Network Working Group
Internet Draft
Category: Informational
Expires: April, 2002

G. Ash
AT&T Labs
October, 2001

Traffic Engineering & QoS Methods for IP-, ATM-, & TDM-Based Multiservice Networks

<draft-ietf-tewg-qos-routing-04.txt>

ABSTRACT

This is an informational document submitted in response to a request from the IETF Traffic Engineering Working Group (TEWG) for service provider uses, requirements, and desires for traffic engineering best current practices. As such, the work sets a direction for routing and traffic performance management in networks based on traffic engineering (TE) and QoS best current practices and operational experience, such as used in the AT&T dynamic routing/class-of-service network. Analysis models are used to demonstrate that these currently operational TE/QoS methods and best current practices are extensible to Internet TE and packet networks in general. The document describes, analyzes, and recommends TE methods which control a network's response to traffic demands and other stimuli, such as link failures or node failures. These TE methods include:

a) traffic management through control of routing functions, which include call routing, connection routing, QoS resource management, routing table management, and dynamic transport routing.
b) capacity management through control of network design, including routing design.
c) TE operational requirements for traffic management and capacity management, including forecasting, performance monitoring, and short-term network adjustment.

TABLE OF CONTENTS

1.0 Introduction

This is an informational document submitted in response to a request from the IETF Traffic Engineering Working Group (TEWG) for service provider uses, requirements, and desires for traffic engineering best current practices. As such, the work sets a direction for routing and traffic performance management in networks based on traffic engineering (TE) and QoS best current practices and operational experience, such as used in the AT&T dynamic routing/class-of-service network [A98]. Analysis models are used to demonstrate that these currently operational TE/QoS methods and best current practices are extensible to Internet TE and packet networks in general.

TE is an indispensable network function which controls a network's response to traffic demands and other stimuli, such as network failures. TE encompasses

- traffic management through control of routing functions, which include number/name translation to routing address, connection routing, routing table management, QoS resource management, and dynamic transport routing.
- capacity management through control of network design.

Current and future networks are rapidly evolving to carry a multitude of voice/ISDN services and packet data services on internet protocol (IP), asynchronous transfer mode (ATM), and time division multiplexing (TDM) networks. The long awaited data revolution is occurring, with the extremely rapid growth of data services such as IP-multimedia and frame-relay services. Within these categories of networks and services supported by IP, ATM, and TDM protocols have evolved various TE methods. The TE mechanisms are covered in the document, and a comparative analysis and performance evaluation of various TE alternatives is presented. Finally, operational requirements for TE implementation are covered.

The recommended TE methods are meant to apply to IP-based, ATM-based, and TDM-based networks, as well as the interworking between these network technologies. Essentially all of the methods recommended are already widely applied in operational networks worldwide, particularly in PSTN networks employing TDM-based technology. However, the TE methods are shown to be extensible to packet-based technologies, that is, to IP-based and ATM-based technologies, and it is important that networks which evolve to employ these packet technologies have a sound foundation of TE methods to apply. Hence, it is the intent that the recommended TE methods in this document be used as a basis for requirements for TE methods, and, as needed, for protocol development in IP-based, ATM-based, and TDM-based networks to implement the TE methods.

Hence the TE methods encompassed in this document include:

- traffic management through control of routing functions, which include call routing (number/name translation to routing address), connection routing, QoS resource management, routing table management, and dynamic transport routing.

- capacity management through control of network design, including routing design.

- TE operational requirements for traffic management and capacity management, including forecasting, performance monitoring, and short-term network adjustment.

Results of analysis models are presented which illustrate the tradeoffs between various TE approaches. Based on the results of these studies as well as established practice and experience, TE methods are recommended for consideration in network evolution to IP-based, ATM-based, and/or TDM-based technologies.

We begin this document with a general model for TE functions, which include traffic management and capacity management functions responding to traffic demands on the network. We then present a traffic-variations model which these TE functions are responding to. Next we outline traffic management functions which include call routing (number/name translation to routing address), connection or bearer-path routing, QoS resource management, routing table management, and dynamic transport routing. These traffic management functions are further developed in ANNEXES 2, 3, 4, and 5. We then outline capacity management functions, which are further developed in ANNEX 6. Finally we briefly summarize TE operational requirements, which are further developed in ANNEX 7.

In ANNEX 2, we present models for call routing, which entails number/name translation to a routing address associated with service requests, and also compare various connection (bearer-path) routing methods. In ANNEX 3, we examine QoS resource management methods in detail, and illustrate per-flow versus per-virtual-network (or per-traffic-trunk or per-bandwidth-pipe) resource management and the realization of multiservice integration with priority routing services. In ANNEX 4, we identify and discuss routing table management approaches. This includes a discussion of TE signaling and information exchange requirements needed for interworking across network types, so that the information exchange at the interface is compatible across network types. In ANNEX 5 we describe methods for dynamic transport routing, which is enabled by the capabilities such as optical cross-connect devices, to dynamically rearrange transport network capacity. In ANNEX 6 we describe principles for TE capacity management, and in ANNEX 7 we present TE operational requirements.

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119.

2.0 Definitions

Alternate Path Routing:	a routing technique where multiple paths, rather than just the shortest path, between a source node and a destination node are utilized to route traffic, which is used to distribute load among multiple paths in the network;
Autonomous System:	a routing domain which has a common administrative authority and consistent internal routing policy. An AS may employ multiple intradomain routing protocols and interfaces to other ASs via a common interdomain routing protocol;
Blocking:	refers to the denial or non-admission of a call or connection-request, based for example on the lack of available resources on a particular link (e.g., link bandwidth or queuing resources);
Call:	generic term to describe the establishment, utilization, and release of a connection (bearer path) or data flow;
Call Routing:	number (or name) translation to routing address(es), perhaps involving use of network servers or intelligent network (IN) databases for service processing;

