



ATIS-0100016

END-TO-END SERVICE AVAILABILITY: GENERAL DEFINITION

TECHNICAL REPORT



The Alliance for Telecommunication Industry Solutions (ATIS) is a technical planning and standards development organization that is committed to rapidly developing and promoting technical and operations standards for the communications and related information technologies industry worldwide using a pragmatic, flexible and open approach. Over 1,100 participants from over 300 communications companies are active in ATIS' 22 industry committees and its Incubator Solutions Program.

< <http://www.atis.org/> >

Notice of Disclaimer & Limitation of Liability

The information provided in this document is directed solely to professionals who have the appropriate degree of experience to understand and interpret its contents in accordance with generally accepted engineering or other professional standards and applicable regulations. No recommendation as to products or vendors is made or should be implied.

NO REPRESENTATION OR WARRANTY IS MADE THAT THE INFORMATION IS TECHNICALLY ACCURATE OR SUFFICIENT OR CONFORMS TO ANY STATUTE, GOVERNMENTAL RULE OR REGULATION. AND FURTHER, NO REPRESENTATION OR WARRANTY IS MADE OF MERCHANTABILITY OR FITNESS FOR ANY PARTICULAR PURPOSE OR AGAINST INFRINGEMENT OF INTELLECTUAL PROPERTY RIGHTS. ATIS SHALL NOT BE LIABLE, BEYOND THE AMOUNT OF ANY SUM RECEIVED IN PAYMENT BY ATIS FOR THIS DOCUMENT, WITH RESPECT TO ANY CLAIM, AND IN NO EVENT SHALL ATIS BE LIABLE FOR LOST PROFITS OR OTHER INCIDENTAL OR CONSEQUENTIAL DAMAGES. ATIS EXPRESSLY ADVISES ANY AND ALL USE OF OR RELIANCE UPON THIS INFORMATION PROVIDED IN THIS DOCUMENT IS AT THE RISK OF THE USER.

NOTE - The user's attention is called to the possibility that compliance with this standard may require use of an invention covered by patent rights. By publication of this standard, no position is taken with respect to whether use of an invention covered by patent rights will be required, and if any such use is required no position is taken regarding the validity of this claim or any patent rights in connection therewith.
--

ATIS-0100016, *End-to-End Service Availability: General Definition*

Is an ATIS Standard developed by the **Network Reliability Task Force (REL)** under the **ATIS Network Performance, Reliability, and Quality of Service Committee (PRQC)**.

Published by

Alliance for Telecommunications Industry Solutions
1200 G Street, NW, Suite 500
Washington, DC 20005

Copyright © 2008 by Alliance for Telecommunications Industry Solutions
All rights reserved.

No part of this publication may be reproduced in any form, in an electronic retrieval system or otherwise, without the prior written permission of the publisher. For information contact ATIS at 202.628.6380. ATIS is online at < <http://www.atis.org/> >.

Printed in the United States of America.

END-TO-END SERVICE AVAILABILITY: GENERAL DEFINITION

Secretariat

Alliance for Telecommunications Industry Solutions

Approved September 2007

Abstract

This Technical Report (TR) provides a general definition of end-to-end service availability for application to a service supported by a telecommunications network. This TR is intended as a framework and guide in the development of service-specific definitions of this concept. Service-specific definitions of this concept will be provided in separate forthcoming TRs.

FOREWORD

The Alliance for Telecommunication Industry Solutions (ATIS) serves the public through improved understanding between carriers, customers, and manufacturers. The Network Performance, Reliability, and Quality of Service Committee (PRQC) -- formerly T1A1 -- develops and recommends standards, requirements, and technical reports related to the performance, reliability, and associated security aspects of communications networks, as well as the processing of voice, audio, data, image, and video signals, and their multimedia integration. PRQC also develops and recommends positions on, and foster consistency with, standards and related subjects under consideration in other North American and international standards bodies.

This Technical Report (TR) provides a general definition of end-to-end service availability for application to a service supported by a telecommunications network. This TR is intended as a framework and guide in the development of service-specific definitions of this concept. Service-specific definitions of this concept will be provided in separate forthcoming TRs.

This technical report is intended for providers of IP-based communications networks and services, communications equipment suppliers, and government agencies responsible for addressing emergency situations.

