ATIS-0100028

NETWORK RESILIENCY PLANNING FOR ENTERPRISE CUSTOMERS

TECHNICAL REPORT

ATIS is the leading technical planning and standards development organization committed to the rapid development of global, market-driven standards for the information, entertainment and communications industry. More than 200 companies actively formulate standards in ATIS' Committees, covering issues including: IPTV, Cloud Services, Energy Efficiency, IP-Based and Wireless Technologies, Quality of Service, Billing and Operational Support, Emergency Services, Architectural Platforms and Emerging Networks. In addition, numerous Incubators, Focus and Exploratory Groups address evolving industry priorities including Smart Grid, Machine-to-Machine, Networked Car, IP Downloadable Security, Policy Management and Network Optimization.

ATIS is the North American Organizational Partner for the 3rd Generation Partnership Project (3GPP), a member and major U.S. contributor to the International Telecommunication Union (ITU) Radio and Telecommunications' Sectors, and a member of the Inter-American Telecommunication Commission (CITEL). ATIS is accredited by the American National Standards Institute (ANSI). For more information, please visit < http://www.atis.org >.

ATIS-0100028, *Network Resiliency Planning for Enterprise Customers*

Is an ATIS Standard developed by 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

ATIS-0100028

Technical Report on

NETWORK RESILIENCY PLANNING FOR ENTERPRISE CUSTOMERS

Alliance for Telecommunications Industry Solutions

Approved April, 2010

Abstract

This Technical Report provides an overview of the network resiliency design process for Enterprise Customers. The design process examines a variety of resiliency options for all customer sites depending on site reliability requirements. The key metric driving this process is site availability. A service provider can then provide a range of resiliency options for connecting all customer sites together over the service provider's network.

FOREWORD

The Alliance for Telecommunication Industry Solutions (ATIS) serves the public through improved understanding between providers, customers, and manufacturers. The Network Performance, Reliability, and Quality of Service Committee (PRQC) 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.

The mandatory requirements are designated by the word *shall* and recommendations by the word *should*. Where both a mandatory requirement and a recommendation are specified for the same criterion, the recommendation represents a goal currently identifiable as having distinct compatibility or performance advantages. The word *may* denotes a optional capability that could augment the standard. The standard is fully functional without the incorporation of this optional capability.

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 Technical Report, had the following members:

M. Neibert, PRQC Chair (Telcordia)

P. Tarapre, PRQC Vice Chair (AT&T)

C. Underkoffler, ATIS Chief Editor

M. Linnell, Technical Editor (Telcordia)

S. Makris, Technical Editor (Telcordia)

Active Participants:

C. Dvorak

E. Geelen

Y. Kogan

B. Linick

M. Linnell

S. Makris

A. Morton

M. Niebert

A. Nguyen

H. Pant

E. Rojek

J. Schiavone

P. Tarapore

A. Webster

TABLE OF CONTENTS

TABLE OF FIGURES

TABLE OF TABLES

Technical Report on

Network Resiliency Planning for Enterprise Customers

1 SCOPE AND PURPOSE

Availability is a key measure in Service Level Agreements (SLA) between service providers and their customers as well as their vendors and suppliers. Several availability efforts have been addressed by various standards bodies. They can be briefly described as follows.

In the area of network availability, the access availability of routers in IP-based networks has been defined in terms of the availability of customer-facing ports [T1.TR.78-2003] and, more recently, in terms of the availability of customer-facing line cards [ATIS-0100025].

In the area of service availability:

- The availability of Voice over IP (VoIP) services was defined in terms of a Defects Per Million (DPM) metric [A-0100008]. An estimation methodology outline is also provided in this standard.
- In the ITU-T, general definitions of IP service availability [Y.1540] and MPLS service availability [Y.1561] have been defined. In addition, Ethernet service availability definition work is also ongoing.

From the perspective of service availability, a service provider needs to determine how best to deliver the desired availability to customers based on requirements specified in SLAs. Customers who agree to SLA-specified availability requirements include:

- Government Agencies: Agencies at the Federal, state, and local governments typically buy several types of telecommunications services from service providers and hence, they specify many requirements including service availability. The US Federal Government also specifies general requirements that dictate how large service providers are supposed to provide basic telephone service to the population.
- Enterprise Customers: Large corporations, financial institutions, universities, etc, are examples of Enterprise Customers. Such customers agree to very specific requirements for the type of service desired. Service availability is a key requirement for such customers.

