
  TCP Connections
```

Figure 6.1: An example of SNMP information.

Factors to Consider with SNMP

SNMP versions 1 and 2 do not encrypt the transmitted data. This means that management information is passed in the clear and is quite readable by humans. There's a security risk in allowing critical management information to pass in the clear, even inside the enterprise network.

Because the different versions of SNMP are not compatible, use of SNMP for network management is often relegated to the lowest version supported in the network. For most enterprises today, that is SNMPv2, which does not support encrypted messages. Upgrading an enterprise network to SNMPv3 has often proven impractical. Existing routers and other network elements often cannot support the newest version. Although upgrading the OS might seem like a simple solution, often hardware replacement is the only viable means to upgrade to SNMPv3. The benefits of the latest protocol standard may not be a powerful enough business driver to warrant necessary hardware upgrades.

Autodiscovery

SNMP tools are widely used by malicious intruders for reconnaissance purposes. Many SNMP tools allow the simple use of subnet masking to run a scan across not just an individual network element but also a subnet or full network to discover what devices are listening for SNMP commands.

SNMP is a very simple network discovery tool. One of the features of SNMP tools is an automatic discovery feature, through which new devices discovered in the network are polled automatically. Most implementations will allow for a quick scan that yields tremendous information. Even if the public and private community strings have been set to a secure string that is not the default, the simple act of allowing SNMP enables discovery that quickly identifies working IP addresses on the network and the domain or network names associated with each.

Negative Implications

SNMP may be the intruder's easiest and most friendly tool. Software utilities are abundant for free downloading. Many are point-and-click operations, requiring no technical skill. It's quite common within an enterprise for employees to be curious and use these simple tools for network reconnaissance and exploration. Employees are typically within a trusted environment, so it's natural that they may have access to view a great deal of information. There is a danger of network topography being mapped from within because of this implied trust relationship.

Vendor's approaches to SNMP implementation vary widely. For some vendors, it isn't an element of the core product design but a feature that has been added or incorporated later in the product development cycle. Since the tree structure and data indexing techniques may vary, the internal data structures any particular vendor has implemented may vary. As a result, querying the network equipment with SNMP can produce in unwanted problems, like increased CPU utilization. Large routing tables, like those often found in Border Gateway Protocol (BGP) or Interior Gateway Protocol (IGP), are one example of a situation where this problem is likely to occur.

The lack of encryption capability in versions 1 and 2 introduce the threat of simple packet sniffing/capture, easily revealing the plain-text SNMP community string. No versions of SNMP use a challenge/response approach to authentication. That leaves all versions vulnerable to both brute-force and dictionary attacks. An assortment of both free and commercial software tools to instigate these attacks are readily available.

Because SNMP is UDP-based, it's connectionless in nature. This leaves SNMP vulnerable to IP spoofing attacks. Effectively restricting access to SNMP requires extensive access control list implementation across multiple network elements in many corporate networks.

It's noteworthy that SNMP has frequently surfaced in the SANS Institute's Top 10 Most Critical Security Threats as a result of the default community strings being set to *public* and *private*.

> ✎ SNMP has frequently surfaced in the SANS Institute's Top 10 Most Critical Security Threats as a result of the default community strings being set to public and private.

> 🕮 For more information about SNMP security implications, the US-CERT maintains an excellent SNMP Vulnerabilities FAQ at http://www.cert.org/tech_tips/snmp_faq.html.

ICMP

ICMP is a foundation of the TCP/IP suite. It is mainly used by networked computers' OSs to send error messages—indicating, for instance, that a requested service is not available or that a host or router could not be reached. In the connectionless, packet environment of IP, each host and router acts autonomously. Packet delivery is on a best-effort basis. Everything functions just fine as long as the network is working correctly, but what happens when something goes wrong within the subnet? As a connectionless service, IP has no direct mechanism to tell higher-layer protocols that something has gone awry. Furthermore, IP does not even have a method for peer IP entities to exchange information; if an IP host receives a packet, it attempts to hand it off to a higher-layer protocol. ICMP has been defined for exactly this purpose—IP-to-IP communication, usually about some abnormal event within the network. ICMP messages are carried in IPv4 packets with a protocol value of 1.

ICMP is defined in RFC 792 and is part of STD 5, which defines IP; this strongly suggests that ICMP is an integral part of IP. There are several types of ICMP messages. The following list highlights the most commonly used ICMP messages:

- Destination unreachable—Indicates that packets cannot be delivered because the destination cannot be reached. The reason is also provided. Examples include:
    - Host or network unreachable or unknown
    - Protocol or port is unknown or unusable
    - Fragmentation is required but not allowed (DF-flag is set)
    - Network or host is unreachable for this type of service