Circuit Switching	denotes the transfer of an individual set of bits within a TDM time-slot over a connection between an input port and an output port within a given circuit-switching node through the circuit-switching fabric (see Switching)
Class of Service	characteristics of a service such as described by service identity, virtual network, link capability requirements, QoS & traffic threshold parameters;
Connection:	bearer path, label switched path, virtual circuit, and/or virtual path established by call routing and connection routing;
Connection Admission Control (CAC)	a process by which it is determined whether a link or a node has sufficient resources to satisfy the QoS required for a connection or flow. CAC is typically applied by each node in the path of a connection or flow during set-up to check local resource availability;
Connection Routing:	connection establishment through selection of one path from path choices governed by the routing table;
Crankback	a technique where a connection or flow setup is backtracked along the call/connection/flow path up to the first node that can determine an alternative path to the destination node;
Destination Node:	terminating node within a given network;
Flow:	bearer traffic associated with a given connection or connectionless stream having the same originating node, destination node, class of service, and session identification;
GoS (grade of service)	a number of network design variables used to provide a measure of adequacy of a group of resources under specified conditions (e.g., GoS variables may be probability of loss, dial tone delay, etc.)
GoS standards	parameter values assigned as objectives for GoS variables
Integrated Services:	a model which allows for integration of services with various QoS classes, such as key-priority, normal-priority, & best-effort priority services;
Link:	a bandwidth transmission medium between nodes that is engineered as a unit;
Logical Link:	a bandwidth transmission medium of fixed bandwidth (e.g., T1, DS3, OC3, etc.) at the link layer (layer 2) between 2 nodes, established on a path consisting of (possibly several) physical transport links (at layer 1) which are switched, for example, through several optical cross-connect devices;
Node:	a network element (switch, router, exchange) providing switching and routing capabilities, or an aggregation of such network elements representing a network;
Multiservice Network	a network in which various classes of service share the transmission, switching, queuing, management, and other resources of the network;
O-D pair:	an originating node to destination node pair for a given connection/bandwidth-allocation request;
Originating Node:	originating node within a given network;
Packet Switching	denotes the transfer of an individual packet over a connection between an input port and an output port within a given packet-switching node through the packet-switching fabric (see Switching)
Path:	a concatenation of links providing a connection/bandwidth-allocation between an O-D pair;
Physical Transport Link:	a bandwidth transmission medium at the physical layer (layer 1) between 2 nodes, such as on an optical fiber system between terminal equipment used for the transmission of bits or packets (see transport);
Policy-Based Routing	network function which involves the application of rules applied to input parameters to derive a routing table and its associated parameters;
QoS (quality of service)	a set of service requirements to be met by the network while transporting a connection or flow; the collective effect of service performance which determine the degree of satisfaction of a *user* of the *service*
QoS Resource Management	network functions which include class-of-service identification, routing table derivation, connection admission, bandwidth allocation, bandwidth protection, bandwidth reservation, priority routing, and priority queuing;
QoS Routing	see QoS Resource Management;
QoS Variable	any performance variable (such as congestion, delay, etc.) which is perceivable by a user
Route:	a set of paths connecting the same originating node-destination node pair;
Routing	the process of determination, establishment, and use of routing tables to select paths

	between an input port at the ingress network edge and output port at the egress network edge; includes the process of performing both call routing and connection routing (see call routing and connection routing)
Routing Table:	describes the path choices and selection rules to select one path out of the route for a connection/bandwidth-allocation request;
Switching	denotes connection of an input port to an output port within a given node through the switching fabric
Traffic Engineering	encompasses traffic management, capacity management, traffic measurement and modeling, network modeling, and performance analysis;
Traffic Engineering Methods	network functions which support traffic engineering and include call routing, connection routing, QoS resource management, routing table management, and capacity management;
Traffic Stream:	a class of connection requests with the same traffic characteristics;
Traffic Trunk:	an aggregation of traffic flows of the same class which are routed on the same path (see logical link)
Transport	refers to the transmission of bits or packets on the physical layer (layer 1) between 2 nodes, such as on an optical fiber system between terminal equipment (note that this definition is distinct from the IP-protocol terminology of transport as end-to-end connectivity at layer 4, such as with the Transport Control Protocol (TCP))
Via node:	an intermediate node in a path within a given network;

3.0 Traffic Engineering Model

Figure 1.1 illustrates a model for network traffic engineering. The central box represents the network, which can have various architectures and configurations, and the routing tables used within the network. Network configurations could include metropolitan area networks, national intercity networks, and global international networks, which support both hierarchical and nonhierarchical structures and combinations of the two. Routing tables describe the path choices from an originating node to a terminating node, for a connection request for a particular service. Hierarchical and nonhierarchical traffic routing tables are possible, as are fixed routing tables and dynamic routing tables. Routing tables are used for a multiplicity of traffic and transport services on the telecommunications network.

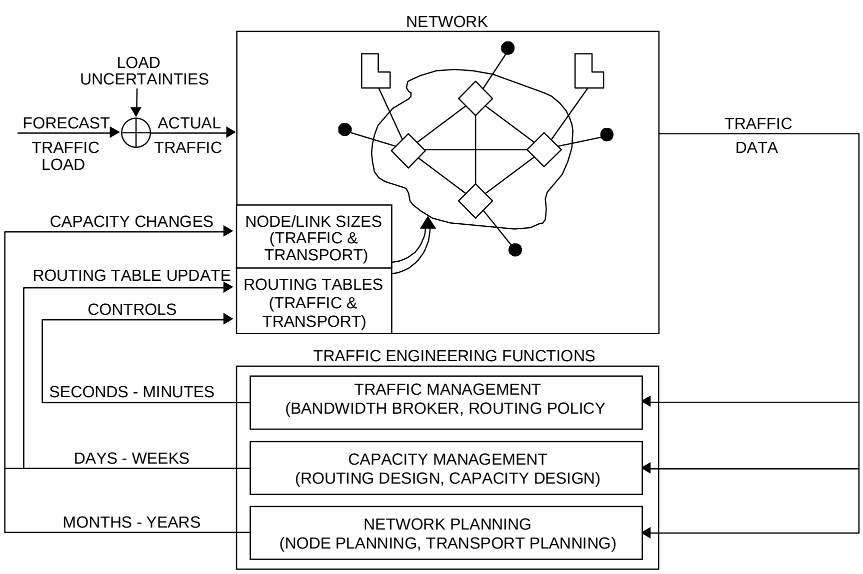

Figure 1.1 - Traffic Engineering Model

The functions depicted in Figure 1.1 are consistent with the definition of TE employed by the Traffic Engineering Working Group (TEWG) within the Internet Engineering Task Force (IETF):

Internet Traffic Engineering is concerned with the performance optimization of operational networks. It encompasses the measurement, modeling, characterization, and control of Internet traffic, and the application of techniques to achieve specific performance objectives, including the reliable and expeditious movement of traffic through the network, the efficient utilization of network resources, and the planning of network capacity.

This definition of TE methods is somewhat inconsistent with the traditional telecom usage of the term “traffic engineering”, which has more to do with network dimensioning and capacity planning. While these functions are encompassed by the above definition of TE, the scope of the analysis and recommendations here goes well beyond dimensioning and includes call and connection routing, QoS resource management, routing table management, dynamic transport routing, capacity management, and operational requirements.