The delivery of specific services to such customers typically is classified as “Enterprise Services”. A service provider can manage availability as follows:

- Design appropriate network resources with appropriate reliability features to deliver the required level of availability.
- Implement necessary measurement capabilities to monitor and estimate the availability of the customer’s service over specified time periods (e.g., monthly, annually, etc).

Both actions require significant investment in terms of planning and cost. Service providers may choose to strongly address both options. Alternately, a service provider may choose to invest heavily in the planning stage whereby significant efforts are dedicated to ensure that sufficient resource allocation, coupled with appropriate level of diversity and redundancy, is built into the network design [ATT-OFC] such that the desired service availability can be satisfactorily achieved. This may mitigate the need for sophisticated and expensive monitoring and measurement capabilities for estimating service availability on the “back end”.

This Technical Report provides an overview of the network resiliency design process for Enterprise Customers.

2 NORMATIVE REFERENCES

The following standards contain provisions which, through reference in this text, constitute provisions of this American National Standard. At the time of publication, the editions indicated were valid. All standards are subject to revision, and parties to agreements based on this American National Standard are encouraged to investigate the possibility of applying the most recent editions of the standards indicated below.

[ATIS-0100025] ATIS-0100025 (2009), *Estimating Availability of Access IP Routers in Terms of Customer Facing Line Card Availability.*[1]

[A-0100008] ATIS-01000082007, *DPM Metric for Transaction Services such as VoIP*. [1]

[ATT-OFC] AT&T Presentation at Optical Fiber Communications Exposition (2008), B. H. Linick & M. A. Lazer, *Peeling the Reliability Onion: Telecommunications Service Reliability*.

[T1.TR.78-2003] T1.TR.78 (2003), *Access Availability of Routers in IP Networks*. [1]

[ATIS-0900105] ATIS-0900105.01.2000(R2010), *Telecommunications – Synchronous Optical Network (SONET) – Automatic Protection Switching*. [1]

[ATIS-0300231] ATIS-0300231.2003(R2007), *Digital Hierarchy – Layer 1 In-Service Digital Transmission Performance Monitoring*. [1]

[ATIS-0300231.04] ATIS-0300231.04.2003), *SONET – Layer 1 In-Service Digital Transmission Performance*. [1]

[G.8080] ITU-T Recommendation G.8080/Y.1304 (2003), *Architecture of the Automatic Switched Optical Network (ASON)*.[2]

[Y.1540] ITU-T Recommendation Y.1540 (2007), *IP Packet Transfer and Availability Performance Parameters*. [2]

[Y.1561] ITU-T Recommendation Y.1561 (2004), *Performance and Availability Parameters for MPLS Networks*. [2]

[RFC4090] IETF RFC 4090 (2005), *Fast Reroute Extensions to RSVP-TE for LSP Tunnels*.[3]

3 DEFINITIONS

Enterprise Customer Site: An Enterprise Customer can have multiple site locations. These sites can range from:

- National Center: Customer Headquarters or Large Data Center
- Regional Center: Regional Headquarters
- Local Office: Local center serving a handful of field agents or possibly a single agent home office

Site Availability: For the purpose of this document, the availability of an Enterprise Customer's site is defined as the fraction of operating time over which the site is successfully connected to the edge of the service provider's network.

---

[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/ >

[3] This document is available from the Internet Engineering Task Force (IETF). < http://www.ietf.org >

4 ACRONYMS & ABBREVIATIONS

ADM	Add Drop Multiplexer
ATM	Asynchronous Transfer Mode
CPE	Customer Premise Equipment
DCS	Digital Cross-Connect Switch
DPM	Defects Per Million
DRP	Disaster Recovery Plan
DS-1/DS-3	Digital Signal Level-1/Digital Signal Level-3
IP	Internet Protocol
LATA	Local Access and Transport Area
LSO	Local Serving Office
MPLS	Multi-Protocol Label Switching
MTTF	Mean-Time-To-Failure
NE	Network Element
POP	Point of Presence
SLA	Service Level Agreement
SPoF	Single Point of Failure
STS-n	Synchronous Transport Signal-n
VoIP	Voice over IP

5 INTRODUCTION

Increasingly, Enterprise Customers run applications critical to their businesses on services procured from telecommunications service providers and they are requesting high availability services. Telecommunications service providers meet these needs by providing several high resiliency service options to the enterprise customer.