- Time exceeded—The packet has been discarded because the Time to Live (TTL) field decremented to 0 or because all fragments of a packet were not received before the fragmentation timer expired.
- Parameter problem—There was a problem with something in the packer header preventing a router or host from processing the packet.
- Source quench—Indicates that a router along the path is experiencing congestion and is discarding packets. This is usually caused by limitations in buffer space.
- Redirect—If a router receives a packet that should have been sent to another router, the router will forward the packet appropriately and let the sending host know the address of the appropriate router for the next packet.

The remaining ICMP messages are used to query the network for information:

- Echo and Echo Reply—Is used to confirm whether systems are active. One host sends an Echo message to the other. The destination system must respond with an Echo Reply with the same data that it received. These messages are the basis for the TCP/IP ping command.
- Timestamp and Timestamp Reply—These messages provide more information than the simple Echo messages. A timestamp, with granularity to the millisecond, is inserted in the messages. This provides a mechanism for measuring how long remote systems spend buffering and processing packets. It can also be used as a clock synchronization tool between hosts.
- Address Mask Request and Address Mask Reply—Can be used by network nodes to determine their address mask when assigned an IP address.
- Information Request and Information Reply—This field is now obsolete.

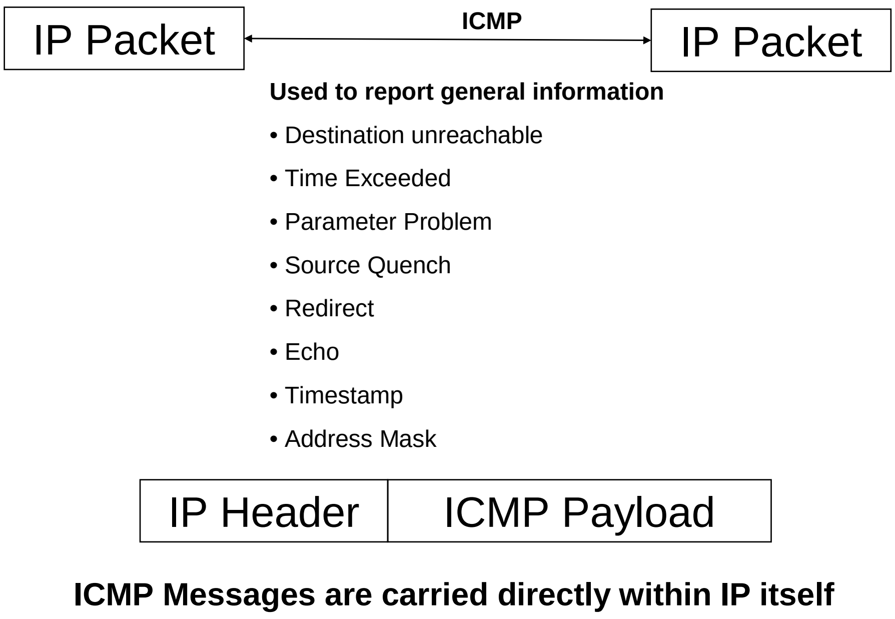

Figure 6.2: ICMP.

ICMP Message Format

Figure 6.3 shows the general format of an ICMP message. The following list highlights the first four bytes of all ICMP ("error" and "query") messages:

- Type—Indicates the type of ICMP message, including Echo Reply (0), Destination Unreachable (3), Source Quench (4), Redirect (5), Echo (8), Time Exceeded (11), Parameter Problem (12), Timestamp (13), Timestamp Reply (14), Address Mask Request (18), and Address Mask Reply (19).
- Code—Additional information specific to the message type. In the Time Exceeded message, for example, the Code field indicates whether the TTL counter was exceeded (0) or if the fragment reassembly timer expired (1).
- Checksum—16-bit checksum similar to that used in IP.

The next four bytes are labeled *miscellaneous.* They're used differently by different messages. In most ICMP "error" messages (for example, Destination Unreachable, Source Quench, Redirect, Time Exceeded, and Parameter Problem), these 32 bits are unused and set to 0. In the Parameter Problem message, however, the first byte is used as a pointer to the byte where the parameter problem was detected; in the Redirect message, these four bytes contain the address of the router to which future traffic should be directed.

The final field shown in the diagram contains the IP packet header plus the first 64 bits of the packet's Data field (or payload) in the offending packet. The receiving host uses this information to match the message to the appropriate CPU process. The 64 bits of user data are returned so that at least part of the header of any upper-layer protocol, including any port numbers, gets back to the original sender.

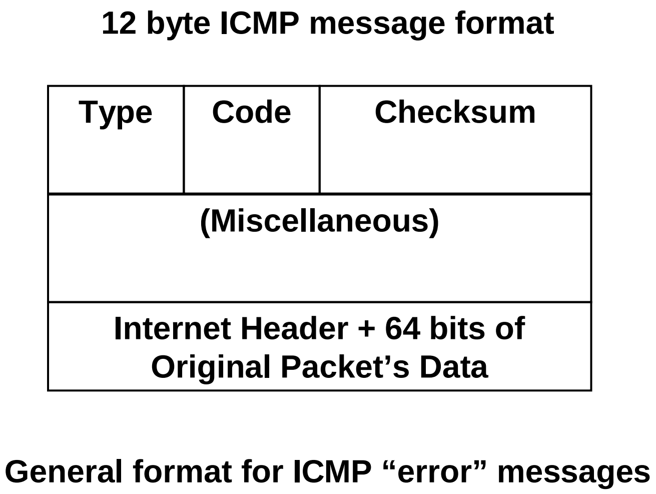
12 byte ICMP message format
Type
Code
Checksum
(Miscellaneous)
Internet Header + 64 bits of Original Packet's Data
General format for ICMP "error" messages

Figure 6.3: The ICMP message format.

ICMP differs in purpose from TCP and UDP in that it is usually not used directly by user network applications. One exception is the ping tool, which sends ICMP Echo Request messages (and receives Echo Response messages) to determine whether a host is reachable and how long packets take to get to and from that host.

ICMP Technical Details

ICMP is part of the TCP/IP suite as defined in RFC 792. ICMP messages are normally generated in response to errors in IP packets, per RFC 1122 specifications, or for diagnostic or routing purposes. The version of ICMP for IP version 4 is also known as ICMPv4, as it is part of IPv4. IPv6 has an equivalent protocol, ICMPv6.

ICMP messages are constructed at the IP, or network, layer. They are usually built from a normal IP packet that has generated some type of ICMP response. The appropriate ICMP message is encapsulated in IP with an IP header in order to return the ICMP message to the originating host.

For example, every router in the network that forwards an IP packet must decrement the TTL field of the packet header by 1. If the TTL reaches 0, an ICMP TTL Exceeded message will be sent to the source from that router.

Every ICMP message is directly encapsulated in a single IP packet. Like UDP, ICMP does not provide any delivery guarantees.

Although ICMP messages are contained inside standard IP packets, ICMP messages are usually processed as a special case. They're not normally treated as an IP sub-process because it's often necessary to inspect the contents of the ICMP message, then deliver the appropriate error message to the originating host and application.

Reachability Testing

One of the most crucial tests for network monitoring is the simple determination as to whether a system or network element is reachable via the network. The two most common tools for determining reachability are ping and traceroute.

ping

Ping is perhaps the most widely utilized tool on all TCP/IP systems. It allows users to determine the status of other systems. It also provides a tool for measuring the expected round-trip delay between the local system and a remote network element. Ping is useful for many reasons. Prior to attempting to establish a TCP virtual circuit, a local system might first ping the intended destination to verify that it is up and reachable.

Ping uses ICMP Echo and Echo Reply messages. It has the following general format (where items in square brackets [] are optional):

