Reliability of Emerging Internet-Based Services

By

Dr. H. Eslambolchi, Ph.D., M. Daneshmand, Ph.D.,

Table of Contents

Abstract

The Internet:
"No one owns it. And no one in particular runs it. Yet more than half a billion people rely on it as they do a light switch"

- The New York Times (January 10, 2002)

In this paper we introduce a systematic approach for monitoring, quantification, and improvement of the reliability of emerging Internet-based services. Our focus is on the end-users' perspective of the reliability - the way the "half a billion people" view the reliability of their Internet services. We begin with the definition of reliability - as the US Department of Defense in 1952 first originated it. We expand the definition and formulation, initially intended for measurement of reliability of relatively simple electronic devices of the post war, to the measurement of end-users' perception of reliability of the relatively complex Internet services of the new century. We integrate the reliability concepts developed by the computer scientists and Internet inventors over the past 10 years, with those developed and practiced by the telecommunications industry over the past 50 years. We build the Internet service reliability based on the rich experience of the telephone service reliability. We show how users' experience of quality of an Internet service can be mapped into a single Defect Per Million (DPM) scale from which service providers may monitor and improve the reliability of their offerings. In particular, we provide practical guidelines to measure and monitor the reliability of emerging Internet-based applications such as e-mail, and web applications for e-commerce/e-business. We demonstrate how DPM monitoring and trending can help service providers to execute the P3 paradigm of Predictive, Proactive, and Preventive in order to meet the service reliability expectations of their end-users.

1. Introduction

There are more than 150 million hosts on the global Internet. The number of Internet users is over 500 million, worldwide[1]. The Internet has been doubling annually since 1988. The size of the global Internet is estimated to exceed the size of the global telephone network by 2006. It is estimated that e-commerce on the Internet will reach somewhere between $1.8T and $3.2T by 2003[1]. There is a need for a reliable and secure Internet. Users expect the Internet to be accessible, reliable and secure all the time.

Reliability and security continue to be the number one concern of the users of Internet-based services. Recent surveys and focus group studies indicate that security and reliability continue to be the two most serious limiting factors in Electronic-Commerce/Electronic-Business (EC/EB). Users consistently rank reliability and security as "very important". Internet businesses lose billions of dollars each year because of slow and failed web services. As the e-business/e-commerce grow exponentially so does the reliability and security concerns of both service users as well as service providers. With stock, real state, and many other online commodities now selling for thousands of dollars, a more reliable Internet exchange are becoming increasingly important.

In this paper we focus on the reliability aspects of the Internet-based telecommunication services. In particular, we demonstrate that reliability is a stochastic concept and show how Defect Per Million (DPM) approach can be used to quantify end-to-end users' perspective of the reliability of complex IP-Based applications. We develop a mathematical structure from which service providers can measure and monitor the reliability of web applications as being experienced by the end users and conduct near real time root cause analysis of any reliability degradation. We build on the P3 (Predictive, Preventive, and Proactive) paradigm and show how Internet service providers can execute the P3 and reduce DPM (Defect Per Million) and gain competitive edge.

In section 2 we provide a short historical background as well as initial definition of reliability and expand this definition to include reliability of Internet-based services. Section 3 provides a short description of components of the global Internet and World Wide Web (WWW) supporting the end-to-end Internet services and applications. Section 4 describes and classifies real-time and non-real time Internet-based applications. The reliability framework of Accessibility, Continuity, and Fulfillment is introduced in section 5. Section 6 covers the fundamentals of DPM (Defect per Million) and provides a step-by-step practical guidelines for DPM deployment. Finally, we demonstrate the application of the methodology to major Internet services of "email" and "web-hosting/e-commerce" as well as Intelligent Content Distribution (ICD) in sections 7 and 8.

[1] Internet Engineering Task Force (IETF), RFC2026, January 1, 2002.

2. Reliability Definition

The increasing complexity of military electronic systems was generating failure rates, which resulted in reduction of the availability and increased costs. In 1952, the US Department of Defense and the electronics industry jointly set up the Advisory Group on Reliability of Electronic Equipment (AGREE). The AGREE report [1] defined the reliability and provided scientific guidelines to quantify, monitor, and improve reliability of electronics products. In 1978, American National Standard Institute (ANSI) and American Society for Quality Control (ASQC) expanded the definition to include all products and services [2]. The classic definition of reliability provided by AGREE, ANSI & ASQC is:

2.1 Classic Definition of Reliability:

Reliability is the *probability* of a *product* performing without *failure* specified function under given conditions for a specified period of *time*.

The definition specifies six key components of reliability that need to be considered in any reliability management program:

1. Probability
Failure occurrence is a stochastic phenomenon and therefore reliability is quantified in terms of probability,

2. Product
Reliability is measured for a specific product (service), and thus, a proper specification of the product or service under consideration is required,

3. Failure
Each usage of the product or service would lead to either a successful or failure performance. There is a need for a clear and unambiguous definition of a "failed" or "successful" usage,

4. Function
The function that the product or service is required to perform must be specified. The specification should include perspectives of both "provider" and "user",

5. Conditions
The environment and condition under which the product/service must perform the required function need to be specified,

6. Time
The period of time during which the product/service is expected to perform the required function without failure need to be determine.

2.2 Reliability of IP-Based Services:

The classic definition of reliability does not readily expand to the IP-Based telecommunication networks and services. An IP-Based telecommunications network is a complex system and thus specification and definition of the six key components of reliability continues to be a challenging task:

1. Probability

For simple systems, calculations of the probability of performing without failure are relatively simple. However, calculations of the probabilities of failure of telecommunication products or services are extremely complex. In current years there has been a concerted effort to develop methodologies for quantification of reliability and security of emerging Internet-based services. In this chapter we will introduce a generic framework for quantification of reliability of various IP-Bases services. Our generic framework will naturally lead to DPM (Defect Per Million) metric adopted by the Standards [3] bodies and currently being used within the industry to monitor, diagnose, and resolve reliability degradation of the IP services.

2. Product/Service

Numerous IP-Based services exist and are being developed. These services include non-real time applications such as Email, E-Commerce/E-Business, Web-Hosting, Media Streaming, Intelligent Content Distribution (ICD), and real-time applications such as Voice and Video. The reliability requirement of each of these services/applications is different.

3. Failure

Successful performance varies from service to service. Further, for a given service, the user's perspective of a successful performance could be different than that of the service provider.

4. Function

Clearly different services are required to perform different functions, and thus, measurement of reliability would depend on the specific IP-Based service.

5. Conditions

Each Internet service is required to function under its own pre-specified condition. These conditions on an operational profile of a service feature.

6. Time

The time period over which the service reliability needs to be measured has to be specified. Depending on the application the time period may range from a 15 minutes to hourly/daily/weekly intervals.