Suggestions for improvement of this document are welcome. They should be sent to the Alliance for Telecommunications Industry Solutions, PRQC Secretariat, 1200 G Street NW, Suite 500, Washington, DC 20005.

At the time it approved this document, PRQC, which is responsible for the development of this TR, had the following members:

M. Neibert, PRQC Chair
 N. Seitz, PRQC Vice-Chair
 C. Underkoffler, ATIS Chief Editor
 J. Bennet, PRQC Technical Editor

Organization Represented	Name of Representative
Alcatel USA Inc.	Ken Biholar Tim Pantis (Alt)
AT&T	Percy Tarapore Charles A. Dvorak (Alt)
Department of Defense	Chris Fitzgerald
Embarq Corporation	Carl M. Coopage John M. Heinz (Alt)
Ericsson Incorporated	Mustafa Kocaturk Susan Sabater-Maroto(Alt)
ETRI	Tae-Soo Chung Sung-Soo Kang (Alt)
Intelsat	Mark Neibert
National Communications Systems	An Nguyen Carol-Lyn Taylor (Alt)
National Telecom & Info Admin	Neal B. Seitz

Organization Represented	Name of Representative
- NTIA	Arthur Webster (Alt)
Nokia Siemens Networks	Nagaraja Rao David Francisco (Alt)
Nortel	Joseph A. Zearth
Qwest	Steve Showell Michael Fargano (Alt)
Sprint	Steve Oliva
Telcordia Technologies	Spilios Makris Cliff Halevi (Alt)
Tellabs Operations, Inc.	William A. Walker Kevin Stodola (Alt)
Verizon Communications	John Colombo Greg Cermak (Alt)

The Network Reliability Task Force (REL) was responsible for the development of this document, and had the following officers and participants:

A. McCain, REL Chair
 S. Makris, REL Vice Chair
 F. Kaudel, PRQC Chief Editor
 J. Bennett, REL Technical Editor
 P. Tarapore, REL Technical Editor

C. Bailey
 J. Bennett - Telcordia
 C. Dvorak - AT&T
 P. Kimbrough - AT&T

S. Makris - Telcordia
 A. McCain - AT&T
 A. Nguyen - NCS
 R. Paterson - Nortel

N. Seitz - NTIA
 P. Tarapore - AT&T
 A. Webster - NTIA/ITS
 R. Wohlert - AT&T

TABLE OF CONTENTS

1 PURPOSE	1
2 SCOPE.....	1
3 APPLICATION	1
4 RELATED WORK.....	2
5 GENERAL DEFINITION OF END-TO-END SERVICE AVAILABILITY	2
5.1 THE COMMUNICATIONS NETWORK.....	3
5.2 WEIGHTS FOR ACCESS POINT PAIRS	4
5.3 SERVICE AVAILABILITY BETWEEN TWO ACCESS POINTS	4
5.4 EXAMPLE CALCULATION OF END-TO-END SERVICE AVAILABILITY.....	5
6 EQUIVALENT DEFINITION OF END-TO-END SERVICE AVAILABILITY	7
6.1 DEFINITION IN TERMS OF WEIGHTED DOWNTIME	7
6.2 EXAMPLE OF PREDICTION OF END-TO-END SERVICE AVAILABILITY.....	7
7 ESTIMATION OF END-TO-END SERVICE AVAILABILITY	8
8 CONCLUSION	9
9 ABBREVIATIONS AND ACRONYMS	9

ATIS Standard on –

End-to-End Service Availability: General Definition

1 PURPOSE

End-to-end service availability has become the major metric for the reliability of a communications network. Service providers are using this metric as a basis for network design and for Service Level Agreements (SLAs). However, there is no industry-standard definition for a metric that describes the availability of a service over a given network taken as a whole.

This Technical Report (TR) provides a general definition of end-to-end service availability for application to a service supported by a telecommunications network. This TR is intended as a framework and guide in the development of service-specific definitions of this concept. Service-specific definitions of this concept will be provided in separate forthcoming TRs.

2 SCOPE

This definition of network availability encompasses several concepts that have not been addressed in the past. It weights the various point-to-point access pairs across the network in estimating an overall network availability. It allows for network growth in the time period in which network availability is measured or estimated. It allows for the estimation not only of the mean network availability, but also the variability of network availability measures and estimates.

3 APPLICATION

This work has application to SLAs and other statements made concerning network availability. Currently, the lack of a standard definition for network availability makes any statements of network availability open to wide and diverse interpretation, potentially leading to serious misunderstandings.