```
Usage: ping [-t] [-a] [-n count] [-l size] [-f] [-i TTL] [-v TOS]
            [-r count] [-s count] [[-j host-list] | [-k host-list]]
            [-w timeout] target_name
```

In the first test that Figure 6.4 shows, the test pings the host *www.yahoo.com* to determine whether it is up and running. This demonstrates the simplest use of the ping command and uses none of the optional parameters.

The second test in the figure uses the optional -t parameter to tell the workstation to send an ICMP Echo message continuously. The optional size and quantity parameters are not specified, so ping uses the default values—64-byte messages. These are sent continuously until the program is interrupted using Control-C to break the continuous cycle. The second test results in a list of the round-trip delays experiences by each Echo message sent.

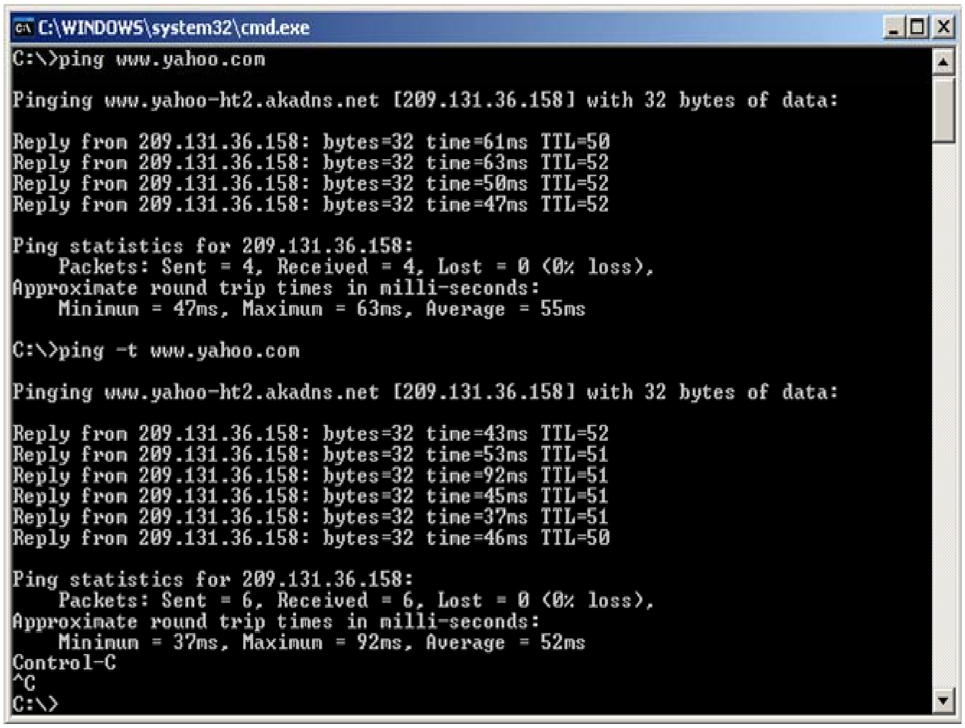