It is seen therefore, that separate service reliability programs may need to be deployed for separate services. Further, a successful reliability program must carefully specify the

service for which the reliability will be measured, the unit to be observed, the function to be performed, the condition under which the service need to perform, the time interval over which reliability will be measured, and a clear criteria for defective/effective performance classification. In short, an IP-Based service reliability program requires meeting of the minds on:

(1) reliability as a probability measure or DPM (Defect Per Million) metric,
(2) the service for which the reliability is quantified,
(3) the environment and condition under which the service must perform,
(4) a definition of a successful ("effective") service performance.

In this chapter we define and formulate reliability of several IP-Based services.

Reliability of an Internet service is quantified in terms of the probability of failure/"Defective" (or success/"Effective") of user's "transaction". A transaction will be successful ("effective") if and only if all the systems/subsystems that the transaction touches during an Internet application operate in harmony and remain operative at the required level of performance during the whole period needed to complete the transaction. Otherwise, the transaction will be classified as "defective".

The reliability of an Internet service, therefore, is intertwined with the reliability of a collection of systems/subsystems that a user's transaction touches during an Internet service application. These include desktop, browser, access infrastructure, Internet Service Providers (ISPs), backbones, peering, application and content servers, databases, gateways & gatekeepers, and other hardware/software that run these systems. In this chapter, we introduce methodological foundations necessary for quantifying and monitoring the reliability each of the Internet services. Further, we demonstrate how these techniques will naturally lead to enhancement of service reliability as being experienced by the users, and meeting customers' end-to-end reliability requirement.

3. The Internet

In this section we provide a short description of components of the global Internet supporting the end-to-end Internet services and applications. Our focus is on components that impact the end-to-end performance and reliability of Internet services and applications as being experienced by users.

3.1. The Internet Infrastructure

The Internet is a conglomeration of thousands of interconnected networks. It is a medium for information exchange between computers all around the world. Consumers use Internet to exchange electronic mail (email), pay bills, buy and sell stocks, shop and conduct research. Businesses use Internet to sell, place orders, train employees, conduct web conferences/meetings and provide customer service. Voice, video, music, and fax are transmitted over the Internet. A high-level Internet infrastructure is shown in Figure 1 below.

Backbone

The interconnected networks communicate with each other over a suite of standardized protocols called Transmit Control Protocol/Internet Protocol (TCP/IP). TCP/IP breaks up the data into "envelopes" called "Packets". Packets are the fundamental unit of transmission in any Internet application. The networks transmit Internet traffic, packets, over thousands of interconnected backbones. Backbones, each operated under distinct administrative control, consist of high-speed routers and links that transmit traffic at gigabit speeds. The different carriers that operate Internet backbone, exchange traffic with each other at *metropolitan area exchanges (MAEs)* and *network access points (NAPs)* also known as *peering*.

Peering

Different Internet networks exchange data at network access points (NAPs) called "peering" sites. NAPs enable users and sites on different networks to send and receive data to and from each other. These NAPs are distributed worldwide in major metropolitan cities including Washington, D.C., Chicago, San Francisco, Dallas, London, Amsterdam, and Frankfurt. Peering sites are operated by commercial organizations. Internet Service Providers (ISPs) and Network Service Providers (NSPs) lease ports at peering sites and connect to their routers. Different peering arrangements are made between different network operators, meeting certain levels of reliability, service quality, and exchange speed.

Network Service Providers (NSPs) own and operate high-speed links that exchange data at peering sites and make up the Internet. At the peering sites, routers transfer messages between backbones owned and operated by many network service providers. To prevent traffic congestion at major exchange centers, network service providers have arranged private peering sites. Private peering sites are direct exchange places that carriers agree on many aspects of the data exchange including the performance and reliability related parameters of delay, service level, and amount of the data to be transferred. The reliability can be monitored more closely at the private peering exchanges.

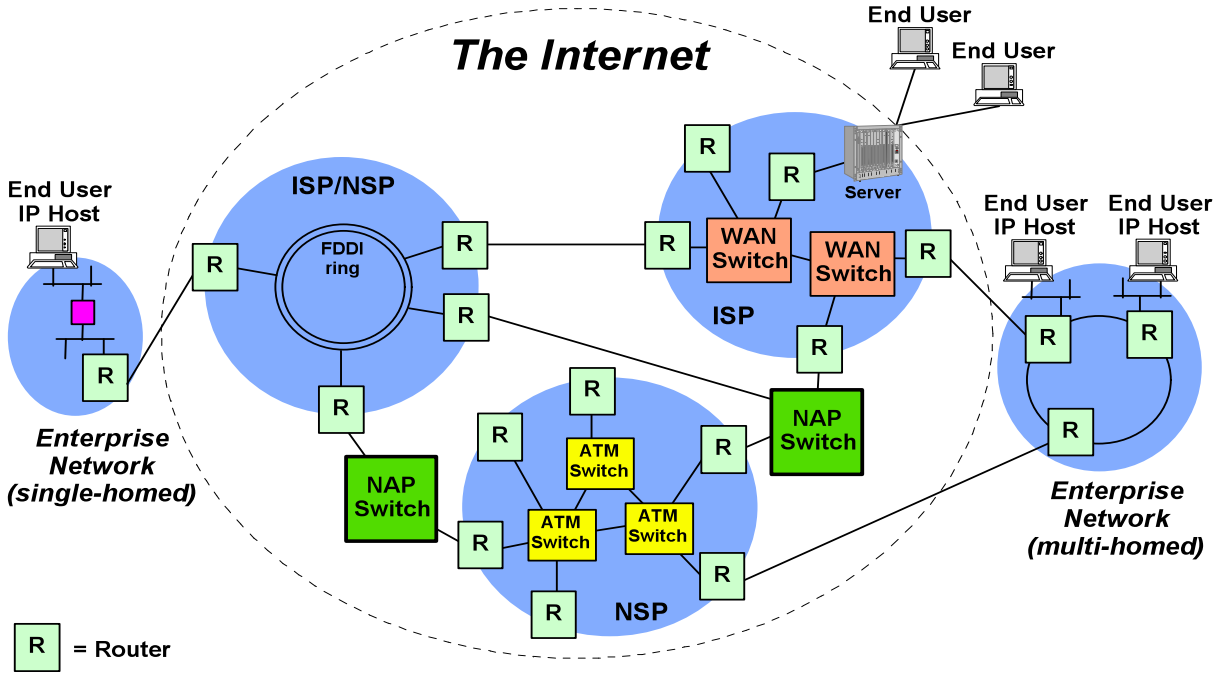

Figure 1. A High-Level Architecture of Internet Infrastructure

3.2. The World Wide Web

The World Wide Web (WWW) is a vehicle for multimedia presentation of information in terms of audio, video, and text. It enables users to send and hear sound, see video, colors, and graphical representation. The basic idea is that and Personal Computer (PC) should be able to find information without needing to know a particular computer language. As long as they use WWW browsers (such as Netscape or Internet Explorer), all personal Computers (PCs) are compatible with the Web. The advent of the WWW (1989) and browsers (1993) created the point-and-click environment by which users can easily navigate their way from computer to computer on the Internet and access information resources.