This TR provides a high-level definition of network availability. It assumes perfect knowledge of service outages and their effects on a network's customer base. Obtaining such information to such a perfect degree would undoubtedly require resources far beyond the levels achievable at this time by service providers. As such, the definition should be viewed as a goal to reach toward. It should be used as a guide in directing the efforts of service providers to measure and estimate the availability of service across their networks. It is unlikely that service providers would be able to reach this goal, but they are encouraged to try to attain it as closely as their resources permit, and understand the limitations and uncertainties that remain in their measurement of network availability.

The definition is applicable to the design of networks, as well as to the monitoring of existing networks. The definition allows a network designer to estimate overall network availability for various network

design options. This capability allows the designer to consider not only the expected availability, but also the variability in availability observed. For example, it is possible that one Network Architecture A has a higher expected availability than Network Architecture B, but the variability of the availability measure in B is lower than A. (This can happen when an architecture has greater dependency on one part of the network that has generally high reliability, but when it fails, impacts the availability greatly.) The definition described here allows for considering this aspect of network availability. For example, it is possible to design the network to achieve the best possible result for the 10% quantile rather than the mean expected availability (i.e., guaranteeing that network availability is no lower than x 10% of the time.)

4 RELATED WORK

Work related to this topic has been and continues to be performed by various standards bodies including ITU-T and ATIS. This section cites examples of such work.

T1.TR.78-2003, *Access Availability of Routers in IP-Based Networks*.¹

T1.513-2003 (R2008), *Frame Relay Data Communication Service - Access, User Information Transfer, Disengagement, and Availability Performance Parameters*.¹

ITU-T Recommendation Y.1540, *IP Packet Transfer & Availability Performance Parameters*, December 2002.²

ITU-T Recommendation Y.1540 Amendment 1, *Background on IP Service Availability*, August 2003.

ATIS Contribution PRQC-2008-004R2, *Draft Standard - Availability Metric for IP-Based Networks and Services* (G. Choudhury, Y. Kogan, A. Morton, E. Rojek, & P. Tarapore; February 2008, work in progress).

5 GENERAL DEFINITION OF END-TO-END SERVICE AVAILABILITY

A *communications network* is an entity that connects users who access the network through ingress and egress points. The network is composed of various network elements and links that connect them. However, to define end-to-end service availability at the highest conceptual level, the network can be treated as a single entity with ingress and egress points that allow access to the network by its users. *Service availability* is the fraction of time that a service is considered available as supported by a system. In the case of end-to-end service availability for a network, the system is a communications network and the definition must be extended to represent all access points to the network.

End-to-end service availability for a communications network is the fraction of time that service is available between an arbitrarily-specified ingress point and arbitrarily-specified egress point to the communications network. Typically, an access point is expected to be an edge network element;

¹ This document is available from the Alliance for Telecommunications Industry Solutions (ATIS), 1200 G Street N.W., Suite 500, Washington, DC 20005. < <https://www.atis.org/docstore/default.aspx> >

² This document is available from the International Telecommunications Union. < <http://www.itu.int/ITU-T/> >

however, the level of granularity can be adjusted as needed. For example, an access point could be an access port if the calculation or data demands that level of granularity.

Given a network with a set of access (ingress and egress) points, a specified service, and a specified time window, the definition of end-to-end service availability for the network supporting that service over the duration of the time window is:

$$\text{End-to-End Service Availability} = \frac{\sum_{i=1}^N \sum_{j=1}^N w_{ij} f_{ij}}{\sum_{i=1}^N \sum_{j=1}^N w_{ij}}$$

where f_{ij} is the fraction of time that service is available between ingress i and egress j (see 4.3) and N is the number of access points to the network.

Note that the distinction between f_{ij} and f_{ji} allows for the possible difference in originating and terminating service. The w_{ij} weights are needed when the access points are different sizes (e.g., different numbers of users, different communication bandwidths) and/or have been in-service for different periods of time. The inclusion of separate w_{ij} and w_{ji} weights allows for differences in the direction of service; for example, if service goes from i to j but not j to i , then $w_{ij} > w_{ji} = 0$.

The following sections describe the key elements of the definition:

1. Specifying the communications network and its access points.
2. Specifying the w_{ij} weights for access point pairs.
3. Specifying service availability f_{ij} between access points.