Terminology used in the document, as illustrated in Figure 1.2, is that a link is a transmission medium (logical or physical) which connects two nodes, a path is a sequence of links connecting an origin and destination node, and a route is the set of different paths between the origin and destination that a call might be routed on within a particular routing discipline. Here a call is a generic term used to describe the establishment, utilization, and release of a connection, or data flow. In this context a call can refer to a voice call established perhaps using the SS7 signaling protocol, or to a web-based data flow session, established perhaps by the HTTP and associated IP-based protocols. Various implementations of routing tables are discussed in ANNEX 2.

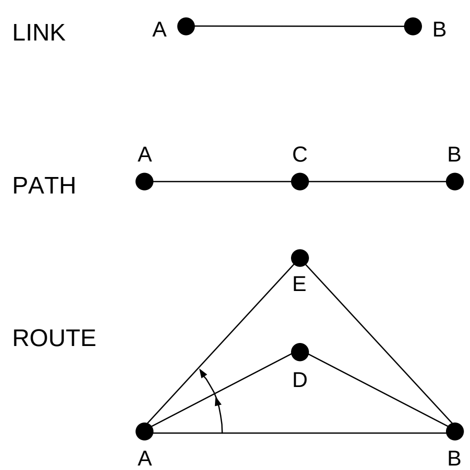

Figure 1.2 - Terminology

Traffic engineering functions include traffic management, capacity management, and network planning. Traffic management ensures that network performance is maximized under all conditions including load shifts and failures. Capacity management ensures that the network is designed and provisioned to meet performance objectives for network demands at minimum cost. Network planning ensures that node and transport capacity is planned and deployed in advance of forecasted traffic growth. Figure 1.1 illustrates traffic management, capacity management, and network planning as three interacting feedback loops around the network. The input driving the network ("system") is a noisy traffic load ("signal"), consisting of predictable average demand components added to unknown forecast error and load variation components. The load variation components have different time constants ranging from instantaneous variations, hour-to-hour variations, day-to-day variations, and week-to-week or seasonal variations. Accordingly, the time constants of the feedback controls are matched to the load variations, and function to regulate the service provided by the network through capacity and routing adjustments.

Traffic management functions include a) call routing, which entails number/name translation to routing address, b) connection or bearer-path routing methods, c) QoS resource management, d) routing table management, and e) dynamic transport routing. These functions can be a) decentralized and distributed to the network nodes, b) centralized and allocated to a centralized controller such as a bandwidth broker, or c) performed by a hybrid combination of these approaches.

Capacity management plans, schedules, and provisions needed capacity over a time horizon of several months to one year or more. Under exceptional circumstances, capacity can be added on a shorter-term basis, perhaps one to several weeks, to alleviate service problems. Network design embedded in capacity management encompasses both routing design and capacity design. Routing design takes account of the capacity provided by capacity management, and on a weekly or possibly real-time basis adjusts routing tables as necessary to correct service problems. The updated routing tables are provisioned (configured) in the switching systems either directly or via an automated routing update system. Network planning includes node planning and transport planning, operates over a multiyear forecast interval, and drives network capacity expansion over a multiyear period based on network forecasts.

The scope of the TE methods includes the establishment of connections for narrowband, wideband, and broadband multimedia services within multiservice networks and between multiservice networks. Here a multiservice network refers to one in which various classes of service share the transmission, switching, management, and other resources of the network. These classes of services can include constant bit rate (CBR), variable bit rate (VBR), unassigned bit rate (UBR), and available bit rate (ABR) traffic classes. There are quantitative performance requirements that the various classes of service normally are required to meet, such as end-to-end blocking, delay, and/or delay-jitter objectives. These objectives are achieved through a combination of traffic management and capacity management.

Figure 1.3 illustrates the functionality for setting up a connection from an originating node in one network to a destination node in another network, using one or more routing methods across networks of various types. The Figure illustrates a multimedia connection between two PCs which carries traffic for a combination of voice, video, and image applications. For this purpose a logical point-to-point connection is established from the PC served by node a1 to the PC served by node c2. The connection could be a CBR ISDN connection across TDM-based network A and ATM-based network C, or it might be a VBR connection via IP-based network B. Gateway nodes a3, b1, b4, and c1 provide the interworking capabilities between the TDM-, ATM-, and IP-based networks. The actual multimedia connection might be routed, for example, on a path consisting of nodes a1-a2-a3-b1-b4-c1-c2, or possibly on a different path through different gateway nodes.

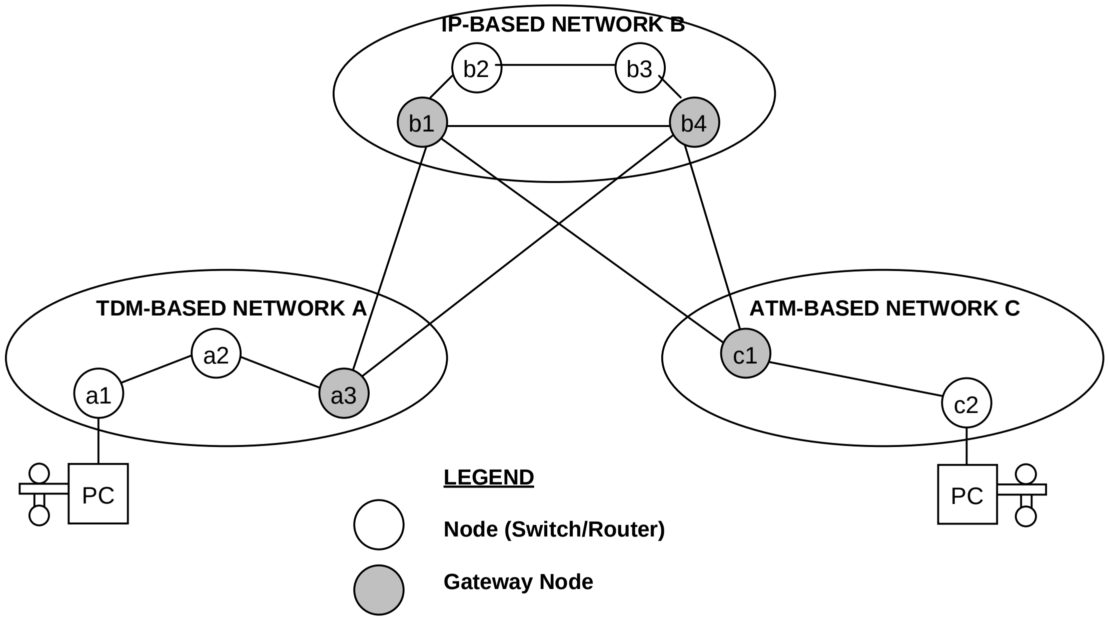

Figure 1.3 Example of Multimedia Connection across TDM-, ATM-, and IP-Based Networks

We now briefly describe the traffic model, the traffic management functions, the capacity management functions, and the TE operational requirements, which are further developed in ANNEXES 2-7 of the document.