Service resilience has always been an area of critical focus for the carriers. Historically, the industry has coined the expression “carrier grade” referring to the availability demands placed by carriers on the equipment they deploy. Carrier grade equipment was sufficiently reliable to support the availability commitments carriers made to their customers. This commitment often takes the form of a promise to provide “five nines” of availability. Five nines, or availability of 99.999%, translate to about 5 minutes of outage per year [ATIS-0300231], [ATIS-0300231.04]. This takes into account all the incidents that could disrupt communications services including hardware failures, fiber cuts, or software failures. The notion of availability can be expanded to address those instances where a customer may have redundant access to network services. In this case, it would be appropriate to define an availability metric to reflect site availability which is the ability of the site to reach the network even when some of the redundant access connections have failed. Additional measures that reflect the probability of failures of some duration may also be useful when describing service resilience for enterprise customers.

This Technical Report explores the different aspects of “carrier grade” networks and services, by peeling the reliability onion, one layer at a time, and then provides guidance on the network design process necessary to achieve the desired degree of service availability for Enterprise Customers.

In today’s environment, customers’ business operations and foot prints are increasingly decentralized and reliant on a single converged telecommunications network. Failures in a customer’s enterprise network could cause interruptions of operations with severe and far reaching consequences. Customers now expect that their key services will be available nearly 100% of the time, and service providers must meet this expectation in order to be successful in today’s marketplace.

Meeting these expectations requires careful planning and design at multiple levels. It entails multiple reliability layers starting with the outermost layers of customer's business applications progressing inwards towards the innermost layer of the core of carrier's network.

Five major areas for consideration are:

- The design of the enterprise network and the selection of the proper reliability options.
- The definition and application of reliability metrics that allow enterprise customers to understand the potential for service and business operations disruptions.
- The design of a set of service offerings with a range of reliability options available.
- The design of business processes consistent with business continuity requirements.
- The availability of a robust infrastructure supporting these reliable services offerings.

These five areas encompass the network resiliency design process – "peeling the reliability onion". This process is outlined in Clause 6 and an illustrative example is provided in Clause 7.

6 RESILIENCY DESIGN PROCESS FOR ENTERPRISE CUSTOMERS

Government agencies, large corporations, financial institutions, universities, etc, are examples of Enterprise Customers. Such customers specify strict SLAs for reliability and resiliency for the type of service desired. One key requirement is a high degree of Service Availability. It is incumbent on the service provider to implement the necessary mechanisms that ensure high levels of network reliability and resiliency in order to deliver the agreed upon level of service availability to Enterprise Customers. This Clause provides an overview of the necessary steps by which a service provider "peels the reliability onion" [ATT-OFC] to accomplish this task.

Typically, Enterprise Customers are characterized by their locations or sites which then need to be connected together in a reliable manner. It is up to the service provider to craft a suitable network design that meets the reliability and resiliency criteria subject to the customer's telecommunications budget.

To illustrate the resiliency design process, it is assumed that a typical Enterprise Customer operates the following types of customer sites or locations:

- National Centers: This type of site is typically a headquarters type of location or a large data center for the customer's network. Such a site would have the most stringent requirements for service availability in terms of being connected to all other customer sites.
- Regional Centers: Typically such sites serve a given region for the customer's network and would require reliable connectivity to the national center as well as to all local offices in the region.
- Local Offices: Typically these sites would be small centers that may house a handful of customer agents or perhaps even a single agent home office. Reliability requirements for connecting such offices to regional or national centers are important; however they may not be as stringent as those for the regional or national centers.

The Enterprise Customer is thus expected to have a range of site availability requirements depending on the number and type of site locations. The service provider needs to create a design process that meets these site availability requirements for the Enterprise Customer.

Large service providers design their core backbone networks with high levels of redundancy built in for network elements (e.g., routers, cross-connects, etc) as well as transport facilities linking these elements. Hence, core backbone network failures typically do not result in extended service outages. The exception may be large scale disasters such as earthquakes, hurricanes, etc, or terrorist attacks

that may cause substantial damage to a network over a large region[4]. Bottlenecks occur at the edges of the network under normal conditions and hence, the main issue is to ensure accessibility from all the Enterprise Customer sites to the edge of the core backbone of the service provider network.

Clauses 6.1 – 6.3 provide descriptions on various methodologies and processes by which networks in general and core backbone networks in particular, can achieve high levels of reliability and resiliency. Clause 6.4 defines a key availability metric that is utilized for designing suitable network access to the core backbone network for all customer sites. Clause 6.5 describes the design process for connecting customer sites to the core backbone network.