```
C:\>ping www.yahoo.com

Pinging www.yahoo-ht2.akadns.net [209.131.36.158] with 32 bytes of data:

Reply from 209.131.36.158: bytes=32 time=61ms TTL=50
Reply from 209.131.36.158: bytes=32 time=63ms TTL=52
Reply from 209.131.36.158: bytes=32 time=50ms TTL=52
Reply from 209.131.36.158: bytes=32 time=47ms TTL=52

Ping statistics for 209.131.36.158:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
    Minimum = 47ms, Maximum = 63ms, Average = 55ms

C:\>ping -t www.yahoo.com

Pinging www.yahoo-ht2.akadns.net [209.131.36.158] with 32 bytes of data:

Reply from 209.131.36.158: bytes=32 time=43ms TTL=52
Reply from 209.131.36.158: bytes=32 time=53ms TTL=51
Reply from 209.131.36.158: bytes=32 time=92ms TTL=51
Reply from 209.131.36.158: bytes=32 time=45ms TTL=51
Reply from 209.131.36.158: bytes=32 time=37ms TTL=51
Reply from 209.131.36.158: bytes=32 time=46ms TTL=50

Ping statistics for 209.131.36.158:
    Packets: Sent = 6, Received = 6, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
    Minimum = 37ms, Maximum = 92ms, Average = 52ms
Control-C
^C
C:\>
```

Figure 6.4: Ping.

Traceroute

Traceroute is another common TCP/IP tool that lets users learn about round-trip delays and the network routing between systems. Traceroute works by sending a sequence of UDP packets with an invalid port identifier to the destination system. The first three packets have the TTL field value set to 1; this causes the first router in the path to send back an ICMP message reporting that the TTL has expired. Then three more UDP messages are sent, each with the TTL value set to 2, which causes the second router to send ICMP replies. This process repeats, incrementing the number of router hops until the message actually reaches the destination. Traceroute identifies a completed cycle when it detects an “invalid port” error reply.

Figure 6.5 shows the route from a workstation on the author's network to *www.yahoo.com.* The route that is displayed tells the following:

- The first hop is through a system called GE-1-1-ur01.olympia.wa.seattle.comcast.net. This is the first point on the path in the Comcast provider network. The second hop is further along the way on the Comcast network.
- Hop 3 shows a hop across the AT&T provided backbone network in Seattle.
- Hops 7 – 10 traverse the Level3 backbone network from Seattle to San Jose.
- Hop 11 hits a network node that only provides an IP address, 4.71.112.14.
- Finally, on the Yahoo network, we hop across to the final destination at 209.131.36.158, which has the DNS name f1.www.vip.sp1.yahoo.com. This is the system responding to the request for www.yahoo.com.

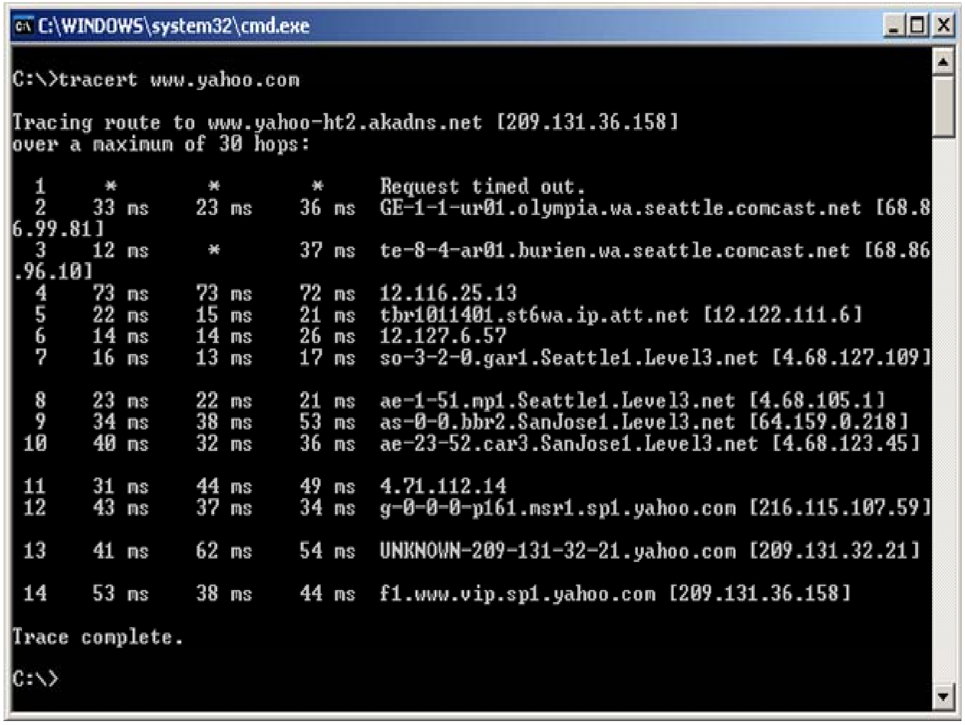

```
C:\>tracert www.yahoo.com

Tracing route to www.yahoo-ht2.akadns.net [209.131.36.158]
over a maximum of 30 hops:

  1     *        *        *     Request timed out.
  2    33 ms    23 ms    36 ms  GE-1-1-ur01.olympia.wa.seattle.comcast.net [68.8
6.99.81]
  3    12 ms     *       37 ms  te-8-4-ar01.burien.wa.seattle.comcast.net [68.86
.96.10]
  4    73 ms    73 ms    72 ms  12.116.25.13
  5    22 ms    15 ms    21 ms  tbr1011401.st6wa.ip.att.net [12.122.111.6]
  6    14 ms    14 ms    26 ms  12.127.6.57
  7    16 ms    13 ms    17 ms  so-3-2-0.gar1.Seattle1.Level3.net [4.68.127.109]

  8    23 ms    22 ms    21 ms  ae-1-51.mp1.Seattle1.Level3.net [4.68.105.1]
  9    34 ms    38 ms    53 ms  as-0-0.bbr2.SanJose1.Level3.net [64.159.0.218]
 10    40 ms    32 ms    36 ms  ae-23-52.car3.SanJose1.Level3.net [4.68.123.45]

 11    31 ms    44 ms    49 ms  4.71.112.14
 12    43 ms    37 ms    34 ms  g-0-0-0-p161.msr1.sp1.yahoo.com [216.115.107.59]

 13    41 ms    62 ms    54 ms  UNKNOWN-209-131-32-21.yahoo.com [209.131.32.21]

 14    53 ms    38 ms    44 ms  f1.www.vip.sp1.yahoo.com [209.131.36.158]

Trace complete.