4.0 Traffic Models

In this section we discuss load variation models which drive traffic engineering functions, that is traffic management, capacity management, and network planning. Table 1.1 summarizes examples of models that could be used to represent the different traffic variations under consideration. Traffic models for both voice and data traffic need to be reflected.

Work has been done on measurement and characterization of data traffic, such as web-based traffic [FGLRRT00, FGHW99, LTWW94]. Some of the analysis suggests that web-based traffic can be self-similar, or fractal, with very large variability and extremely long tails of the associated traffic distributions. Characterization studies of such data traffic have investigated various traditional models, such as the Markov modulated Poisson Process (MMPP), in which it is shown that MMPP with two parameters can suitably capture the essential nature of the data traffic [H99, BCHLL99].

Modeling work has been done to investigate the causes of the extreme variability of web-based traffic. In [HM00], the congestion-control mechanisms for web-based traffic, such as window flow control for transport-control-protocol (TCP) traffic appear to be at the root cause of its extreme variability over small time scales. [FGHW99] also shows that the variability over small time scales is impacted in a major way by the presence of TCP-like flow control algorithms which give rise to burstiness and clustering of IP packets. However, [FGHW99] also finds that the self-similar behavior over long time scales is almost exclusively due to user-related variability and not dependent on the underlying network-specific aspects.

Regarding the modeling of voice and date traffic in a multiservice model, [HM00] suggests that the regular flow control dynamics are more useful to model than the self-similar traffic itself. Much of the traffic to be modeled is VBR traffic subject to service level agreements (SLAs), which is subject to admission control based on equivalent bandwidth resource requirements and also to traffic shaping in which out-of-contract packets are marked for dropping in the network queues if congestion arises. Other VBR traffic, such as best-effort internet traffic, is not allocated any bandwidth in the admission of session flows, and all of its packets would be subject to dropping ahead of the CBR and VBR-SLA traffic. Hence, we can think of the traffic model consisting of two components:

- the CBR and VBR-SLA traffic that is not marked for dropping constitute less variable traffic subject to more traditional models
- the VBR best-effort traffic and the VBR-SLA traffic packets that are marked and subject to dropping constitute a much more variable, self-similar traffic component.

Considerable work has been done on modeling of broadband and other data traffic, in which two-parameter models that capture the mean and burstiness of the connection and flow arrival processes have proven to be quite adequate. See [E.716] for a good reference on this. Much work has also been done on measurement and characterization of voice traffic, and two-parameter models reflecting mean and variance (the ratio of the variance to the mean is sometimes called the peakedness parameter) of traffic have proven to be accurate models. We model the large variability in packet arrival processes in an attempt to capture the extreme variability of the traffic.

Here we reflect the two-parameter, multiservice traffic models for connection and flow arrival processes, which are manageable from a modeling and analysis aspect and which attempt to capture essential aspects of data and voice traffic variability for purposes of traffic engineering and QoS methods. In ANNEX 2 we introduce the models of variability in the packet arrival processes.

Table 1.1
Traffic Models for Load Variations of Connection/Flow Arrival Processes

Traffic Variations Time Constant	Load Variation Examples For Traffic Management	Illustrative Traffic Model for Capacity Management	Capacity Impacts
minute to minute	real-time random traffic fluctuations; bursty overflow traffic; Focused overloads (e.g., caused by radio/TV call-ins, natural disasters, etc.); general overloads (e.g., caused by peak-day calling); traffic congestion caused by network failure (e.g., fiber cut or node failure); traffic shifts due to price variations for transit traffic, arbitrage, and bulk re-sale.	stochastic model; normally with 2 parameters (mean and variance); focused and general overload traffic excluded; network failure traffic excluded	busy-hour traffic load capacity (excludes focused overload, general overload, and network failure traffic)
hour to hour	business traffic day peak; web-based (consumer) traffic evening peak; mobile traffic (consumer) weekend/evening peak	deterministic model; 20-day average time-varying mean; multihour design	multihour capacity
day to day	Monday morning busiest for business day traffic	stochastic model; normally with 2	day-to-day capacity

	compared to average morning; Sunday evening busiest for web-based traffic compared to average evening; Friday evening busiest for mobile traffic compared to average evening;	parameters (mean and variance); several levels of variance modeled for low/med./high day-to-day variations	
week to week	winter/summer seasonal variations; forecast errors	stochastic model; normally with 2 parameters (mean and variance); maximum flow routing & capacity design	reserve capacity

For instantaneous traffic load variations, the load is typically modeled as a stationary random process over a given period (normally within each hourly period) characterized by a fixed mean and variance. From hour to hour, the mean traffic loads are modeled as changing deterministically; for example, according to their 20-day average values. From day to day, for a fixed hour, the mean load can be modeled, for example, as a random variable having a gamma distribution with a mean equal to the 20-day average load. From week to week, the load variation is modeled as a random process in the network design procedure. The random component of the realized week-to-week load is the forecast error, which is equal to the forecast load minus the realized load. Forecast error is accounted for in short-term capacity management.

In traffic management, traffic load variations such as instantaneous variations, hour-to-hour variations, day-to-day traffic variations, and week-to-week variations are responded to in traffic management by appropriately controlling number translation/routing, path selection, routing table management, and/or QoS resource management. Traffic management provides monitoring of network performance through collection and display of traffic and performance data, and allows traffic management controls, such as destination-address per-connection blocking, per-connection gapping, routing table modification, and path selection/reroute controls, to be inserted when circumstances warrant. For example, a focused overload might lead to application of connection gapping controls in which a connection request to a particular destination address or set of addresses is admitted only once every x seconds, and connections arriving after an accepted call are rejected for the next x seconds. In that way call gapping throttles the calls and prevents overloading the network to a particular focal point. Routing table modification and reroute control are illustrated in ANNEXES 2, 3, 5, and 7.

Capacity management must provide sufficient capacity to carry the expected traffic variations so as to meet end-to-end blocking/delay objective levels. Here the term blocking refers to the denial or non-admission of a call or connection request, based for example on the lack of available resources on a particular link (e.g., link bandwidth or queuing resources). Traffic load variations lead in direct measure to capacity increments and can be categorized as (1) minute-to-minute instantaneous variations and associated busy-hour traffic load capacity, (2) hour-to-hour variations and associated multihour capacity, (3) day-to-day variations and associated day-to-day capacity, and (4) week-to-week variations and associated reserve capacity.