6.1 Service Offerings and Reliability Options

In general, services for Enterprise Customers must provide a high degree of resiliency. These services are designed to meet the desired level of service availability (e.g., the "5 nines" availability guarantees the industry is famous for). Resiliency in service design relies on a number of fundamental principles. Chief among these are the notions of *"Single Point of Failure"* (SPoF) avoidance and of self healing networks. A service SPoF is defined in terms of support for any given service. A SPoF is a network component (either software or hardware) whose failure will disrupt that service until the failure is automatically restored or physically repaired. Good examples of this are the ingress and egress nodes for a given service path. In the rare event that one of these nodes is lost, the self healing capabilities of the network would not be able to restore those service paths (see Figure 1).

If the SPoFs introduce more than acceptable risk, the risk may be mitigated by using diversity based service options. Such services allow for the complete physical separation of groups of circuits so that no single failure will disrupt more than any one of the group. Services based on routing diversity generally imply that the customer contracts for twice as many paths as are required by the traffic load. In the ideal case, one set of paths is fully diverse[5] from the other. In some instances, one set of paths serves as a backup while the other carries all traffic and the path switching is under customer control. In other cases, traffic may be load shared across both sets of paths. Diversity arrangements can be designed to avoid most SPoFs, except the customer location (see Figure 1), but they are highly resource intensive requiring considerable excess capacity. Diversity options are useful for designing access from the customer site to the core backbone network.

[4] Re-routing traffic seamlessly with automated recovery mechanisms is not feasible in such cases and the only mitigation is to install disaster recovery techniques (see Clause 6.4).

[5] In some cases, depending on the available topology (e.g., physical spur) complete physical diversity may not be possible.

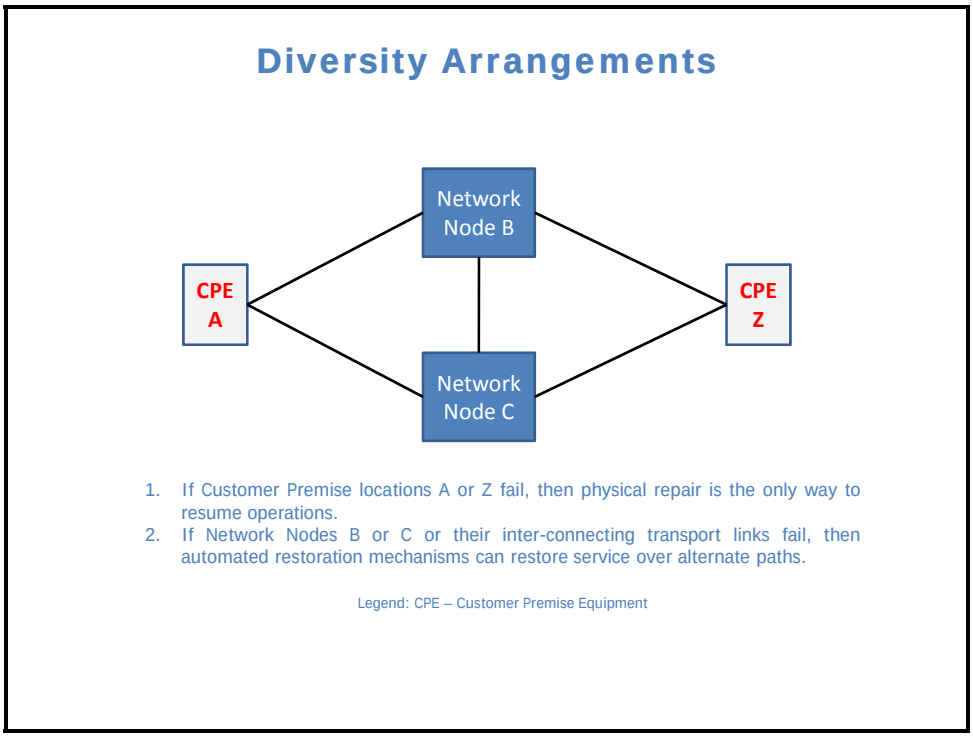

Figure 1 - Illustrative Example of Diversity Arrangements

Self-healing services are designed to react to a failure by automatically avoiding the failed elements. These mechanisms exist at multiple network layers and are based on automatic restoration techniques. Restoration mechanisms widely deployed to support self-healing services are based on self-healing ring architectures, or on mesh networks with fast re-routing capabilities around failures. The notion of re-routing is applicable to transport as well as voice and data networks. SONET Ring architectures are useful for network access while mesh restoration methods are typically deployed in core backbone networks.