Browser

Browser is a program that requests and receives the information from the web and displays it to the user. Browsers installed on PCs provide simple interactive interfaces for the users to access and navigate the web. Browsers are easy to use and have facilitated public use of the Internet.

Client - Server

The World Wide Web is based on client -server model. The client is a service-requesting computer, usually a PC, which generates request, receives and reads information from the service-providing computers (servers) within the Web. Clients utilize browsers to communicate with servers through the Internet and create graphical interfaces to access and navigate WWW.

The web server is the service-providing computer on which the web information is located. The server provides service to clients. It receives, processes, and responds to the clients' (users') requests. The server locates the requested information and copies it to the network connection to be sent to the client. The server needs to meet the requirement of the application it is serving. Depending on their function servers may be classified into "web server", "application server", and "database server". The application server works in conjunction with the web server. It may process the user's orders, and check the availability of desired goods. The database server provides access to a collection of data stored in a database. A database server holds information and provides this information to application server when requested.

3.3. Internet Access

Individuals, homes, small-office users, corporations, and institutions connect to the Internet via many types of telecommunication services. These include analog dial-up lines, Integrated Services Digital Network (ISDN), Digital Subscriber Line (DSL) services, Cable Modem (CM), T1, T3, and most recently wireless access services. Dial-up, ISDN, and DSL access utilizes the exiting use Public Switched Telephone Network

(PSTN) infrastructure. Cable Modem access uses the cable TV network or CATV (community area TV) infrastructure. T1 and T3 access are private (dedicated or leased) connection lines.

Home or small-office consumers connect to the Internet using a web browser running on a PC that connects through a modem to the local ISP. On the other end of the line, ISPs aggregate traffics from many users and send it to the Internet backbones over high-speed lines. Corporate consumers use browsers running on computers in a local area network that connects to the Internet through routers and purchased or leased permanent T1/T3 links. In any event, the dial-up or dedicated connections to the Internet via network service providers (NSPs) or ISPs facilitate access to the information resources of the worldwide Internet consisting of thousands of separate and independent networks and millions of individual computers. The connection speed depends on the access technology used and it has essential impact on the user's expectation of the performance as well as the user's perception of an "effective" (or "defective") Internet experience.

Dial-up Access

Dial-up access utilizes PSTN's existing infrastructure in analog (copper) loop. Dial-up users, program their PCs to dial the local telephone number associated with their ISP. The call is routed to the local telephone company's central office, which in turn sends the call to the central office connected to the subscriber's ISP equipment. The connection to the local central office is through two modems between which an analog "voice call" has been established. ISPs have banks of modems in remote access server devices that handle calls received from Telephone Company. The ISP has direct high-speed links to a regional data center that is connected to the Internet. Today's low-speed dial-up phone line modem connections provide speeds up to 64 kb/s (kilo bits per second) and constitute majority of homes and small businesses Internet access connections. Dials up accesses impose challenging performance problems because of their variable and asymmetric bit rate.

ISDN Access

ISDN utilizes a technology known as Digital Subscriber Line (DSL) that provides a higher speed and a more stable connection between user and PSTN over the local loops. A simplified view of the ISDN connection between a user and local exchange is the network termination equipment that provides ISDN modem functionality at either end of the traditional analog local loop. The digital link between network terminations is always on. ISDN provides point-to-point connectivity in terms of one or more 64 kb/s. Typical ISDN Internet access provides speeds up to 128 kb/s.

DSL Access

A new DSL technology, generically called xDSL, offers higher-speed Internet access to homes and businesses. A popular version of the DSL-based access technology is ADSL (Asymmetric DSL). A distinguishing feature of the ADSL service from the DSL-based ISDN is higher speed. ADSL is capable of delivering up to 9 million bit/s downstream (to the subscriber site - downloading) and up to 1.5 million bit/s upstream (from the subscriber site - uploading). The actual speeds depend on the quality of loop over which the ADSL modems operate - longer or older loops lead to lower speeds. The number of high-speed subscriber lines has increased from 1.8 million to 3.3 million users by the end of the 2001 (Yankee Group).

Cable Modem Access

In recent years the cable industry has been working to leverage the cable TV network infrastructure for residential and business high-speed Internet access. A cable TV network, initially designed for TV signals, utilizes fiber optics from the "head end" to the neighborhood and coaxial cable for the last hop into the customer premises. The hybrid fiber coax (HFC) technology provides an alternative high-speed Internet access. Using cable modems, subscribers connect their PCs to the cable TV network for high-speed full-time (always on) Internet access. Cable modems technology provide high-speed data transfer rates of up to 10 million bit/s depending on the number of users and the particular cable modem configuration. By year's end (2001), cable companies are expected to have 7 million residential broadband subscribers, up from 3.7 million in 2000, according to the Yankee Group.

The high-speed DSL access or CM accesses are generally called broadband services. According to the Yankee Groups[5], by January 2002 some 10.7 million of the US households have a broadband service or about 16% of all households online.

Dedicated Access

Large and medium businesses with high volume of traffic use dedicated, private lines to access the Internet. The ISP arranges for the dedicated T1/T3 line to run from the customer site to the ISP point of presence (PoP). T1 and T3 speeds are 1.54 million bit/s and 44 million bit/s respectively.

The end-to-end users' experience of reliability and performance of a given Internet application depends on reliability and performance of a number of distinct networks, and their associated hardware/software, systems, and subsystems that make up an end-to-end Internet application. Furthermore, different operators often administer the Internet distinct networks. An end-to-end Internet application therefore is impacted by a number of Autonomous Systems (AS) in the Internet. Figure 2 below shows a high-level architecture of these autonomous systems [12]. The challenge is to ensure that all these systems and subsystems function together in harmony and create an "effective" user's experience.