Design methods within the capacity management procedure account for the mean and variance of the within-the-hour variations of the offered and overflow loads. For example, classical methods [e.g., Wil56] are used to size links for these two parameters of load. Multihour dynamic route design accounts for the hour-to-hour variations of the load and, hour-to-hour capacity can vary from zero to 20 percent or more of network capacity. Hour-to-hour capacity can be reduced by multihour dynamic routing design models such as the discrete event flow optimization, traffic load flow optimization, and virtual trunking flow optimization models described in ANNEX 6. As noted in Table 1.1, capacity management excludes non-recurring traffic such as caused by overloads (focused or general overloads), or failures. This process is described further in ANNEX 7.

It is known that some daily variations are systematic (for example, Monday morning business traffic is usually higher than Friday morning); however, in some day-to-day variation models these systematic changes are ignored and lumped into the stochastic model. For instance, the traffic load between Los Angeles and New Brunswick is very similar from one day to the next, but the exact calling levels differ for any given day. This load variation can be characterized in network design by a stochastic model for the daily variation, which results in additional capacity called day-to-day capacity. Day-to-day capacity is needed to meet the average blocking/delay objective when the load varies according to the stochastic model. Day-to-day capacity is nonzero due to the nonlinearities in link blocking and/or link queuing delay levels as a function of load. When the load on a link fluctuates about a mean value, because of day-to-day variation, the mean blocking/delay is higher than the blocking/delay produced by the mean load. Therefore, additional capacity is provided to maintain the blocking/delay probability grade-of-service objective in the presence of day-to-day load variation.

Typical day-to-day capacity required is 4--7 percent of the network cost for medium to high day-to-day variations, respectively. Reserve capacity, like day-to-day capacity, comes about because load uncertainties---in this case forecast errors---tend to cause capacity buildup in excess of the network design that exactly matches the forecast loads. Reluctance to disconnect and rearrange link and transport capacity contributes to this reserve capacity buildup. At a minimum, the currently measured mean load is used to adjust routing and capacity design, as needed. In addition, the forecast-error variance component in used in some models to build in so-called protective capacity. Reserve or protective capacity can provide a cushion against overloads and failures, and generally benefits network performance. However, provision for reserve capacity is not usually built into the capacity management design process, but arises because of sound administrative procedures. These procedures attempt to minimize total cost, including both network capital costs and operations costs. Studies have shown that reserve capacity in some networks to be in the range of 15 to 25 percent or more of network cost [FHH79]. This is further described in ANNEXES 5 and 6.

5.0 Traffic Management Functions

In ANNEXES 2-5, traffic management functions are discussed:

a) Call Routing Methods (ANNEX 2). Call routing involves the translation of a number or name to a routing address. We describe how number (or name) translation should result in the E.164 ATM end-system addresses (AESA), network routing addresses (NRAs), and/or IP addresses. These addresses are used for routing purposes and therefore must be carried in the connection-setup information element (IE).

b) Connection/Bearer-Path Routing Methods (ANNEX 2). Connection or bearer-path routing involves the selection of a path from the originating node to the destination node in a network. We discuss bearer-path selection methods, which are categorized into the following four types: fixed routing (FR), time-dependent routing (TDR), state-dependent routing (SDR), and event-dependent routing (EDR). These methods are associated with routing tables, which consist of a route and rules to select one path from the route for a given connection or bandwidth-allocation request.

c) QoS Resource Management Methods (ANNEX 3). QoS resource management functions include class-of-service derivation, policy-based routing table derivation, connection admission, bandwidth allocation, bandwidth protection, bandwidth reservation, priority routing, priority queuing, and other related resource management functions.

d) Routing Table Management Methods (ANNEX 4). Routing table management information, such as topology update, status information, or routing recommendations, is used for purposes of applying the routing table design rules for determining path choices in the routing table. This information is exchanged between one node and another node, such as between the ON and DN, for example, or between a node and a network element such as a bandwidth-broker processor (BBP). This information is used to generate the routing table, and then the routing table is used to determine the path choices used in the selection of a path.

e) Dynamic Transport Routing Methods (ANNEX 5). Dynamic transport routing combines with dynamic traffic routing to shift transport bandwidth among node pairs and services through use of flexible transport switching technology, such as optical cross-connects (OXCs). Dynamic transport routing offers advantages of simplicity of design and robustness to load variations and network failures, and can provide automatic link provisioning, diverse link routing, and rapid link restoration for improved transport capacity utilization and performance under stress. OXCs can reconfigure logical transport capacity on demand, such as for peak day traffic, weekly redesign of link capacity, or emergency restoration of capacity under node or transport failure. MPLS control capabilities are proposed for the setup of layer 2 logical links through OXCs [ARDC99].

6.0 Capacity Management Functions

In ANNEX 6, we discuss capacity management methods, as follows:

a) Link Capacity Design Models. These models find the optimum tradeoff between traffic carried on a shortest network path (perhaps a direct link) versus traffic carried on alternate (longer, less efficient) network paths.

b) Shortest Path Selection Models. These models enable the determination of shortest paths in order to provide a more efficient and flexible routing plan.

c) Multihour Network Design Models. Three models are described including i) discrete event flow optimization (DEFO) models, ii) traffic load flow optimization (TLFO) models, and iii) virtual trunking flow optimization (VTFO) models. DEFO models have the advantage of being able to model traffic and routing methods of arbitrary complexity, for example, such as self-similar traffic.

d) Day-to-day Load Variation Design Models. These models describe techniques for handling day-to-day variations in capacity design.

e) Forecast Uncertainty/Reserve Capacity Design Models. These models describe the means for accounting for errors in projecting design traffic loads in the capacity design of the network.

7.0 Traffic Engineering Operational Requirements

In ANNEX 7, we discuss traffic engineering operational requirements, as follows:

a) Traffic Management. We discuss requirements for real-time performance monitoring, network control, and work center functions. The latter includes automatic controls, manual controls, code controls, cancel controls, reroute controls, peak-day controls, traffic management on peak days, and interfaces to other work centers.

b) Capacity Management – Forecasting. We discuss requirements for load forecasting, including configuration database functions, load aggregation, basing, and projection functions, and load adjustment cycle and view of business adjustment cycle. We also discuss network design, work center functions, and interfaces to other work centers.

c) Capacity Management – Daily and Weekly Performance Monitoring. We discuss requirements for daily congestion analysis, study-week congestion analysis, and study-period congestion analysis.

d) Capacity Management – Short-Term Network Adjustment. We discuss requirements for network design, work center functions, and interfaces to other work centers.

e) Comparison of off-line (TDR) versus on-line (SDR/EDR) TE methods. We contrast off-line TE methods, such as in a TDR-based network, with on-line TE methods, such as in an SDR- or EDR-based network.