Resilient service options need to be carefully designed to afford customers the appropriate degree of service continuity. In the transport services world, this means offering services that are guaranteed to be quickly restored in the event of a failure (e.g., SONET Ring re-routing around a cable cut/intrusion). Services that ride on top of the transport layer may add additional layers of service specific resilience.

Ensuring end-to-end service resilience support may be fairly complex as the end-to-end service path includes multiple network segments (see Figure 2). The end-to-end path can be viewed as two access segments on either side of a core backbone (service) network. Note that the access segment may include a backhaul segment to reach the provider network edge.

The access segments provide connectivity from the customer premise to the edge of the core backbone network. Resiliency within the core backbone can be provided by a variety of mechanisms including redundancy in elements and fast rerouting. Resiliency in the connectivity to the service network can be achieved via diverse or resilient access coupled with dual homing options, thus avoiding the creation of a SPoF at edge elements. Different dual-homing options allow the customer's traffic to be homed to either separate ports on the same switch, to separate switches in the same location, or to switches in separate locations, depending on the customer requirements. Maximum resiliency is achieved when each network segment is in and of itself resilient, and there are no SPoFs between them. As service SPoFs are removed by combining dual homing with resilient access and resilient core backbone networks, customers are provided with the highest possible availability.

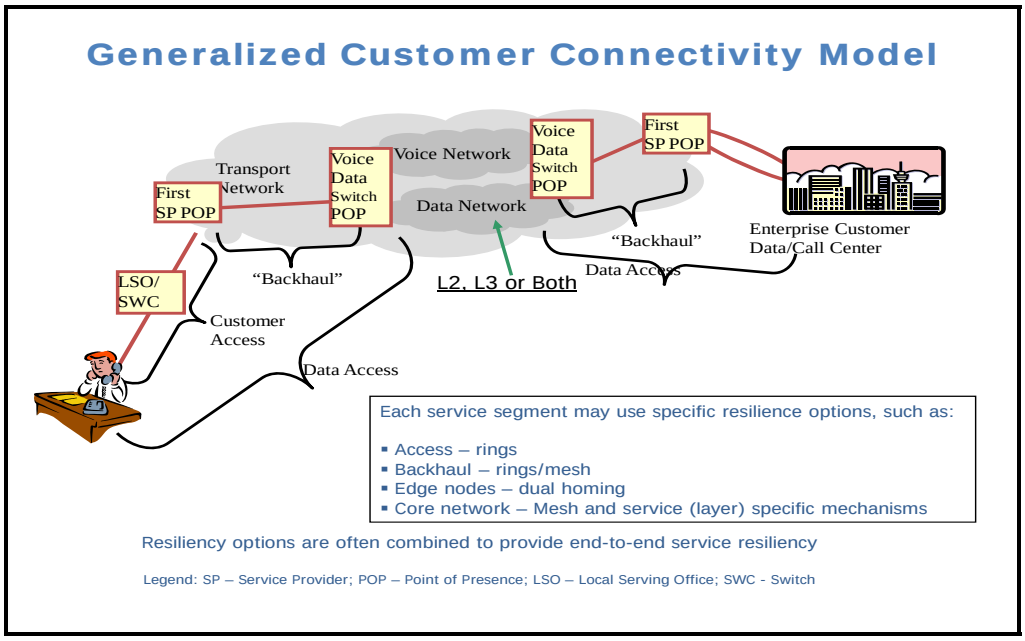

Figure 2 - Generalized Enterprise Customer Connectivity Segments

6.2 Business Processes

The realization of an enterprise network that meets customer needs can only be achieved when there is a partnership between the customer and the telecommunications services provider. Key to this partnership is the sharing of information between carrier and customer. The customer shares information about its business process designs, key sensitivities, and customer site information - number and type of customer sites/offices and the availability requirements for each site.

The service provider shares data about its services, service options and their cost, and service capabilities specific to geographic location. It then creates a network design that seeks to optimize the cost to the Enterprise Customer such that the site availability requirements of the customer are satisfactorily met. To do this, the service provider takes into consideration a range of service reliability options for the critical segments of the customer's network. Some of the high-resiliency options can be resource intensive, and as such carry a higher cost for the Enterprise Customer.

Customers must be active participants in developing end to end solutions that will meet their business needs. This entails the design and execution of an appropriate set of business processes. The customer sites that house critical business processes can be important SPoFs. Customers must ensure that critical sites are not SPoFs by providing alternate locations and putting into place all the processes needed to ensure that multiple sites have updated data bases, are continually ready, and are properly staffed. Disaster recovery plans are needed to allow the transfer of business operations from one site to the other.