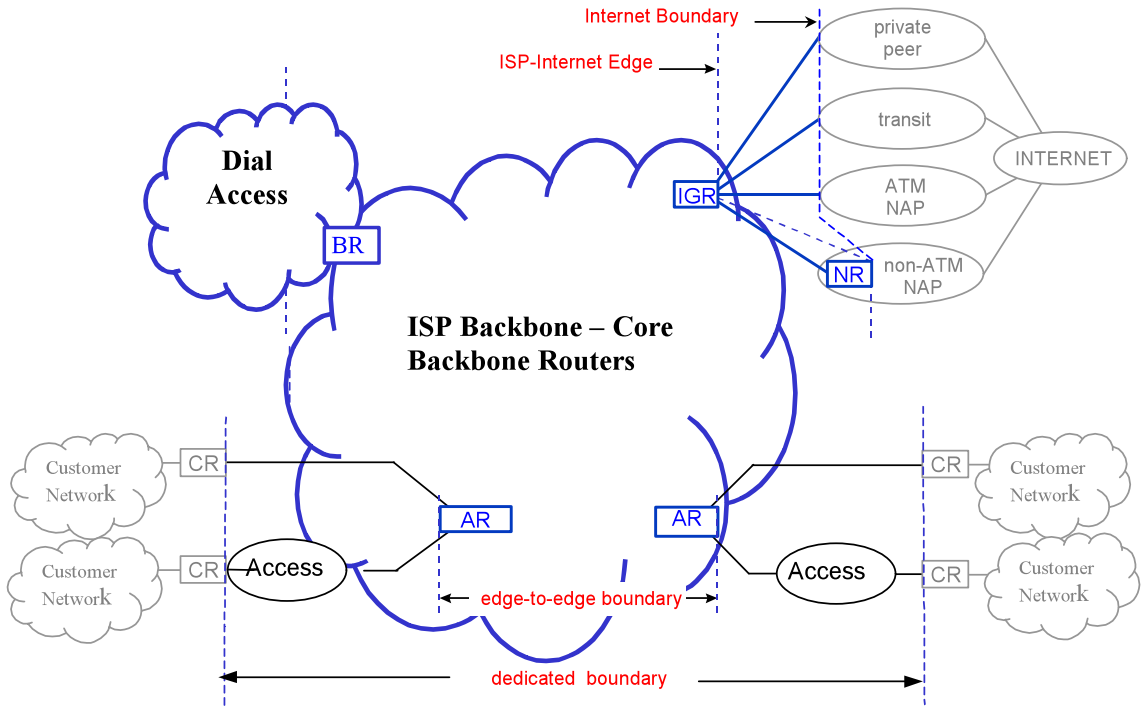

Figure 2 - A High-Level Architecture of Internet Autonomous Systems

4. Internet Services and Applications

The user's experience of an "effective" Internet application varies from application to application. To formulate the concept of reliability from a user's experience we need to understand and consider user's application. Understanding characteristics of applications is an important aspect of identification, definition and formulation of reliability of Internet services. In this section we provide a brief list and description of current services and applications. For the purpose of reliability and performance the Internet services may be classified into two general categories of "real-time" and "non-real-time" applications.

4.1. Real -Time Applications:

As it was indicated earlier, Internet is a pcket-based network. An Internet application digitizes the information into bits, forms a collection of bits into an IP Packet, and transmits the IP Packet from source to destination(s) over the global Internet network. Packets, therefore, are the fundamental units of information transmission in any Internet application.

Real-time applications are those applications in which a distinct notion of timelines and ordering is associated with their packet delivery. Packets must reach the destination within a bounded time period and within bounded variations. Examples include applications such as audio, video, media streaming, real-time fax, and broadcasting. The total end-to-end time period within which a packet gets delivered from a source to a destination is called one-way "delay" and its variations is called delay variations or "jitter". Real-time applications are sensitive to delay and jitter. The degree of sensitivity depends on the type of the application. Real-time applications are further classified into two categories of "interactive" (intolerant) and "non-interactive" (tolerant) [4]. The differences between different categories of applications boil down to the differences in how applications react to packet delay and jitter.

Interactive (Intolerant)

Two-way conversational applications such as voice and video are called real-time interactive applications. In these applications audio/video signals are packetized at the source and transported through the network to the receiver. The temporal ordering as well as the end-to-end latency of these packets, as they cross the network, has fundamental impact on service quality as being experienced by the two end users. These applications are intolerant to the jitter and packet delay. End-to-end delays and jitter beyond certain limits could degrade the signal an unacceptable quality rendering the application unfunctional ("defective").

Non-Interactive (Tolerant)

Real time applications such as streaming audio/video (web-based broadcasting) are essentially one-way information transfer and, contrary to the interactive applications, users are not overly concerned if they see or hear events few milliseconds (or sometimes seconds) after their occurrence. They are tolerant to delay and jitter. The requirement for the total latency and jitter for tolerant applications is less stringent than tolerant applications. Given levels of latency and jitter that may lead to an unacceptable interactive audio may not be acceptable for a non-interactive audio streaming.

4.2. Non-real Time (Elastic) Applications

Non-real time, or elastic, applications wait for the packets to arrive before the application processes the data. In contrast to the real-time applications that drop delayed packets and do not re-transmit lost packets, elastic applications re-transmit lost or delayed packet until it gets to the destination. Examples of elastic applications include all other Internet services such as electronic mail (email), IP-based fax, File Transfer Protocol (FTP), and

web-browsing/e-commerce related applications. In short, elastic encompasses applications that are less sensitive to delay, jitter, and packet loss. They process packets immediately after their arrival.

Real -Time versus Non-Real Time

The fundamental distinction between the two applications lies in flow-control and error-handling mechanism. In real-time applications flow-control is self-paced, in that the rate at which packets are sent is determined by the characteristics of the sending application. The network then attempts to minimize the latency and packet loss based on condition imposed by the sender application. On the other hand, elastic applications use an adaptive flow-control algorithm in which packet flow rate is based on sender's view of the available network capacity and available space at the receiver's end - at that point of time.

With respect to the error handling, a real-time application has timeliness factor associated with it. Packets must be delivered within a certain elapsed time frame and within application specified sequencing. Otherwise, the packet becomes irrelevant and can be discarded. Packets in error, out of sequence packets, and packets with latency beyond the application requirement must be discarded. A non-real time application uses error - detection and retransmission recovery technique to recover the transmission error.

5. Reliability Assessment Framework

In this section we introduce a framework for understanding end-user's expectations of Internet services. The framework will assist service providers to identify, define, and measure parameters associated with each of the Internet-based services - at each of the stages of a user's experience. We demonstrate similarities between the Quality of Service Framework (QSF) introduced by the telecommunication community for the PSTN-based services with that of the Accessibility, Continuity, Fulfillment framework introduced in recent years by the computer/communication community for the Internet-based services. We also show the evolution of the service reliability from a "service provider centric" approach to a "user centric" approach.

Prior to the Bell System divestiture of 1984, quality and reliability of telecommunications services were determined largely from the service provider's perspective. The divestiture and competition changed the service quality and reliability paradigm from "service provider's view" to "user's/customer's view". The attention was focused on defining and measuring user's perspective of quality and reliability of communications services. In 1998, Richter's and Dvorak of AT&T Labs introduced a Quality of Service Framework (QSF) for defining the quality and reliability of communications services [6]. The QSF consisted of a matrix of "rows" and "columns". Rows were made up of communication functions performed by the users when using a service, and columns were made up of quality and reliability criteria perceived by users. Communication functions were

classified into three major categories of: *Connection Establishment, User Information Transfer*, and *Connection Release*. Quality and reliability criteria included *Speed, Accuracy, Availability, Reliability, Security,* and *Simplicity*. The ITU (International Telecommunications Unit) later accepted the framework. And specific QSF for each of the services including voice and facsimile were developed.