8.0 Traffic Engineering Modeling & Analysis

In ANNEXES 2-6 we use network models to illustrate the traffic engineering methods developed in the document. The details of the models are presented in each ANNEX in accordance with the TE functions being illustrated.

In the document, a full-scale 135-node national network node model is used together with a multiservice traffic demand model to study various TE scenarios and tradeoffs. Typical voice/ISDN traffic loads are used to model the various network alternatives. These voice/ISDN loads are further segmented in the model into eight constant-bit-rate (CBR) virtual networks (VNETs), including business voice, consumer voice, international voice in and out, key-service voice, normal and key-service 64-kbps ISDN data, and 384-kbps ISDN data. The data services traffic model incorporates typical traffic load patterns and comprises three additional VNET load patterns. These include a) a variable bit rate real-time (VBR-RT) VNET, representing services such as IP-telephony and compressed voice, b) a variable bit rate non-real-time (VBR-NRT) VNET, representing services such as WWW multimedia and credit card check, and c) an unassigned bit rate (UBR) VNET, representing best-effort services such as email, voice mail, and file transfer multimedia applications. The cost model represents typical switching and transport costs, and illustrates the economies-of-scale for costs projected for high capacity network elements in the future.

Many different alternatives and tradeoffs are examined in the models, including:

1. centralized routing table control versus distributed control
2. off-line, pre-planned (e.g.,TDR-based) routing table control versus on-line routing table control (e.g., SDR- or EDR-based)
3. per-flow traffic management versus per-virtual-network (or per-traffic-trunk or per-bandwidth-pipe) traffic management
4. sparse logical topology versus meshed logical topology
5. FR versus TDR versus SDR versus EDR path selection
6. multilink path selection versus two-link path selection
7. path selection using local status information versus global status information
8. global status dissemination alternatives including status flooding, distributed query for status, and centralized status in a bandwidth-broker processor

Table 1.2 summarizes brief comparisons and observations, based on the modeling, in each of the above alternatives and tradeoffs (further details are contained in ANNEXES 2-6).

Table 1.2
Tradeoff Categories and Comparisons
(Based on Modeling in ANNEXES 2-6)

Tradeoff Category	Traffic Management Performance Comparisons	Routing Table Management Comparisons	Capacity Management Comparisons
TE methods applied vs. no TE methods applied	TE methods considerably improve performance	control load comparable	comparable design efficiency
centralized vs. distributed routing table control	Distributed control performance somewhat better (more up-to-date status information)	control load comparable on per-node basis	comparable design efficiency
off-line/ pre-planned (TDR) vs. on-line (SDR, EDR) routing table control	On-line control somewhat better performance	TDR and EDR control load less than SDR	SDR & EDR comparable design efficiency; both better than TDR
FR vs. TDR vs. SDR vs. EDR path selection	EDR/SDR performance better than TDR better than FR	FR/TDR/EDR have lower control load than SDR	EDR/SDR design efficiency better than TDR better than FR
multilink path selection vs. two-link path selection	Multilink path selection better under overload; Two-link path selection better under failure; Two-link path selection lower call set-up delay	multilink path selection control load generally less than two-link path selection	multilink design efficiency better than two-link
sparse logical topology vs. meshed logical topology	Sparse topology better under overload; Meshed topology better under failure	sparse topology control load generally less than meshed topology	sparse topology design efficiency somewhat better than meshed
single-area flat topology vs. multi-area	Single-area performance better than multi-area	SDR case: multi-area control load less than	single-area topology design efficiency

hierarchical topology		single-are control load EDR case: total control load comparable	somewhat better than multi-area
local status information vs. global status information	Local status performance somewhat better than global (more up-to-date information)	local status control load less than global status control load	comparable design efficiency
status dissemination: status flooding vs. distributed query-for-status vs. centralized status in BBP	Distributed query-for-status somewhat better than status flooding & centralized status (more up-to-date information)	centralized BBP and distributed query-for-status comparable on per-node basis; status flooding considerably higher control load	comparable design efficiency
per-flow traffic management vs. per-virtual-network (per-traffic-trunk) traffic management	Comparable performance	per-virtual-network control load less than per-flow control load	per-flow design efficiency somewhat better than per-virtual-network
integrated voice & data network vs. separate voice & data networks	Integrated network performance better than separate network performance	Total control load comparable	integrated network design efficiency better than separate network

9.0 Conclusions/Recommendations

Following is a summary of the main conclusions/recommendations reached in the document.

9.1 Conclusions/Recommendations on Call Routing & Connections Routing Methods (ANNEX 2)

- TE methods are recommended to be applied, and in all cases of the TE methods being applied, network performance is always better and usually substantially better than when no TE methods are applied

- Sparse-topology multilink-routing networks are recommended and provide better overall performance under overload than meshed-topology networks, but performance under failure may favor the 2-link STT-EDR/DC-SDR meshed-topology options with more alternate routing choices.

- Single-area flat topologies are recommended and exhibit better network performance and, as discussed and modeled in ANNEX 6, greater design efficiencies in comparison with multi-area hierarchical topologies. As illustrated in ANNEX 4, larger administrative areas can be achieved through use of EDR-based TE methods as compared to SDR-based TE methods.

- Event-dependent-routing (EDR) TE path selection methods are recommended and exhibit comparable or better network performance compared to state-dependent-routing (SDR) methods.

  a. EDR TE methods are shown to an important class of TE algorithms. EDR TE methods are distinct from the TDR and SDR TE methods in how the paths (e.g., MPLS label switched paths, or LSPs) are selected. In the SDR TE case, the available link bandwidth (based on LSA flooding of ALB information) is typically used to compute the path. In the EDR TE case, the ALB information is not needed to compute the path, therefore the ALB flooding does not need to take place (reducing the overhead).

  b. EDR TE algorithms are adaptive and distributed in nature and typically use learning models to find good paths for TE in a network. For example, in a success-to-the-top (STT) EDR TE method, if the LSR-A to LSR-B bandwidth needs to be modified, say increased by delta-BW, the primary LSP-p is tried first. If delta-BW is not available on one or more links of LSP-p, then the currently successful LSP-s is tried next. If delta-BW is not available on one or more links of LSP-s, then a new LSP is searched by trying additional candidate paths until a new successful LSP-n is found or the candidate paths are exhausted. LSP-n is then marked as the currently successful path for the next time bandwidth needs to be modified. The performance of distributed EDR TE methods is shown to be equal to or better than SDR methods, centralized or distributed.

c. While SDR TE models typically use available-link-bandwidth (ALB) flooding for TE path selection, EDR TE methods do not require ALB flooding. Rather, EDR TE methods typically search out capacity by learning models, as in the STT method above. ALB flooding can be very resource intensive, since it requires link bandwidth to carry LSAs, processor capacity to process LSAs, and the overhead can limit area/autonomous system (AS) size. Modeling results show EDR TE methods can lead to a large reduction in ALB flooding overhead without loss of network throughput performance [as shown in ANNEX 4].

d. State information as used by the SDR options (such as with link-state flooding) provides essentially equivalent performance to the EDR options, which typically used distributed routing with crankback and no flooding.

e. Various path selection methods can interwork with each other in the same network, as required for multi-vendor network operation.