Ideally, the selection of sites for critical operations centers includes telecommunications resiliency considerations. This requires a partnership with the service provider to identify service and resiliency options available, and access alternatives. Often, LATA boundaries, tariffs and the availability of fiber are factors in assessing a location from a telecommunications reliability perspective. In some instances, Enterprise Customers contract with multiple service providers to mitigate the risk of business operations interruptions. However, in many situations this has proven to be a less effective method than a strong partnership between the Enterprise Customer and its network provider. To a large extent, this is because many right-of-ways are shared between carriers and common conditions will cause failures for multiple providers' networks. A good example of this situation is the 2006 earthquake in the Luzon

Strait off the coast of Taiwan, in which facilities for multiple providers of trans-Pacific networks were impacted.

6.3 Infrastructure Support

With more and more critical applications running globally, across cities, countries or continents, it becomes imperative that networking services deliver the reliability sought by enterprises. As previously noted, service reliability is supported by risk mitigation at different levels. High-end communications services often come with a slew of options for mitigating the risk factors (from simple fiber cuts, to catastrophic events leading to loss of central offices or enterprise locations) for improving service reliability. These services must be carried in an infrastructure equally resilient to failure. Several key network resiliency-related mechanisms of redundancy, restoration and rerouting employed in modern telecommunications networks are briefly discussed here. These mechanisms are widely deployed in core backbone networks.

6.3.1 Redundancy Mechanisms

There are multiple types of redundancy used extensively in the telecommunications networks including equipment and facility redundancy. While redundant equipment and facilities are most often thought of being used to protect against failures, they are also useful in preserving service when scheduled maintenance activities are required.

Equipment redundancy has been used extensively to protect against equipment outages which may occur due to software or hardware failures. In order to maximize network resilience, network element architectures have evolved to provide complete redundancy for critical core subsystems such as power supplies and common controllers. Service supporting customer facing line-cards may also be configured in a 1+1 or 1:N protection scheme based upon the level of resiliency required. Network elements may be deployed in a redundant manner, coupled with dual connectivity between nodes and advanced routing mechanisms. Often dual connectivity is coupled with diverse routing of the connecting links to provide end-to-end protection. Facility redundancy requires route diversity and is generally used either in conjunction with restoration or dual connectivity as described below.

6.3.2 Restoration and Re-Routing

At Layer 1, restoration is used primarily to mitigate degradation or loss of signal transmission, such as that which would result from fiber intrusions or other failures. As a practical point, restoration schemes use re-routing capabilities in network nodes to accomplish transmission restoration. They use the network capacity more efficiently than an approach based solely on using diverse connectivity between nodes. In addition, restoration can also mitigate risks associated with losing intermediate nodes, by providing rerouting capabilities around a failed node. Note however that restoration is never instantaneous – there is added value in having diversity arrangements.

Facility restoration mechanisms have been in place and have evolved over many years. SONET/SDH self-healing ring technology is bandwidth intensive and is constrained by a maximum ring circumference and number of nodes that can be accommodated while still meeting restoration time targets of 50 ms (for span switching) and 200 ms (for loopback switching) [ATIS-0900105]. Technology and economic considerations have prompted development of alternatives based on mesh topology for core backbone networks. Intelligent cross-connect nodes capable of supporting computational intensive re-routing deliver sub-second restoration in mesh networks [G.8080]. Regardless of the technique, effective facility restoration requires careful capacity planning and monitoring to ensure the presence of sufficient restoration capacity. This is easiest for self-healing rings, but often requires modeling and simulation for mesh-based restoration schemes.

There are similar mechanisms used at higher network layers as well. Resilient Packet rings use self-healing ring mechanisms for Ethernet traffic. Ethernet switches, ATM switches and IP routers use mesh based re-routing to by-pass failed links or failed nodes in the layer 3 network. Mesh-based re-routing such as MPLS Fast Re-route [RFC4090] is available from several equipment vendors. Restoration

times vary, as the algorithms tend to be proprietary and restoration performance may degrade when large networks suffer major outages.

With restoration available and in–use at multiple layers, it is essential that restoration is coordinated properly across multiple layers. In some instances higher layer devices such as routers or ATM switches are configured with a delay interval that ensures that the higher layer re-routing occurs only if the lower layer restoration mechanisms have not provided a restoration path. In other instances, the lower layer mechanisms may be by-passed entirely.