A similar framework can be developed to define, measure and monitor the end users' experience of quality and reliability of Internet-based services. Users would like to be able to *access* the desired application, initiate a transaction, *continue* using the application with no interruption for a desired duration of time, and *fulfill* the initiated transaction at a desired quality. Thus, for the Internet-based applications, user's perspective of quality and reliability is characterized by *Accessibility (A), Continuity (C),* and *Fulfillment (F)* framework defined below:

Accessibility:

Accessibility is the ability of a user to access the application and initiate a transaction. For instance, the first step in an email application consists of accessing mail server(s). For a dial-up user, accessibility consists of two succeeding stages - first, the user need to be able to access the dial platform and - second, the user need to access the appropriate mail server (given the accessibility to the dial platform has been successful).

Continuity:

Continuity is the ability of delivering the service with no interruption, given that accessibility has been successful and a transaction has been initiated. For the email application, for instance, if the message gets sent or received completely with a proper connection closure, it is considered a successful transaction. If not, it is considered a failed transaction.

Fulfillment:

Fulfillment is the ability of delivering the service meeting the user's quality expectations. In the case of email, for instance, the quality of transaction can be principally expressed in terms of throughput or speed. Once an access has been established for an email application, the time that takes to complete sending or receiving email determines the fulfillment. Email transactions taking longer than user's expectations will be considered as a "defective" transaction.

We demonstrate application of the framework for two major Internet services namely email and web browsing.

6. Fundamentals of Defects per Million (DPM)

In this section we introduce the fundamentals of the DPM concept and provide a general systematic process for DPM calculations. We list the set of tasks need to be undertaken towards in the process of assessment of service reliability as being experienced by users.

As indicated earlier, reliability is quantified in terms of probability. Accordingly, reliability must be specified and measured as a true proportion of defective "cases". The basic idea is to count number of defective items among a "specified" number of items. A popular count is "number of defectives per 100" - i.e., percentages. In many engineering fields, such as electronic equipment, a "defective" item is a rare event. For this reason the concept of Parts Per Million (PPM) - the number of defectives per one million - has been introduced and became increasingly popular way of measuring reliability. In late 1980s, the Committee on Component Quality and Reliability of the EIA (Electronic Industry Association) standardized the concept of Parts Per Million (PPM)[7] as a measure of component defectiveness to be used by both buyers and sellers of electronic components. The standard provided guidelines and methods to define "items" and calculate PPM for reliability measurement and monitoring.

In recent years, all the telecommunications industry have introduced the concept of Defect Per Million (DPM). Similar to the PPM, DPM is a measure of "defectiveness" normalized to one million. The DPM has been standardized [8] and can be applied to any service, product, component, system, or network. Early telecom applications include call blocking within the PSTN infrastructure and IP-packet loss, access network, and backbone DPM within the Internet infrastructure. However, application of DPM to the end-to-end user's perspective of the reliability of the emerging Internet-based services is not so obvious. These applications involve many challenges including definition of an "item" for a service, classification of "effective" / "defective" items, measurement approach ("active" or "passive"), DPM calculations and trending. Clearly users view of service reliability is based on "transaction" rather than network elements - users buy transactions. Therefore, our DPM method must reflect user's experience of a transaction rather than how the service is provided. To eliminate any ambiguity, we take a "transaction" as a unit of study, i.e., we observe users' transactions and classify them into "defective" / "effective" categories. Assuming that we have observed *N* transactions of which *D* have been defective, the general formula for DPM, then, is:

$$DPM = (D/N) * 10^6 \qquad (6.1)$$

6.1. Practical Steps for DPM Deployment

In practice, there are several essential tasks that need to be addressed to arrive at a program capable of measuring service reliability in terms of DPM. Bellow, we list the set of essential tasks common to all Internet-based services:

Task 1 - Specify the Internet Service

Specify the Internet service for which a DPM metric needs to be deployed. Specifically, these services could be any of the real-time services such as voice and video or non-real time applications such as email, web browsing, streaming, fax over IP and so forth.

Task 2 - Define End Users' Transaction

Define end user's transaction - this is a unit of study to be observed and classified into a "defective" or "effective" experience. For instance, for email service, a transaction would consist of "Get-Mail" experience - that is getting a mail from a mail server. For a fax over IP service, a transaction could consist of sending five pages. For web browsing service a transaction could consist of the experience of downloading the main page of a web site.

Task 3 - Define a Defective Transaction

Define a "defective" transaction. The DPM concept is based on the premise that from a user's perspective, a given application (such as email) can be classified as either "defective" or "effective" (failed or succeeded) transaction. Transparent to the user, a typical web application will go through several operationally distinct stages. For instance a dial up email transaction would consist of several distinct stages such as "accessing the dial platform", "accessing the mail servers", "sending or getting email with no interruption", and "delivering the service within an acceptable time period". Each of these stages itself is made up of a number of sub-stages. For instance, "access dial platform" consists of "telephone line dial up", "modem sync", and "authentication". Failure in any of the above stages would lead to user's experience of a "defective" email transaction. It is essential to provide a clear definition of a defective transaction.

Task 4 - Decide on Measurement Approach

Decide on measurement approach. There are two distinct measurement approaches called "active" and "passive".

- **Active measurement:** In active measurement, a transaction request is initiated with the sole purpose of measuring the performance. An active approach emulates users' experience and provides an end-to-end perspective of service reliability. It generates the transaction needed to make the measurement and allows much more direct observation and analysis. An active approach necessitates a carefully designed sampling plan including a representative transaction; a geographical balanced sampling distribution, and appropriate sampling frequency.

- **Passive measurement:** In passive measurement, performance data are collected from network elements including routers, server log files, and switches. An independent device passively monitors and collects performance data as the traffic traverses within the network. Since the measurement is made within a network it would be impossible to extract a complete end-to-end user's experience.

The decision on measurement approach could include both or either active or passive approaches.

Task 5 - Sampling Plan and Data Collection

Develop a sampling plan and collect the data. This task involves many activities including the desired confidence level and accuracy and validity of the collected data [9] (see [9]).

Task 6 - DPM Calculations and Trend Analysis

In practice many DPMs, each representing different segment of an end-to-end reliability might be calculated. Trend analysis of the DPMs over time requires advanced statistical techniques that need to be specified and formulated at this stage. Six Sigma [10] and Statistical Process Control [11] are appropriate techniques for DPM trend analysis (see {10] and [11]).

Task 7 - Set DPM Thresholds

Set the threshold against which the DPM process needs to be controlled. Thresholds may be set for each of the DPMs, and different thresholds could be set for a given DPM. A threshold could set to meet the operations requirements, another threshold could be set based on the historical DPM data, yet another threshold could be set to meet Service Level Agreement (SLA). The DPM trend will be plotted against the specified thresholds. When a threshold is exceeded a corrective action takes place and a root cause analysis is conducted.

Task 8 - Execute the P3 Paradigm

Execute the P3 paradigm of Predictive, Proactive, and Preventive. A continuos monitoring and analysis of the DPM enables service providers, network operators, and system/equipment/product vendors the ability to detect, diagnose, and resolve failures impacting users' transactions. The DPM program is the essence of the P3 paradigm and meeting end-users' expectations of reliability of the emerging Internet-based services.