C:\>
```

Figure 6.5: Traceroute.

📖 For more information about traceroute, see RFC 1393.

Syslog, Data Logging, and the Console

The term *syslog* is used to describe both the application that sends syslog messages and the syslog protocol itself. The syslog protocol, defined in RFC 3164, is very simple. A syslog sender transmits a small text message to the syslog receiver or server.

Syslog is available almost universally to aid in systems management and security auditing. Although syslog has several shortfalls, it is widely supported by almost every element of the network. Because it's nearly ubiquitous, syslog can be used to integrate log data from many different types of systems into a central data store for event correlation and analysis.

Syslog is used for network management and security auditing. Syslog itself is quite simple. It may seem simple for auditing use, but its broad availability is a great advantage. It allows a centralized, corporate syslog server to become the central data repository for audit and event correlation information. Syslog data is in plain-text format, so it's easy to manipulate with standard simple tools and scripts. Most organizations start with scripts and spreadsheets for analyzing syslog data. Larger organizations, monitoring many devices, may find that this approach is too labor intensive to be effective. Large log files and large numbers of log files may require adopting scalable commercial tools and developing automated processes to ease the work involved.

Technology alone can't solve the anomaly detection problem. A great deal of syslog and event monitoring is tied not just to performance but also to network security. What gets monitored, how log data is used, and how the organization responds to events at the time of detection are all a critical part of the cycle of network management, monitoring, and defense. Administrators employ detection mechanisms because they offer notification as quickly as possible when a network anomaly, intrusion, or other malicious event occurs. Network threats mutate quickly. Worms spread almost instantaneously. The threat of *zero-day* attacks will not allow for weak incident management prevention and detection processes. Effective incident management tools and processes ensure quick reaction and recovery when an event does occur.

The syslog protocol provides a transport mechanism that allows devices to send event notification messages over the network to syslog servers. These servers are often simply message collectors that don't return any acknowledgement.

The syslog protocol is very simple. The sender transmits a text message that is less than 1024 bytes. The syslog server (often referred to as syslogd or the syslog daemon) appends the message to the file. These messages can be transmitted using either UDP or TCP. Normally, syslog data is transmitted as plaintext, but there are tools that use an SSL wrapper to add encryption for increased security.

Although TCP can be used, syslog doesn't require a three-way handshake. Given the small size of the messages, UDP port 514 is the most commonly used communication. As UDP is connectionless, no acknowledgments are provided. At the application layer, syslog servers normally don't send any acknowledgments back to the sender either. Thus, devices transmitting syslog messages never know whether the syslog server has received the messages. Most sending devices will send syslog messages even if there is no syslog server in place.

Syslog packets are limited to 1024 bytes and carry the following information:

- Facility
- Severity
- Hostname
- Timestamp
- Message

Syslog messages are categorized based on the generating source. These sources can be the OS running a device, a syslog process (or service), or an application. To learn more about syslog, see the IETF document at http://www.ietf.org/internet-drafts/draft-ietf-syslog-protocol-19.txt for further technical details.

Integrating Tools for Event Correlation

An NMS is a combination of hardware and software used to monitor and administer the addressable and manageable elements of the network. In converged service networks, VoIP and video services introduce a new set of manageable network elements that perform telecommunications service functions. These elements typically include gateways, call management servers, emergency responders, voicemail servers, media gateways or servers, and so on.

General network management involves functions such as network planning, traffic routing, user authorization, configuration management, fault management, security management, performance management, and accounting management. Many protocols exist to support network and network device monitoring and management. As we discussed, SNMP is a common network protocol, but others that may come into play include Common Management Information Protocol (CMIP), Web-Based Enterprise Management (WBEM), Common Information Model (CIM), Transaction Language 1 (TL1), and Java Management Extensions (JMX). We won't probe these protocols in depth here.

When implementing the converged network, NMSs take on a new, crucial role in enterprise service delivery. Enterprises need to bolster their management capabilities to test and manage QoS, performance, and availability in performance metrics, especially with VOIP services. To get started, companies should analyze their business requirements and determine key performance and QoS metrics.

A comprehensive, enterprise-wide data collection mechanism is required to provide effective service assurances. Collecting as much data about the network as possible will aid in the ability to ensure call quality and consistency of service.

Network Monitoring

An NMS constantly monitors and notifies the network administrator via email, pager, or other alarms in the event of outages or anomalies that exceed defined thresholds. Monitoring is vital to service assurance and VoIP management. An NMS continually monitors the network for problems that result from overloaded and/or crashed servers, network connections, or other devices. For example, to determine the status of a Web server, monitoring software may simply ping the server periodically to check for a response. A more comprehensive NMS technique is to send an HTTP request to fetch a specific Web page; testing email servers might involve sending a periodic test messages to ensure the email services (SMTP, POP3, and IMAP, for example) are up and running properly.

Status request failures—like those found when a *ping* fails, the Web page can't be retrieved, or another unexpected condition is encountered—can be configured with most NMS platforms to activate some predefined response. These responses can vary from event to event. In some cases, an alarm might be sent to the systems administrator's email, pager, or mobile phone so that human intervention can follow. Highly evolved systems might trigger some automatic failover system mechanism for continuity of operations. Or a non-critical server experiencing problems might simply be removed from service until a suitable time is available for repair.

Some of the most important characteristics of network elements monitored in the IP network include CPU utilization, physical memory, disk space usage, virtual memory, and fans and power supplies. Many systems monitor temperature to ensure a proper operating environment is maintained. Monitoring of system backups is incorporated to ensure positive confirmation that backup jobs run as scheduled. Many organizations monitor Web server software (typically Apache or Internet Information Services—IIS), directory services systems, and Domain Name Service (DNS) servers. Security monitoring is often incorporated into the monitoring performed in this network operations center environment.

Managing VoIP services raises the need to monitor both the VoIP service elements and QoS facets of network performance to ensure acceptable call quality. In traditional IP networks, the infrastructure elements are monitored. VoIP introduces new infrastructure elements including voice processing systems, signaling servers, gateways to other networks, border controllers supporting SIP trunking, and voicemail systems.

VoIP Service Elements to Monitor