6.3.3 Disaster Recovery

Any carrier's network and customer business operations may face interruptions caused by natural disasters (e.g., earthquakes, fires, floods, volcanic eruptions) and man made disasters (e.g. sabotage, terrorism or civil unrest). Events such as these are likely to cause multiple failures overwhelming the normal capabilities of the carriers and customers' operational processes, and Disaster Recovery Plans (DRP) need to be executed by both parties.

In order to mitigate potentially very large and lengthy disruptions that may be caused by disaster situations, carriers have developed DRP capabilities. These capabilities support the long-term deployment of specialized equipment and resources. The DRP equipment must be exercised periodically and the staff that would deploy and operate this equipment must be in a constant state of readiness. This is often accomplished by the execution of DRP drills, where equipment is deployed and operated in response to simulated disasters. While actual disasters are rare, well planned and tested DRPs need to be established by carrier and customer alike.

6.4 Resiliency Design Metrics

As mentioned above, a service provider's core backbone network is designed and deployed with multiple layers of redundancy coupled with automated re-routing capabilities resulting in very high levels of network reliability. The key issue for reliable network design for Enterprise Customers is the creation of suitable access options that link the Enterprise Customer's sites to the core backbone network edges.

The key metric that governs the enterprise network design process is the set of site availabilities, for all sites run by the Enterprise Customer. This metric is defined as follows.

Assumptions:

- Service provider's core backbone network is sufficiently robust from a reliability and resiliency standpoint.
- The connectivity from a customer site to the edge of the service provider's network is done via a "Private Line" of a pre-determined bandwidth (e.g., DS-1, DS-3, STS-3, etc). This Private Line connection is either "Up" signifying successful connectivity from the site to the core network or "Down" signifying a failed connection. In the case where connectivity from a critical site (e.g., National Headquarters) to the service provider's network is done via redundant Private Line connections, the "Up" state implies that at least one connection is up and running.

Then:

Availability for Site i:

$$A_i = \frac{T_i}{T}$$

Where:

- *A_i is the Desired Availability for Site i.*
- *T_i is the Total Time over the Period T for which Site i is successfully connected to the edge of the network – Private Line connection is in the "Up" state.*
- *T is the Total Operating Time Period Measured in Months or Years.*

The service provider needs to ensure the design of Private Line connections for the Enterprise Customer is in conformance with the set of Site Availabilities $[A_i]_{i=1,...,N}$, subject to the customer's cost specifications.

Other metrics that are also utilized in assessing the effectiveness of network access designs are as follows.

Mean–Time-To-Failure (MTTF) for:

- Restoration Event: Mean time to occurrence of an event lasting up to *n* seconds, where *n* is based on the restoration technology and the subsequent restoration speed.
- Site Outage: Mean time to occurrence of simultaneous failure of all circuits serving a customer site, for a duration exceeding *n* seconds.

Probability of Site Outage *P(N)*: The probability of a site outage > *N* seconds over Time Period *T*, where *T* is an interval such as one month or one year, and *N* > *n*. Note that per the Site Outage definition above , *n* seconds is considered to be the minimum amount of down time experienced by the site in order to be considered as unavailable.

Probability of Not Meeting Availability Objective: Probability of a total outage occurring during a specified Period T that would cause 99.99%, 99.999% or 99.9999% availability objectives to be violated.

These additional metrics are useful in fine tuning network designs that are determined via the process described in Clause 6.5. Their significance and related impacts are determined in detailed interactions between the service provider and the Enterprise Customer.

6.5 Enterprise Network Design – The Final Step

Typically, core network segments tend to be designed with significant layers of redundancy for network elements (routers, cross-connects, etc) as well as transport facilities.

The degree of redundancy – and hence, resiliency – for individual access and backhaul segments depend on the type of customer site and on the customer's willingness to support extended resiliency from a cost perspective.

Figure 3 depicts a range of access connectivity options of differing levels of resiliency:

- Option 1 has no redundancy and may be deployed for an individual's home office.
- Options 2 and 3 have varying levels of redundancy for the connections' transport legs, which are implemented as SONET Rings, as well as some network element redundancy (Option 3 – dual homing at the provider's POP). These options may be suitable for local offices serving multiple customer agents.
- Option 4 has extended levels of redundancy for the connection's transport legs, which are implemented as SONET Rings, as well as for network elements (dual homing for Local Serving Offices and Provider POPs) and it is recommended for the regional and national centers.