5.1 The Communications Network

In order to define end-to-end service availability, the network must be clearly defined. The limits of the network (i.e., demarcation points) as defined by the access points must be clearly understood. It is recognized that services over IP networks (e.g., VoIP services) may span several provider network domains. In such cases, the assignment of availability for individual providers becomes an issue that must be dealt with. It is proposed that initial work should focus on assessing availability over a single provider's network, and that later efforts can concentrate on tying together end-to-end availability over several network domains.

The definition of a service outage needs to be agreed upon. Specifically, when is a service considered to be unavailable? This depends on the type of service.

In general, it is simplest and most practical to define outage as a binary phenomenon -- i.e., the service is either "up" or "down" based on a specific defined threshold. However, it may be desirable, in a particular evaluation, to consider the possibility of a "partial outage" - e.g., an outage that permits limited service delivery due to inadequate bandwidth or network resources. In such cases, agreement needs to be reached on how to define the f_{ij} fractions in the availability definition above. These fractions could be "pro-rated" according to the degree of partial availability, if reasonable estimates of such pro-rations are available. Alternately, the fractions could be assigned binary values:

- ◆ 1, if X% of traffic is estimated to be conveyed between the end points where the value X = 50, for example.
- ◆ 0 otherwise.

The definition of service outage and partial outage for each service is beyond the scope of this document. However, any calculation of end-to-end service availability should include a description of the definition of service outage used, as well as how partial outage was estimated if the concept was used.

5.2 Weights for Access Point Pairs

Access point weights w_{ij} should be defined based on the sizes of the access points and the time that both access points were in service. One possible method for defining the weight for service between access points i and j is:

$$w_{ij} = s_i s_j D_{ij}$$

where s_i (s_j) is the size of access point i (j) and D_{ij} is the duration of time the network has supported access points i and j in the time period $[T_s, T_E]$. For example, consider a wireline circuit network for voice service where the time window of interest is the year 2004 (i.e., $[T_s, T_E] = [\text{January 1, 2004, December 31, 2004}]$). Two access points could be local switches A and B with 30,000 access lines for A and 20,000 access lines for B. Switch A was in-service for all of 2004 while Switch B was in service for the last 100 days of 2004. In calculating the end-to-end service availability for voice service in this network in 2004, the weight for service availability between switches A and B is:

$$w_{AB} = 30,000 \times 20,000 \times 100 = 60,000,000,000.$$

Access point size may depend on the type of access element, type of network, and/or type of service. This weight definition requires that access points use the same unit dimensions in defining size. This may require conversions for access points with different unit dimensions.

Other definitions for weighting can be considered in the discussions. For instance, in IP networks, port bandwidths can be added together to represent the weight between the two port end points.

5.3 Service Availability Between Two Access Points

The service availability f_{ij} is the fraction of time that service was available between the ingress point i and the egress point j ; the ingress point i is the point at which service was initiated and the egress point j is the point to which ingress point i desires communication. This fraction encompasses any impacts on service, including impacts from loss of point-to-point connectivity or any failure in the network that affects the user's ability to communicate. Given the time period of interest was between times T_s and T_E and given M periods of service outage from ingress point i to egress point j were identified with durations $O_1, O_2, O_3, \dots, O_M$, then the service availability f_{ij} is :

$$f_{ij} = 1 \text{ for } M = 0 \text{ or}$$

$$f_{ij} = 1 - \frac{\sum_{m=1}^M O_m}{D_{ij}} \text{ for } M > 0$$

The definition of service outage period between two access points varies from service to service. These definitions will be addressed separately in each service-specific document. In general, a service outage includes inability to initiate a service attempt, inability to terminate a service attempt, unsatisfactory quality for service in progress, and interruption of service in progress. By varying the definition of service outage, end-to-end service availability at different grades of service can be estimated.

The definitions should consider failure modes to be addressed; the metric could encompass any service outage, including those resulting from hardware failures, software failures, physical damage (e.g., fiber cuts), procedural errors, database errors, maintenance, and traffic congestion.