Voice traffic carries a set of performance expectations users have come to expect through years of telephone use. VoIP services introduce a new range of network elements to monitor. Whenever a device (for example, phone, gateway, and gatekeeper) registers with the network, there will be an auditing entry to review. Problems with device registration, for any reason, can impact service availability. You'll want to be alerted when the number of registration attempts or failures exceeds predefined thresholds. If the number of registered telephones changes dramatically, it could be a signal that there is a problem with the VoIP network. Gateway registration monitoring will help identify new or missing gateway servers.

Call monitoring isn't eavesdropping on individual calls. It's really call-traffic monitoring. It involves monitoring incoming and outgoing call volumes to identify failures. If your VoIP system supports fax calling, attempted fax calls also need to be monitored. Call monitoring typically focuses on four specific areas:

- Calls in progress—When a VoIP phone goes "off hook," a call is deemed in progress until it goes back "on hook." If every call in progress connects successfully, the number of calls in progress will equal the number of active calls. When designing the VoIP network, you'll need to establish an upper-limit threshold for the number of calls that can be in progress at any given time.
- Active calls—Active calls have successfully connected a voice path. Again, when designing the VoIP network, you'll need to establish an upper-limit threshold for the number of active calls that can be handled at any point in time.
- Attempted calls—Designers strive to ensure that all calls attempted will be completed successfully, but such isn't the case in the real world. Monitoring calls attempted over time yields data that aids in identifying peak periods and the busy hour call attempt (BHCA) value.
- Completed calls—A completed call is any successful active call that completes without an abnormal termination code. Monitoring completed calls over time is also useful in identifying peaks periods and the BHCA value.

VoIP services need to interconnect to the PSTN through gateways. In addition to gateway monitoring, it is vital to monitor the PSTN side of the VoIP service network. PSTN connections are frequently established using ISDN Primary Rate Interface (PRI) channels over T-1 circuits. Monitoring active PRI channels, especially over time, can help identify call patterns and busy hour peak call volumes. Baseline data can also be used to identify underutilization of circuits. Data trending helps in capacity planning and the growth and maturation of the VoIP service.

One benefit in deploying VoIP services is the conference bridging capabilities. If your deployment supports conferencing, you must configure the maximum number of audio streams that will be supported. Monitoring will ensure that the number of available audio streams meets acceptable service levels for your organization.

IP phone functionality requires continual monitoring for service assurance. You should monitor IP phones for their registration status, the validity of their dial tones, jitter, latency, and lost packet count. These QoS parameters directly affect service delivery.

Monitoring Bandwidth and QoS

Voice traffic requires specific bandwidth based on the codec used in the VoIP design. G.711 requires about 64Kbps for each direction of a bidirectional call. G.723 and G.729 require significantly less bandwidth due to compression, but congestion can severely impact call quality.

When you add applications to your network, there is always a risk of oversubscribing links. Oversubscription leads to congestion, and congestion may introduce a negative impact on call quality. Packet loss and increased latency are common side effects of congestion and can, when left unchecked, render VoIP services unusable.

For VoIP users to receive an acceptable level of voice quality, VoIP traffic may need to be given some kind of prioritization over other kinds of network traffic, such as data. The main objective of QoS mechanisms is to ensure that each type of traffic—data, voice, and video—receives the preferential treatment it deserves, thereby reducing or eliminating the delay of real-time streaming voice or video packets crossing the network.

The following list highlights examples of metrics that are frequently monitored because of their effect on VoIP call quality:

- Delay or latency is an estimate of the network delivery time expressed in milliseconds. It’s measured as an average value of the difference between the timestamps noted by the senders and the receivers of messages. It is measured when the messages are received. The end-to-end delay, or latency, measured between endpoints is a key factor in determining VoIP call quality.
- Jitter is also called delay variation. It indicates the variance of the arrival rate of packets. Jitter points directly to the consistency or predictability of the network. It is a call quality factor known to adversely affect call quality. Networks can compensate for jitter by implementing jitter buffers to normalize the timing of the traffic flow. Jitter buffer loss occurs when jitter exceeds that which the jitter buffer can hold. Jitter and jitter buffer loss affect call clarity, which affects the overall call quality.
- Packet loss indicates a packet lost during transmission. In VoIP, this could mean the loss of an entire syllable or word during the course of a conversation. Obviously, data loss can severely impair call quality. Monitoring systems measure the number of packets that were expected against the number actually received.
- Mean Opinion Score (MOS) is a subjective measure used in voice telephony, especially when codecs are used to compress the bandwidth requirement of a digitized voice connection from the standard 64Kbps PCM modulation. MOS is generated by averaging the results of a set of standard, subjective tests. In the past, a number of listeners rate the heard audio quality of test sentences read aloud by both male and female speakers then rate each as follows: 1-bad, 2-poor, 3-fair, 4-good, 5-excellent. The MOS is the arithmetic mean of all the individual scores. In current systems, MOS is often determined through software algorithms.

Measurements and Metrics for Voice Quality

Voice quality measurement as part of operational monitoring can be either non-intrusive or intrusive. Non-intrusive tests are typically based on actual voice conversations taking place during daily operations, whereas intrusive testing requires placing test calls across the network.

One approach to evaluating call and voice quality is to assemble a group of participants who will act as judges. A common technique is to have them listen to test calls, and assign scores from 1 to 5, much like the MOS evaluation testing. There are a number of algorithms and methods that might be used, including MOS, Perceptual Analysis Measurement System (PAMS), Perceptual Speech Quality Measurement (PSQM/PSQM+), and Perceptual Evaluation of Speech Quality (PESQ/PESQ-LQ):

- MOS has been adopted from the PSTN and traditional telephony. This historical measure of voice call quality was judged in the US by eight men and eight women who rated voice quality on a scale from 1 at the low end to 5 at the high end. Although this human evaluation of MOS was useful for identifying the quality of experience with calls in the PSTN, it hasn’t proven useful for network monitoring or SLA compliance measurement.
- PAMS was developed by British Telecom. It doesn’t replace MOS but offers a reasonable first attempt at automating the MOS scoring. PAMS compares the original analog voice wave with reproduced speech using a complex weighting method that purports to take into account specific characteristics that are important to the human ear. The PAMS scale is from 0, which represents a perfect match between samples, to 6.5. PAMS values estimate MOS scores with a ±10 to 20% level of accuracy that is inadequate for use in VoIP networks. PAMS values are excellent for benchmarking purposes.
- PSQM/PSQM+ is defined in ITU-T Recommendation P.861. PSQM estimates MOS with a greater level of accuracy than PAMS at ±10%. PSQM and PSQM+ are also good for benchmarking and comparisons. They represent improvements on the earlier PAMS algorithm. PSQM and PSQM+ use the same 0 to 6.5 scoring that PAMS uses.
- PESQ/PESQ-LQ is defined in ITU-T Recommendation P.862. Like other techniques described here, PESQ doesn’t replace MOS. It’s proven excellent for benchmarking and comparisons and is an enhancement and improvement to the PSQM algorithm. It also uses the same 0 to 6.5 scoring.