The remainder of this paper demonstrates application of the DPM methodology to the two major Internet services of "e-mail" and "web/e-e-commerce applications". In section 7 we apply the methodology to generic e-mail, and in particular to the AT&T WorldNet e-mail services. In section 8 we apply the methodology to web-hosting as well as the Intelligent Content Distribution service (ICD).

7. E-mail Application

E-mail is a major Internet application. It is the most popular tool of communication for businesses and consumers. Users have come to expect a high degree of availability and reliability from the Internet mail services. Most of the intra-company email traffic is on

the LAN (Local Area Network) or WAN (Wide Area Network), and most of the inter-company email traffic as well as most of the consumer / residential email traffic flows over the Internet. The growth and use of Internet mail has prompted an interest in the mail services provided by all major ISPs.

In this section we provide a generic architecture of Internet mail and how it works, describe an implementation example - the AT&T WorldNet mail, and formulate measurement of reliability of email services.

7.1. E-mail Architecture

We begin describing email pathway, i.e., how a mail message would travel from one user to another user. To make the case more general, let us assume that the two users are being served by different ISPs. The end-to-end communication between the users can be divided into three parts – 1) mail traveling from the sender's computer to the ISP computer, 2) mail traveling from one ISP computer to the other ISP computer, and 3) the receiver accessing the mail using their computer from their ISP computer. Figure 3 shows a high-level architecture of email [13]

Each of the computers mentioned above run software applications allowing them to perform the tasks of establishing a connection, authentication, and mail transfer as needed. The user computers run the software that allows them access to their respective ISP networks. The access is controlled by the ISPs for billing and security reasons. The user computer thus has to first establish a connection to the ISP access network, from where the user can then access the mail service. The ISP access network connection is made possible either over a normal phone line and modem dial in, or it can be over a broadband connection – either via cable or DSL connection. In the corporate environment, the network access is typically over a LAN.

A mail client is a software application that runs on the end-user machine. The client allows the user to create and send messages, as well as receive messages. A mail client is also known as Mail User Agent (MUA).

Each ISP has a set of computers running mail application within its network. An Internet mail server is a software application that provides mail services of transmitting and receiving mail across the Internet. It is also known as Mail Transfer Agent (MTA). The mail server works with client software to send and receive messages from the users. The machines transfer mail messages using a common protocol. For the first two parts – mail traveling from the sender's computer to the ISP computer and from one ISP computer to another, SMTP (Simple Mail Transfer Protocol) is used. For the third part, the end-user computer accessing mail from the mail server, few different protocols exist. The POP3 (Post Office Protocol Version 3) protocol has been used most commonly, while IMAP (Internet Mail Access Protocol) is now growing in usage.

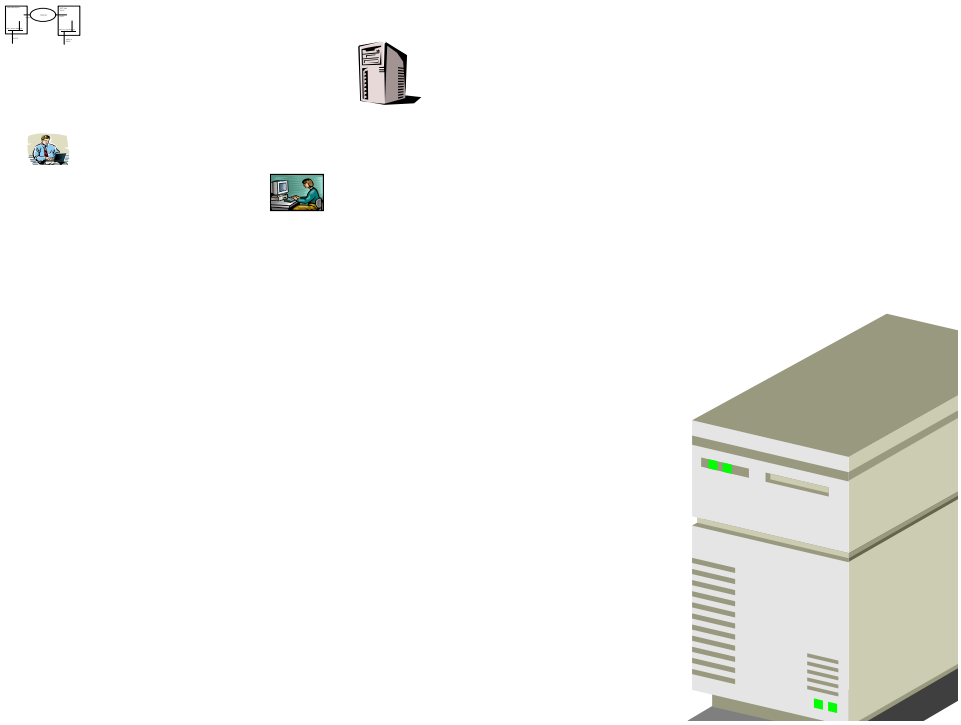

Figure 3: A High-Level E-mail Architecture

With POP (Post Office Protocol), a personal computer user periodically connects to the "server" machine and downloads all of the pending mail to the "client" machine. Thereafter, all mail processing is local to the client machine. POP provides a store-and-forward service, intended to move mail (on demand) from a server to a user machine, usually a PC. Once delivered to the PC, the messages are typically deleted from the POP server. This operation may be referred to as Offline mode because once the mail is downloaded; the user can process the mail without being connected to the network. In this case, the mail is typically available from just one designated computer.

The IMAP protocol is designed to permit manipulation of remote mailboxes as if they were local. With IMAP, the mail client machine does not normally copy it all at once and then delete it from the server. It's more of a client-server model where the IMAP client can ask the server for headers or the bodies of specified messages, or to search for messages meeting certain criteria. With IMAP, all email, including the inbox and all filed mail, remains on the server at all times. The client computer must be connected to the server for the duration of your email session. A temporary copy of the mail or part of it is downloaded to the client computer while you read it or send it, but once the connection is finished, the copy is erased from the client computer; the original remains on the server. IMAP thus works in an online mode, and the mail is accessible from multiple computers.

.

7.2. WorldNet E-mail Service

The generic description of mail service provided above translates into specific implementations by the ISPs. An overview of AT&T's WorldNet Mail service is covered here as an example. The WorldNet service is a dial up service that uses modems and phone line. A broadband service application is very similar after the ISP network access. The main difference between the two is the ISP network access. In the case of broadband, the high-speed ISP network connectivity through a cable or a DSL is available on the user premises.