- Interdomain routing methods are recommended which extend the intradomain call routing and connection routing concepts, such as flexible path selection and per-class-of-service bandwidth selection, to routing between network domains.

9.2 Conclusions/Recommendations on QoS Resource Management Methods (ANNEX 3)

- QoS resource management is recommended and is shown to be effective in achieving connection-level and packet-level GoS objectives, as well as key service, normal service, and best effort service differentiation.

- Admission control is recommended and is the basis that allows for applying most of the other controls described in this document.

- Per-VNET bandwidth allocation is recommended and is essentially equivalent to per-flow bandwidth allocation in network performance and efficiency. Because of the much lower routing table management overhead requirements, as discussed and modeled in ANNEX 4, per-VNET bandwidth allocation is preferred to per-flow allocation.

- Both MPLS QoS and bandwidth management and DiffServ priority queuing management are recommended and are important for ensuring that multiservice network performance objectives are met under a range of network conditions. Both mechanisms operate together to ensure QoS resource allocation mechanisms (bandwidth allocation, protection, and priority queuing) are achieved.

9.3 Conclusions/Recommendations on Routing Table Management Methods & Requirements (ANNEX 4)

- Per-VNET bandwidth allocation is recommended and is preferred to per-flow allocation because of the much lower routing table management overhead requirements. Per-VNET bandwidth allocation is essentially equivalent to per-flow bandwidth allocation in network performance and efficiency, as discussed in ANNEX 3.

- EDR TE methods are recommended and can lead to a large reduction in ALB flooding overhead without loss of network throughput performance. While SDR TE methods typically use ALB flooding for TE path selection, EDR TE methods do not require ALB flooding. Rather, EDR TE methods typically search out capacity by learning models, as in the STT method. ALB flooding can be very resource intensive, since it requires link bandwidth to carry LSAs, processor capacity to process LSAs, and the overhead can limit area/autonomous system (AS) size.

- EDR TE methods are recommended and lead to possible larger administrative areas as compared to SDR-based TE methods because of lower routing table management overhead requirements. This can help achieve single-area flat topologies which, as discussed in ANNEX 3, exhibit better network performance and, as discussed in ANNEX 6, greater design efficiencies in comparison with multi-area hierarchical topologies.

9.4 Conclusions/Recommendations on Transport Routing Methods (ANNEX 5)

- Dynamic transport routing is recommended and provides greater network throughput and, consequently, enhanced revenue, and at the same time capital savings should result, as discussed in ANNEX 6.

a. Dynamic transport routing network design enhances network performance under failure, which arises from automatic inter-backbone-router and access logical-link diversity in combination with the dynamic traffic routing and transport restoration of logical links.

b. Dynamic transport routing network design is recommended and improves network performance in comparison with fixed transport routing for all network conditions simulated, which include abnormal and unpredictable traffic load patterns.

- Traffic and transport restoration level design is recommended and allows for link diversity to ensure a minimum level of performance under failure.

- Robust routing techniques are recommended, which include dynamic traffic routing, multiple ingress/egress routing, and logical link diversity routing; these methods improve response to node or transport failures.

9.5 Conclusions/Recommendations on Capacity Management Methods (ANNEX 6)

- Discrete event flow optimization (DEFO) design models are recommended and are shown to be able to capture very complex routing behavior through the equivalent of a simulation model provided in software in the routing design module. By this means, very complex routing networks have been designed by the model, which include all of the routing methods discussed in ANNEX 2 (FR, TDR, SDR, and EDR methods) and the multiservice QoS resource allocation models discussed in ANNEX 3.

- Sparse topology options are recommended, such as the multilink STT-EDR/DC-SDR/DP-SDR options, which lead to capital cost advantages, and more importantly to operation simplicity and cost reduction. Capital cost savings are subject to the particular switching and transport cost assumptions. Operational issues are further detailed in ANNEX 7.

- Voice and data integration is recommended and

a. can provide capital cost advantages, and

b. more importantly can achieve operational simplicity and cost reduction, and

c. if IP-telephony takes hold and a significant portion of voice calls use voice compression technology, this could lead to more efficient networks.

- Multilink routing methods are recommended and exhibit greater design efficiencies in comparison with 2-link routing methods. As discussed and modeled in ANNEX 3, multilink topologies exhibit better network performance under overloads in comparison with 2-link routing topologies; however the 2-link topologies do better under failure scenarios.

- Single-area flat topologies are recommended and exhibit greater design efficiencies in termination and transport capacity, but higher cost, and, as discussed and modeled in ANNEX 3, better network performance in comparison with multi-area hierarchical topologies. As illustrated in ANNEX 4, larger administrative areas can be achieved through use of EDR-based TE methods as compared to SDR-based TE methods.

- EDR methods are recommended and exhibit comparable design efficiencies to SDR. This suggests that there is not a significant advantage for employing link-state information in these network designs, especially given the high overhead in flooding link-state information in SDR methods.

- Dynamic transport routing is recommended and achieves capital savings by concentrating capacity on fewer, high-capacity physical fiber links and, as discussed in ANNEX 5, achieves higher network throughput and enhanced revenue by their ability to flexibly allocate bandwidth on the logical links serving the access and inter-node traffic.

9.6 Conclusions/Recommendations on TE Operational Requirements (ANNEX 7)

- Monitoring of traffic and performance data is recommended and is required for traffic management, capacity forecasting, daily and weekly performance monitoring, and short-term network adjustment.

- Traffic management is recommended and is required to provide monitoring of network performance through collection and display of real-time traffic and performance data and allow traffic management controls such as code blocks, connection request gapping, and reroute controls to be inserted when circumstances warrant.

- Capacity management is recommended and is required for capacity forecasting, daily and weekly performance monitoring, and short-term network adjustment.