The Do-It-Yourself Approach

Many enterprises create their own management platform suite over time. For many small to mid-sized businesses, a managed service has always been the primary option. These organizations are often resources-constrained and simply don't have staffing to do everything themselves. Larger enterprises often take exactly the opposite approach. These organizations have provided their own telephony services for many years and seem inclined to continue this approach in the emerging multimedia converged networks of today. Although many enterprises view VoIP and video as new service applications on the IP network, the trend seems to be a continued do-it-yourself approach to data, VoIP, and video. Although this is common, it isn't necessary. For many large enterprises, migration to a converged service network may present a perfect opportunity to rethink service delivery and develop partner relationships with service providers for both delivery and management of these services.

The enterprise network is evolving and becoming somewhat organic in nature. This converged network provides a shifting set of real-time, near real-time and non real-time voice, data, and video services. Even an internal SLA for workgroups inside an enterprise needs to be developed to support metrics that are relevant to each of the services provided.

The SLA

An SLA is essentially a contract mechanism that documents the level of service that a customer should expect to receive. An SLA that has been thoroughly thought out will also describe actions that will be taken by the service delivery organization. This is where different service classes will be defined and delivery characteristics identified for each.

Data, voice, and video services will be differentiated and require different traffic/service flows. Thus, the network provides differing Classes of Service (CoS). To assure service delivery, management, monitoring, and analysis tools must be capable of monitoring the service parameters for each service call. The alternative is to implement discrete tools for each service time. This option is often unwieldy and expensive.

SLA Metrics to Consider

The SLA process leads to an informed user within the enterprise. Users will begin to understand multimedia services much the way an automobile purchaser understands performance characteristics of different cars. The MOS, PQSM, R-Value, and other performance metrics described earlier might become as familiar to the integrated services users as MPG, horsepower, and such are to motor vehicle owners. Embedding the right tools within the service delivery organization that support comprehensive monitoring, measurement, and analysis of multimedia traffic compared with established baselines are crucial service delivery support tools.

It’s important to look beyond the metrics that might be specific to VoIP and video. You must also monitor more traditional aspects of the network including loss, delay, jitter, and general availability. These measurements give a common basis of comparison across data, voice, and video services.

In a collaborative network environment between end user workgroups and the service delivery group, the SLA can be used as a tool for continuous improvement as metrics evolve over time to reflect actual network requirements and performance capabilities. The right tools for monitoring and managing your evolving corporate multimedia networks move far beyond the mindset of the Internet as “plumbing” or a simple common architecture for delivering any traffic type. Corporate management must understand, and appropriate necessary funding to support specialized tools to maintain the health of the converged network. Without proper funding of personnel and systems, the IP network may fail with dire consequences on day-to-day operations.

Beyond Implementation—Operational Support

One of the implementation dangers in integrating network services is the inclination to reduce staff without necessary skills transfer. Organizations that are overly focused on cost reduction run the risk of embracing staff cuts followed by the creation of services that remaining staff can’t effectively support. Many organizations have moved forward on the voice integration path with the IP network staff leading the initiative. This can lead to gaps in vision of telephony services and elimination of the traditional telecom support team. This is a risky proposition and the business impacts should be considered at every turn.

It’s surprisingly easy to overlook specialized skills that enterprise telephony engineers bring to the table. IT network engineers often don’t fully appreciate the workings of Erlang-B calculations and the importance of traffic engineering. These technical resources are also the ones that design interoffice trunking facilities, understand call center requirements, and build the organization’s automatic call distribution groups, hunt groups, and call pickup groups. It’s important to remember that although providing these services over IP may be a new approach, these are still core voice services necessary for daily business.

It’s important to partner voice and data specialists in support of new unified communications service. This approach will let you make the best use of every technical resource within the organization. Don’t make the mistake of allowing key institutional knowledge about your services, your network, and your requirements escape through oversight while focusing on cost reduction. In the unified communications network, reductions in cost are won over time. They may not come immediately.

Even Help desk support for telephones is generally viewed as something simple. We assume that everybody knows how to use a phone. In the integrated data, voice, and video network, simple telephony features may well be delivered like they were in the past. It's more likely that the integrated system will bring new facets of data workstation feature management into play. For example, one common feature in VoIP systems is to provide a way for users to manage their own telephone features. Button configuration and speed dial lists may now be managed via a Web interface from the desktop. Some method of oversight, whether through management and monitoring or through specialized vendor-provided tools, will help simplify management of even the simplest changes.

Remember that services converging onto a single infrastructure also mean the corporate Help desk may become the primary support point for all new services. They may begin receiving questions they've never dealt with previously. They too may be unfamiliar with the new integrated services. It's crucial that all support staff, including the Help desk, get the training and management and monitoring tools they need to support daily business operations effectively.

Pros and Cons of Rolling Out Your Own Management Platform

There are distinct advantages in custom-building an enterprise-specific management platform. When an enterprise opts to utilize a suite of managed services, the service provider has some commitment to meet established expectations. In the do-it-yourself approach, you become your own service provider. You might be an internal service provider delivering services to internal departments, divisions, regions, and employees, but end user expectations remain the same. An organization embracing the do-it-yourself approach must implement not just the services but also the tools for measurement, monitoring, and analysis that will ensure service availability.

Vendors and manufacturers often provide the most accurate, most granular, and what might appear to be the most desirable systems embedded within their solutions. These integrated service hardware and software solutions commonly provide a combination of syslog and SNMP-like monitoring. These tools can provide both real-time monitoring and event correlation capabilities.

Many vendors also provide product-specific, proprietary tools to aid in management and monitoring. It's important to identify these tools early in the process. Some may be provided freely as part of the product suite. Some vendors will offer yet another management suite to ease the process. Again, it's important to recognize the costs of adding more management tools to a service that seems to be growing in complexity daily.