5.4 Example Calculation of End-to-End Service Availability

Consider a small network of N=5 access points with the following sizes in terms of thousands of lines:

Access Node i	Size s _i
1	10
2	20
3	30
4	40
5	50

We wish to calculate the end-to-end service availability for this network over the past year. All access nodes except Access Node 1 were in-service for the full year; Access Node 1 was in-service only for the last six months. The in-service time D_{ij} for all access point combinations is shown in the following table in months:

D _{ij} (months)	Access Node 1	Access Node 2	Access Node 3	Access Node 4	Access Node 5
Access Node 1	6	6	6	6	6
Access Node 2	6	12	12	12	12
Access Node 3	6	12	12	12	12
Access Node 4	6	12	12	12	12
Access Node 5	6	12	12	12	12

ATIS-0100016

The weights $w_{ij} = s_i s_j D_{ij}$ for the calculation are:

w_{ij}	Access Node 1	Access Node 2	Access Node 3	Access Node 4	Access Node 5
Access Node 1	600	1200	1800	2400	3000
Access Node 2	1200	4800	7200	9600	12000
Access Node 3	1800	7200	10800	14400	18000
Access Node 4	2400	9600	14400	19200	24000
Access Node 5	3000	12000	18000	24000	30000

The sum of the weights in the table is $\sum_{i=1}^N \sum_{j=1}^N w_{ij} = 252600$.

The availability f_{ij} of service between access point pairs, using empirical data as measured over the year under consideration, is shown in the following table:

f_{ij}	Access Node 1	Access Node 2	Access Node 3	Access Node 4	Access Node 5
Access Node 1	0.99999	0.99990	0.99995	0.99999	0.99999
Access Node 2	0.99991	0.99999	0.99999	0.99999	0.99999
Access Node 3	0.99995	0.99999	0.99999	0.99998	0.99999
Access Node 4	0.99999	0.99999	0.99998	0.99999	0.99997
Access Node 5	0.99999	0.99999	0.99999	0.99997	0.99999

The following table shows the products of availability and weight for each access point pair ($w_{ij} f_{ij}$):

$w_{ij} f_{ij}$	Access Node 1	Access Node 2	Access Node 3	Access Node 4	Access Node 5
Access Node 1	599.994	1199.880	1799.910	2399.976	2999.970
Access Node 2	1199.892	4799.952	7199.928	9599.904	11999.880
Access Node 3	1799.910	7199.928	10799.892	14399.712	17999.820
Access Node 4	2399.976	9599.904	14399.712	19199.808	23999.280
Access Node 5	2999.970	11999.880	17999.820	23999.280	29999.700

The sum of the values in this table is $\sum_{i=1}^N \sum_{j=1}^N w_{ij} f_{ij} = 252595.878$.

The end-to-end service availability for this network over the year considered is:

$$\frac{\sum_{i=1}^N \sum_{j=1}^N w_{ij} f_{ij}}{\sum_{i=1}^N \sum_{j=1}^N w_{ij}} = \frac{252600}{252595.878} = 0.99998368.$$

6 EQUIVALENT DEFINITION OF END-TO-END SERVICE AVAILABILITY

6.1 Definition in Terms of Weighted Downtime

The definition of end-to-end service availability is given in terms of a summation over access point pairs. This is useful conceptually as it defines the end-to-end service availability for the network as a weighted average of the service availability of access point pairs. At times, it can be useful to apply the definition in terms of weighted downtime:

$$\text{End-to-End Service Availability} = \frac{\sum_{i=1}^N \sum_{j=1}^N w_{ij} f_{ij}}{\sum_{i=1}^N \sum_{j=1}^N w_{ij}} = 1 - \frac{\sum_{m=1}^M S_m O_m}{\sum_{i=1}^N \sum_{j=1}^N w_{ij}}$$

where M service outages with durations $O_1, O_2, O_3, \dots, O_M$, occurred in the network during the time period $[T_s, T_E]$. Each outage duration has a corresponding weight S_m reflecting the number of access point pairs affected by the outage and the sizes of the access points affected:

$$S_m = \sum_{i=1}^N \sum_{j=1}^N I_{ijm} S_i S_j$$

where $I_{ijm} = 1$ if service from access point i to access point j was unavailable from outage m and 0 otherwise.