The WorldNet Mail system performs the following functions:

- Receives email from the Internet or from WorldNet users via CBB
- Temporarily stores mail until retrieved by user
- Downloads mail to user's PC on demand
- Routes outgoing email from WorldNet users to the Internet or other WorldNet users
- Allows multiple mailboxes for a single user account
- Stores and updates user account information

7.3. Dial Platform Architecture

A typical user accesses the Dial Network Platform from the computer via a dial up modem by dialing one of the modem pool numbers to first reach AT&T's Dial Platform. Once the modem pool is reached and the modem communication is established, IP/PPP (Point to Point Protocol) connection gets set up to allow the user computer to communicate with the NAAS (Network Authentication and Authorization Services) server. The user computer then provides the login name and password to gain access to the Dial Platform. The access to the mail services from the Dial Platform is depicted in Figure 1 [14]. Once the user has successfully logged on to the Dial Platform, access to mail services can proceed as described below.

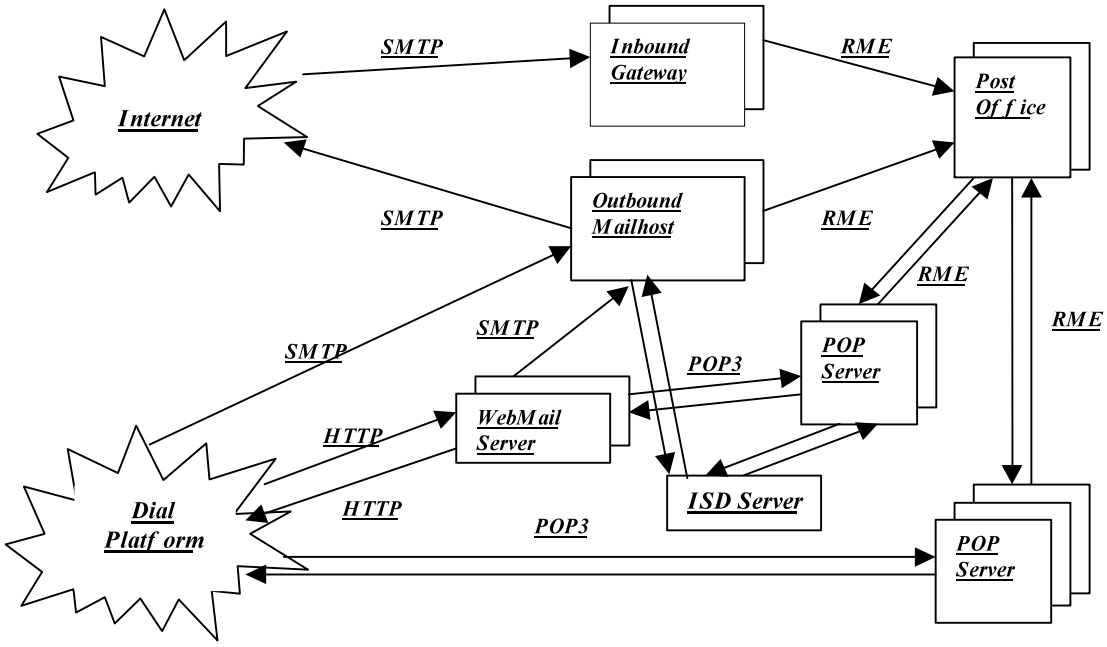

Figure 2: WorldNet Mail System Architecture

The WorldNet Mail system comprises of various server machines running their associated software processes. The category names and basic functions of the various machines with the respective server applications are:

1. **Mail Gateway**: Route incoming mail from the Internet to the Post-Offices, bounce errored mail, and prevent spam.
2. **Mail-host**: Send outgoing mail to the Internet and sends mail for WorldNet customers to the Post-Offices.
3. **Post-Office**: Hold the incoming mail for WorldNet users in their respective mailboxes.
4. **POP-Server**: Authenticate mail access and deliver user mail from the Post-Office mailbox to the user.
5. **Web-Mail Server**: Provide mail access via browser.
6. **ISD (Integrated Services Directory):** ISD provides detailed user account information such as mailbox names and passwords and is used to update all the directory cache databases.

7.4. E-mail Transactions

There are two primary transactions that a user performs with reference to an email application:

- **Send-Mail**: Sending one or more mail messages to other users on the Internet or on WorldNet Service.

- **Get-Mail**: Accessing and/or getting mail downloading messages from the user's mailbox on the Post-Office.

Mail services are accessed via mail client software such as the Outlook Express, Eudora, or a browser residing on the user's computer. The mail client software runs the necessary scripts to communicate with the mail server. The details of this interaction are transparent to the user.

Send-Mail

Using the SMTP protocol, the mail client sends the mail message(s) to the Mail-host. The Mail-host forwards the mail messages to the appropriate Post-Office servers if they are addressed to WorldNet users. If not, they are forwarded to the Internet. The client sends a sequence of commands to server in the process of sending the mail. If any of these commands fail, the server log file creates an error entry.

Get-Mail

Mail from the Internet arrives at the Mail Gateway machines. After screening the mail for errors and spam, it gets forwarded to the appropriate Post-Office servers where the user mailboxes are located.

Let us consider accessing the mail using the POP3 protocol, which is the primary mode of mail access.

The mail client establishes connection with a POP-server, from where the mail is accessed. The user id and password are provided to the POP-server. Based on information from the ISD server database, the user is authenticated and the POP-server knows the user mailbox (Message Store ID). The mail client then issues a sequence of commands to download the mail. The POP server then forwards the mail residing on the Post-Office mailbox to the client. Any failures occurring during the process of retrieving mail are written to a log file residing on the server.

7.5. E-mail DPM Measures

Because the Internet mail service has become important to the users, it has become important to the ISPs to provide a better service; and providing a better service starts with monitoring the existing service.
AT&T continuously measures and monitors the quality and reliability of its services from user's perspective. As it was indicated earlier, user's perspective of quality and reliability of Internet-based applications is characterized within the Accessibility, Continuity, and Fulfillment framework.

1. **Accessibility:** Accessibility is the ability of a user to access the application and initiate a transaction (Getmail or Sendmail). For both Getmail (receiving mail) and Sendmail (sending mail), the access consists of two parts – first an access to the WorldNet Dial Platform, and then the access to the appropriate mail server. Getmail or Sendmail access is classified as "defective" if in attempting a transaction, either one of the above fails.

2. **Continuity:** Continuity is the ability of delivering the service (mail) with no interruption, given that accessibility has been successful and a transaction (mail sending or getting) has been initiated. If the message gets sent or received completely with a proper connection closure, it is considered a successful transaction; otherwise it is classified as "defective".

3. **Fulfillment:** Fulfillment is the ability of delivering the service meeting the customers' quality expectations. In the case of email, the quality of transaction can be principally expressed in terms of throughput or speed. Once an access is established for Getmail or Sendmail, the times required to complete the transaction (send or receive email) determines fulfillment. Thresholds are set for Getmail and Sendmail. Email transactions taking more time than the specified thresholds are classified as "defective".

7.6. Accessibility DPM

For the Accessibility, “active" monitoring is employed, i.e. email tests are conducted to gather data. The test platform initiates test transactions. The test transaction to access mail is monitored at each step for its success or failure, and the data is used to calculate the DPM value. A sampling scheme is incorporated with balanced sampling to assure accuracy of the results. The following measures are calculated and reported as their daily DPM value and their average daily DPM value for each month.