Third-party tools for modeling, monitoring, measuring, and managing are widely available and can provide incredible value to service managers. Some provide broad visibility across a variety of platforms, bringing insight into the network performance aspects that far exceed human capacity to observe or analyze. Some of these third-party tools are commercial products. Some may come from open source libraries.

On the plus side, there is a wealth of available resources. On the negative side, the onus is now kept within the enterprise to select the right tools to manage, monitor, and analyze every facet of the services provided to ensure availability and sustainable service quality.

Freeware and Open Source vs. Commercial Products

When taking the do-it-yourself approach, don’t overlook the sophisticated tools already embedded in the routers, switches, gateways, end systems, and other components of the integrated network. There is a feature-rich toolset built-in to the products your organization will use.

Although manufactures may provide the most accurate tools embedded within their hardware, they may raise a different area of concern. The tightly coupled vendor solutions provide great management granularity and detailed analytical capacity. However, unless the entire network and all services are provided by a single manufacturer, this granular approach may be unable to provide a holistic view of broad services. Manufacture-provided, proprietary tools certainly have their place. They often prove invaluable when optimizing a network or when troubleshooting specific system problems. But they may lack the broader, more standardized and less platform-specific capabilities of systems provided by third parties.

Open source and freeware solutions present a unique set of challenges. Although the price may be right (they’re often free), there may be no support available. This can lead to enterprise engineers providing their own support, developing their own patches, and creating their own performance tweaks. In time, open source software and freeware can morph into a variation on in-house software development.

Another consideration may be the security of these solutions. There are two considerations for open source solutions. Although an enterprise might build a level of trust for a vendor providing commercial products, freeware, and open source tools essentially come from often unknown sources. In many cases, this code has been vetted by hundreds or thousands of well-qualified independent developers. The open availability of the source code means that many people can contribute, each incrementally improving and crafting the code. However, whether appropriate secure coding practices were used in their development may remain completely unknown. And while the danger of these tools being purposefully developed with malicious code elements nested inside has not proven common, these are management platform tools that support the daily enterprise operations. Can you trust the “crown jewels” of the enterprise service network to unsupported software written by unknown third parties? It’s important to weigh the value against any and all potential risks.

Integrated Commercial Solutions

An organization that selects a managed service approach may be able to ignore many of the technical details of the inner workings of service delivery systems. These enterprises may simply need to “check the pulse” and “take the temperature” of the integrated services network. Often service providers will either provide access to these tools or the tools themselves as part of the managed service offering. These tools may simply provide a snapshot of the network status at a given point in time. Some tools may provide a more detailed, near real-time view of the health of the overall service. It’s important to review SLAs, as these tools may only provide a customer view that might not correlate directly with service delivery requirements identified in the SLA.

Multimedia networks supporting voice applications are complex systems that require sophisticated monitoring intelligence. Although individual manufacturer's monitoring tools offer a partial solution, they cannot account for the granular nuance of monitoring and managing a complex, multi-vendor service environment. Generic monitoring and management tools are available, but generally insufficient. To truly be effective, third-party tools have to be assembled with some knowledge of each of the individual manufacturers' systems. This may be the only way to obtain a comprehensive, system-wide, end-to-end view of the converged services network. The following list highlights key areas for monitoring and management.

- Pre-deployment simulation
- Manufacturer-specific monitoring and management
- Real-time business views
- Call detail records
- Calls in progress
- Delay-to-dial-tone rates
- BHCAs
- Busy Hour Call Connects (BHCCs)
- Gateway channel utilization and loading
- Real-time call monitoring
- Phone and multimedia device availability and monitoring
- Poorly performing components
- Service level breaches and SLA compliance
- Real-time interface to manager of managers
- Summary and exception reporting
- Utilization trends over time
- Managed devices by company, department, and location
- Asset tracking
- Capacity planning
- Incoming and outgoing calls
- Loading by route pattern, route group, route list, and gateway
- Bandwidth utilization
- Delay and delay variation
- Packet loss
- Route patterns, utilization, and availability

Using the right mix of sophisticated third-party tools will provide a toolset for documenting results consistently while eliminating any vendor bias.

Pros and Cons of a Packaged Management Platform

As earlier noted, when choosing the do-it-yourself approach, organizations will still have access to the sophisticated tools embedded in network equipment. There are still other tools that can provide insight to the big picture of service management and aid in troubleshooting granular problem areas.

Managed service providers will generally allow end-user organizations, customers, to view a thin slice of their network for performance monitoring. These views pertain directly to a specific customer organization. Service providers often call this service feature Customer Network Management (CNM), which allows a customer to see their own part of the larger, shared virtual network. This approach has been in wide use for at least two decades. Early CNM systems were periodic summary snapshots showing sets of statistics that bore little resemblance to the real world or the SLAs. Today, providers offer services that are much closer to real-time views, with granularity down to the individual connection. In some cases, they allow customers to adjust certain characteristics, such as CoS or available bandwidth, in real time.

Summary

Without an NMS focused on the integrated data, VoIP and video services, companies will find themselves in the dark with a lack of information to support necessary service assurance. VoIP services that lack management will be prone to service delivery and quality problems that cannot be tracked to any specific network elements or service delivery metrics. The appropriate QoS and network performance metrics require constant management, monitoring, and analysis to ensure acceptable service delivery.

Integrated services introduce new complexities and opportunities to simplify the process of moves, adds, and changes to the network. In the past, employees would often move their workstation easily, using Dynamic Host Configuration Protocol (DHCP) to retrieve data network setup parameters. Telephones have always represented a more complex move process. Requiring PBX reprogramming and intercession by a telephone services administrator; convergence to an integrated IP service network may simplify this process. VoIP solutions may also provide new productivity tools to remote workers, but again, management, monitoring, and analysis of network services is crucial to the continuity of daily business operations.

Strategic new business applications appear on the horizon every day. We see network service like voice and video beginning to couple tightly with enterprise resource planning and customer relationship management systems. With a better understanding of how management and monitoring play into the life cycle of integrated services, we have the foundation to move forward to the next chapter, in which we'll explore managing service availability and capacity planning for converged services.

Download Additional eBooks from Realtime Nexus!