6.2 Example of Prediction of End-to-End Service Availability

The above definition of end-to-end service availability as the weighted sum of outage durations is useful as the basis for prediction. Consider a communications network with $N = 50,000$ access points each serving 1,000 users. The network also consists of 200 distribution network elements which provide the access points with an interface to a core backbone IP network. Modeling of the network predicts an average of 3302.2 service outages per year as a result of failures in the network. The table below describes the service outages resulting from these failures:

ATIS-0100016

<i>Failure</i>	Outages	Duration (mins)	Access Point Pairs Affected	Product
<i>Access Point</i>	3000	90	99999	26999730000
<i>Distribution Facility</i>	300	300	99999	8999910000
<i>Distribution Element</i>	2	180	24937500	8977500000
<i>Core</i>	0.2	180	250000000	9000000000
SUM=				53977140000

The first line indicates that modeling predicts an average of 3000 service outages per year, each having a duration of 90 minutes on average; each outage affects 99,999 access point pairs (50,000 pairs originating from the failed access point and 49,999 terminating to the access point, thus eliminating double counting of the failed access point). The other lines of the table are calculated in a similar fashion. The final column is the product of the other three columns. Dividing the sum of this column by the number of access point pair minutes in a year:

$$\text{Access point pair minutes in a year} = 50,000^2 \times 365 \times 24 \times 60 = 131400000000000$$

and subtracting the result from 1 gives a predicted end-to-end service availability of 99.9959%.

NOTE - To simplify the calculation of the example, the access point size of 1,000 users has been ignored; this can be done here since all access points are assumed to have the same size.

Assuming Poisson distributions for the numbers of outages per year and exponential distributions for outage durations, a distribution of end-to-end service availability for this network over a year of service can be estimated using Monte Carlo simulation. The table below shows key quantiles for the distribution of availability in the example:

Quantile	.005	.025	.05	.5	.95	.975	.995
Availability (%)	99.9806	99.9891	99.9915	99.9966	99.9973	99.9973	99.9974

The distribution shows that in about half of the years, end-to-end service availability is slightly better than the average of 99.9959%. However, it also shows that a significant chance exists for this availability to be much lower than the average. The skewed nature of this distribution should be considered when setting SLAs for the network.

7 ESTIMATION OF END-TO-END SERVICE AVAILABILITY

Exact measurement of end-to-end service availability for a network over a specified time period would require continual monitoring of the service capability for each pair of access points in the network. In

most cases, such a measurement process would be a prohibitively difficult endeavor. In such cases, sufficient results can be achieved through sampling service capability over the user base and time. It may be possible to estimate end-to-end service availability based on user experience (e.g., user calls blocked, packets lost from user attempts), but such results may require statistical adjustment for bias.

8 CONCLUSION

It is recognized that end-to-end service availability is a complex issue and will require careful and extensive deliberations. It is proposed that the work should be progressed in the following manner:

- ◆ *Availability Definitions:* Reach agreement on desired end-to-end service availability definitions/metrics, for a variety of services, over a range of different network types. All assumptions and limitations must be recognized and recorded. Given the range of possible services, it is proposed that PRQC start with the known problem of traditional voice PSTN service and then apply the metric towards Voice over IP (VoIP) services. With the experiences gained, other packet-based services such as Virtual Private Network (VPN) services can be addressed later.
- ◆ *Outage Measurement Methodology and Practical Limitations:* The ability to measure service outages and establish their correlation with network element failures must be clearly examined in thorough detail for a variety of network/element types. The pros and cons of all capabilities such as sampling, packet probes, element-based reporting, customer-driven trouble ticket reporting, etc., must be carefully examined. Limitations on short-term outages (e.g., < 1 minute outage durations) and the eventual impact on service availability must be clearly understood. The ability to accurately measure the impact of “partial outages” (see 5.1) must be studied in detail. The goal is to understand the practical limitations of carrying out the various measurements in support of the availability metric definitions agreed to as part of the first step.

It is proposed that the full range of work may require several documents. Given the complexity of this topic, that may be the best way forward. It is suggested that, initially, PRQC should produce these documents in the form of TRs. When greater understanding of the topic has been achieved, the reports can be considered for formal standardization.

9 ABBREVIATIONS AND ACRONYMS

ANSI	American National Standards Institute
ATIS	Alliance for Telecommunications Industry Solutions
ATM	Asynchronous Transfer Mode
IETF	Internet Engineering Task Force
IP	Internet Protocol
ITU-T	International Telecommunications Union - Telecommunications Standardization Sector
PRQC	Network Performance, Reliability, and Quality of Service Committee
PSTN	Public Switched Telephone Network

ATIS-0100016

QoS	Quality of Service
SLA	Service Level Agreement
VoIP	Voice over IP
VPN	Virtual Private Network