1. **Line DPM**: Failure to connect to the modem because of line problems (busy, no carrier, ring-no-answer.) Let A = the number of attempts to access the Dial Platform (i.e. the modems connect, synchronize, and IP/PPP setup and authentication are all successful) and F1 = the number of modem connect failures. Then,

$$LineDPM = \frac{F1}{A} \times 10^6 \qquad \text{----------(7.1)}$$

2. **Modem Synchronization DPM**: Failure of the modems (user’s and Dial Platform) to communicate properly even though connected. Let A = the number of attempts to access the Dial Platform, and F2 = the number of modem synchronization failures. Then,

$$ModemSyncDPM = \frac{F2}{(A - F1)} \times 10^6 \qquad \text{----------(7.2)}$$

3. **IP/PPP Setup and Authentication DPM**: IP or PPP negotiation failed, or authentication to the email service failed (even though the modem connected and synchronized). Let A = the number of attempts to access the Dial Platform, and F3 = the number of IP/PPP setup or authentication failures. Then,

$$IP/PPP\_DPM = \frac{F3}{(A - F1 - F2)} \times 10^6 \qquad \text{----------(7.3)}$$

4. **POPserver Connect DPM**: Failure to connect to the POPserver after Dial Platform connection. Let B = the number of attempts to access the mail (i.e. connect to the POPserver and authenticate) and F4 = the number of failures to connect. Then,

$$POPserverConnectDPM = \frac{F4}{B} \times 10^6 \qquad \text{---------- (7.4)}$$

5. **POPserver Authentication DPM**: Failure to authenticate to the POPserver even though connected. Let B = the number of attempts to access the mail and F5 = the number of failures to authenticate to the POPserver. Then,

$$POPserverAuthenticationDPM = \frac{F5}{(B - F4)} \times 10^6 \qquad \text{----------(7.5)}$$

6. **Mailhost Connect DPM**: Failure to connect to the Mailhost after Dial Platform connection. Let B = the number of attempts to send the mail (i.e. connect to the Mailhost) and F6 = the number of failures to connect. Then,

$$MailhostConnectDPM = \frac{F6}{B} \times 10^6 \qquad \text{----------(7.6)}$$

7.7. Continuity DPM

For the Continuity, “passive" monitoring approach is utilized, i.e., the performance data is collected from network elements (mail server log files). The mail servers keep log files that provide information on each mail transaction as whether or not the transaction was successful or it failed and the reason for failure. The information from all mail servers is gathered on the Log server, from where the DPM server calculates the total number of attempts and failures on each server along with the reasons for failure. The following measures are calculated and reported as their daily DPM value and their average daily DPM value for each month.

1. **Sendmail Continuity DPM**: Failure to complete the transaction of sending mail from the client to the mail server, once the transaction was initiated. Let C1 = the number of initiations to send mail on all Mailhost servers and F7 = the number of failures while sending the mail. Then,

$$SendmailContinuityDPM = \frac{F7}{C1} \times 10^6 \qquad \text{---------- (7.7)}$$

2. **Getmail Continuity DPM**: Failure to complete the transaction of receiving mail from the mail server to the client, once the transaction was initiated. Let C2 = the number of initiations to receive mail on all POPservers and F8 = the number of failures while receiving the mail. Then,

$$GetmailContinuityDPM = \frac{F8}{C2} \times 10^6 \quad \text{---------- (7.8)}$$

7.8. Fulfillment DPM

The Fulfillment DPM uses data from “active" monitoring. Email test transactions are sent and received - just the way a typical end user practice. The time taken to send and receive the test email is recorded. Performance thresholds *X* and *Y* are specified for and Getmail transactions respectively. Test transactions taking more time than the specified thresholds are classified as "defective". Following DPMs are accordingly calculated and reported for each specified time horizon:

1. **Sendmail Fulfillment DPM**: Failure to complete the transaction of sending mail from the client to the mail server within the specified time. Let G1 = the number of “Sendmail” test transactions completed and F9 = the number of transactions that took more than 10.879 seconds to complete. Then,

$$SendmailFulfillmentDPM = \left(\frac{F9}{G1} \times 10^6 \right) \quad \text{---------- (7.9)}$$

2. **Getmail Fulfillment DPM**: Failure to complete the transaction of receiving mail from the mail server to the client within the specified time. Let G2 = the number of “Getmail” test transactions completed and F10 = the number of transactions that took more than 8.098 seconds to complete. Then,

$$GetmailFulfillmentDPM = \left(\frac{F10}{G2} \times 10^6 \right) \quad \text{----------(7.10)}$$

.

ACKNOWLEDGEMENTS:

We would like to Thank many people who have helped us review the reliability white paper and we sincerely appreciate all the people who helped us in this endeavor and our sincere gratitude and appreciation to everyone involved.

[1] Vint Cerf, "The Internet is for Everyone,", IETF, RFC2026, January 1, 2002.

[1] AGREE (1957), Reliability of Military Electronics Equipment, AGREE Task Group Report, U.S., Government Printing Office, Washington, D.C.

[2] ANSI/ASQC Standard A3-1978, "Quality Systems Terminology," American society for Quality Control, Milwaukee, Wisconsin

[3] Committee T1 - Telecommunications, "IP Access Network Availability Defect Per Million (DPM)", Technical Report (Draft January 2002)

[5] K. Hafner, "The Internet's Invisible Hand," The New York Times, January 10, 2002

[12] Y. Kogan, G. Maguluri, G. Ramachandran, "Defect per Million and IP Backbone Availability", Committee T1 Contribution, September 2001.

[4] Barden, R., D. Clark, and S. Shenker. "Integrated Services in the Internet Architecture: An Overview," IETF RFC 1633, June 1994.

[6] J.S. Richters, C.A. Dvorak, "A Framework for Defining the Quality of Communications Services," IEEE Communication Magazine, October 1988.

[7] Robert E. Odeh, D. B. Owen, "Parts Per Million Values for Estimating Quality Levels,", Marcel Dekker, 1988.

[8] Committee T1- Telecommunications, "IP Access Network Defect Per Million,", Draft Technical Report, January 2002.

[9] Gary G. Brush, "How to Choose Proper Sample Size,", American Society for Quality Control, Volume 12.

[10] M. J. Harry, J. R. Lawson, " Six Sigma Producibility Analysis and Process Characterization,", Six Sigma Research Institute and Motorola University Press, Addison -Wesley, 1992.

[11] E. Grant, R. Leavenworth, " Statistical Quality Control,", McGraw-Hill, 1998.

[14] H. Burns, R. Oppenheim, V. Ramaswami, P. Wirth, "End-to-End IP Service Quality:Are DPM Service Transactions the Right Measure?,", Journal of the IBTE, Vol 2, Part 2, April 2001.