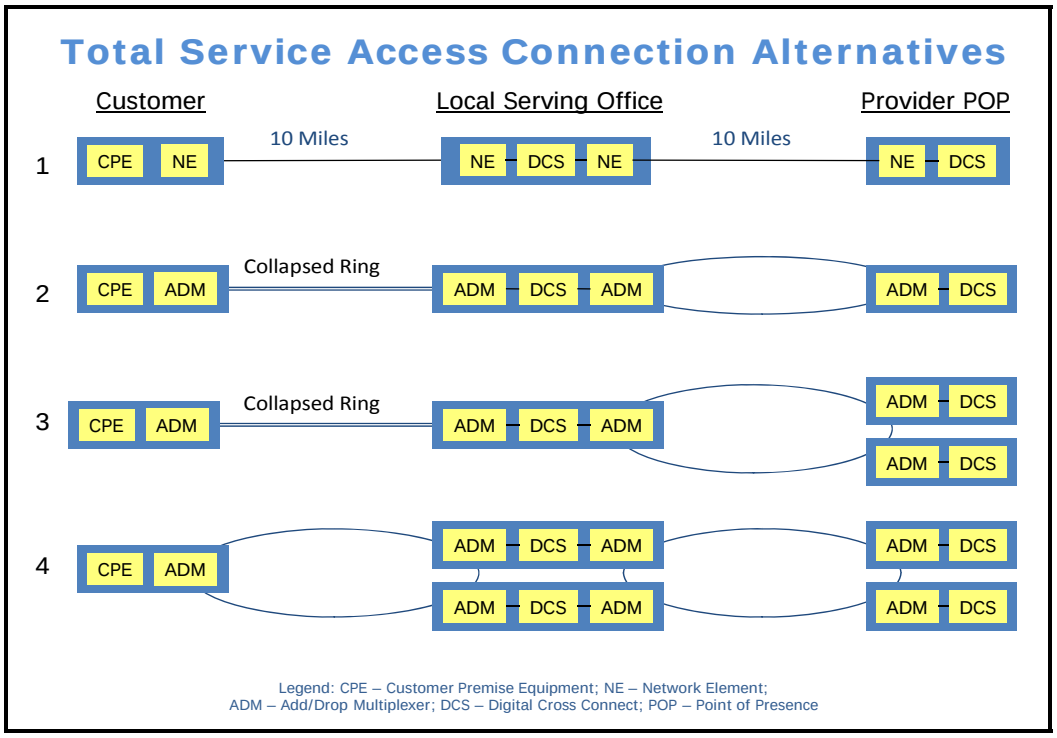

Figure 3 - Access Connectivity Resiliency Options

Similar options can be provided for the backhaul segments of the connections. The design process for the (end-to-core edge) access and backhaul portions of site connections can then be briefly described as follows:

- Select desired reference connections separately for access and backhaul segments based on input from customer and type of customer site.
- Populate reference connections with network elements of desired level of redundancy along with their respective failure and repair rates.
- Specify all applicable failure modes and reliability logic for reference connections.
- Create Markov state transition models for reference connections and execute models to predict customer site availability and other reliability metrics (see Clause 6.4). Other techniques such as simulation modeling may also be utilized as appropriate.
- Pair together appropriate combinations of access and backhaul connection designs.
- Aggregate reliability metrics for pairs of access and backhaul connection designs into metrics for end-to-edge connections.

A service provider can then provide a range of resiliency options for connecting all customer sites together over the service provider's network. It is then up to the customer to balance reliability requirements versus their telecommunications budget and arrive at a suitable solution.

7 ILLUSTRATIVE EXAMPLE

To illustrate the approach at a high level, consider two extreme cases, where one is a single circuit extending from the customer premise to the data switch POP and the other is a pair of circuits traversing a highly resilient architecture.

Case 1 corresponds to Option 1 of Figure 2 for the access segment, with the single circuit also extending 400 miles across a restorable optical mesh network for the backhaul segment.

Case Two corresponds to option 4 of Figure 2 for the access segment, with the pair of circuits also extending 400 miles across a restorable optical mesh network for the backhaul segment, through which the pair of circuits are assumed to be physically diverse. Each of the two circuits connects to the provider's data network via separate provider edge routers at physically diverse locations.

A typical set of metrics and their corresponding values for the two cases is shown in Table 1.

Table 1 - Typical Metrics & Values for the Two Cases

Case	Availability	MTTF in Years for:		Probability of:			
		Backhaul Restoration Event	Site Outage	Site Outage Occurring During 1 Year	Not Meeting Availability Objective During 1 Month		
					0.9999	0.99999	0.999999
1	0.9996	0.36	1.1	0.58	0.057	0.068	0.258
2	0.9999995	0.18	911	0.0011	0.0001	0.0001	0.0008