- Forecasting is recommended and is required to operate over a multiyear forecast interval and drive network capacity expansion.

- Daily and weekly performance monitoring is recommended and is required to identify any service problems in the network. If service problems are detected, short-term network adjustment can include routing table updates and, if necessary, short-term capacity additions to alleviate service problems. Updated routing tables are sent to the switching systems either directly or via an automated routing update system.

- Short-term capacity additions are recommended and are required as needed, but only as an exception, whereas most capacity changes are normally forecasted, planned, scheduled, and managed over a period of months or a year or more.

- Network design, which includes routing design and capacity design, is recommended and is required within the capacity management function.

- Network planning is recommended and is required for longer-term node planning and transport network planning, and operates over a horizon of months to years to plan and implement new node and transport capacity.

10. Recommended TE/QoS Methods for Multiservice Networks

In summary, TE methods are recommended in this Section for consideration in network evolution. These recommendations are based on

- results of analysis models presented in ANNEXES 2-6, which illustrate the tradeoffs between various TE approaches,

- results of operational comparison studies presented in ANNEXES 2-6,

- established best current practices and experience.

10.1 Recommended Application-Layer IP-Network-Based Service-Creation Capabilities

As discussed in ANNEX 4, these capabilities are recommended for application-layer service-creation capabilities:

- Parlay API (application programming interface)
- call processing language (CPL) & common gateway interface (CGI)
- SIP/IN (intelligent network) interworking

10.2 Recommended Call/IP-Flow Control Layer Capabilities

As discussed in ANNEXES 2 and 4, these capabilities are recommended for name translation, call signaling, and split gateway control:

- ENUM/DNS-based name to IP-address translation
- SIP-based distributed call signaling (DCS)
- MGCP/MEGACO for split gateway control

10.3 Recommended Connection/Bearer Control Layer Capabilities

In this Section we summarize the findings in ANNEXES 2, 3, and 4 which give rise to a recommendation for a TE/QoS admission control method for connection/flow admission, which incorporates dynamic QoS routing connection/bearer layer control.

The analysis considered in ANNEXES 2, 3, and 4 investigates bandwidth allocation for the aggregated case ("per traffic-trunk" or per-VNET (virtual network)) versus the per-flow bandwidth allocation. The following recommendations are made on QoS resource management, topology, and connection layer control:

- virtual-network traffic allocation for multiservice network
- MPLS-based virtual-network based QoS resource management & dynamic bandwidth reservation methods
- DiffServ-based priority queuing
- per-virtual-network (per-traffic-trunk) bandwidth allocation for lower routing table management overhead
- sparse-topology multilink routing for better performance & design efficiency
- single-area flat topology (as much as possible, while retaining edge-core architecture) for better performance & design efficiency
- MPLS and DiffServ functionality to meet TE/QoS requirements
- success-to-the-top (STT) event-dependent-routing (EDR) TE path selection methods for better performance & lower overhead

These TE admission control and dynamic QoS routing methods will ensure stable/efficient performance of TE methods and help manage resources for and differentiate key service, normal service, & best effort service, and are now briefly summarized. Figure 1.4 illustrates the recommended QoS resource management methods. As illustrated in the Figure, in the multi-service, QoS resource management network, bandwidth is allocated to the individual VNETs (high-priority key services VNETs, normal-priority services VNETs, and best-effort low-priority services

VNETs). The Figure also illustrates the use of virtual-network traffic allocation for multiservice networks and the means to differentiate key service, normal service, & best effort service. High-priority and normal-priority traffic connections/flows are subject to admission control based on equivalent bandwidth allocation techniques. However, best-effort services are allocated no bandwidth, and all best-effort traffic is subject to dropping in the queuing/scheduling discipline under congestion conditions.

This allocated bandwidth is protected by bandwidth reservation methods, as needed, but otherwise shared. Each ON monitors VNET bandwidth use on each VNET CRLSP, and determines when VNET CRLSP bandwidth needs to be increased or decreased. Bandwidth changes in VNET bandwidth capacity are determined by ONs based on an overall aggregated bandwidth demand for VNET capacity (not on a per-connection demand basis). Based on the aggregated bandwidth demand, these ONs make periodic discrete changes in bandwidth allocation, that is, either increase or decrease bandwidth on the CRLSPs constituting the VNET bandwidth capacity. For example, if connection requests are made for VNET CRLSP bandwidth that exceeds the current CRLSP bandwidth allocation, the ON initiates a bandwidth modification request on the appropriate CRLSP(s). For example, this bandwidth modification request may entail increasing the current CRLSP bandwidth allocation by a discrete increment of bandwidth denoted here as delta-bandwidth (DBW). DBW is a large enough bandwidth change so that modification requests are made relatively infrequently. Also, the ON periodically monitors CRLSP bandwidth use, such as once each minute, and if bandwidth use falls below the current CRLSP allocation the ON initiates a bandwidth modification request to decrease the CRLSP bandwidth allocation by a unit of bandwidth such as DBW.

Therefore the recommendation is to do "per-VNET", or per traffic trunk, bandwidth allocation, and *not* call by call, or "per flow" allocation, as , as discussed in Sections 3.4 and 3.5. This kind of per-VNET bandwidth allocation also applies in the case of multi-area TE, as discussed in Sections 2.8 and 3.8. Therefore some telephony concepts, such as call-by-call set up, are not needed in VoIP/TE. That is, there are often good reasons not to make things look like the PSTN. On the other hand, some principles do still apply to VoIP/TE, but are not used as yet, and should be.

The main point about bandwidth reservation is related to both admission control and queue management. That is, if a flow is to be admitted on a longer path, that is, not the primary path (which is preferred and tried first, but let us assume did not have the available bandwidth on one or more links/queues), then there needs to be a minimum level of available bandwidth, call in RESBW (reserved bandwidth), available on each link and in each queue in *addition to* the requested bandwidth (REQBW). That is, one needs to have RESBW + BEWBW available on each link and queue before admitting the flow on the longer path. On the primary path RESBW is *not* required. The simulation results given in ANNEX 3 are for an MPLS network, and the results show the effect of using bandwidth reservation, and what happens if you do not use bandwidth reservation (see Tables 3.4 and 3.5). Bandwidth allocation and management is done according to the traffic priority (i.e., key, normal, and best effort), as described in ANNEX 3, and is an additional use of bandwidth reservation methods beyond the use in path selection, as in the example above. Bandwidth allocation in the queues is done according to traffic priority, as discussed in Section 3.6. These principles put forth in the document do not depend on whether the underlying technology is IP/MPLS-based, ATM/PNNI-based, or TDM/E.351-based, they apply to all technologies, as is demonstrated by the models.

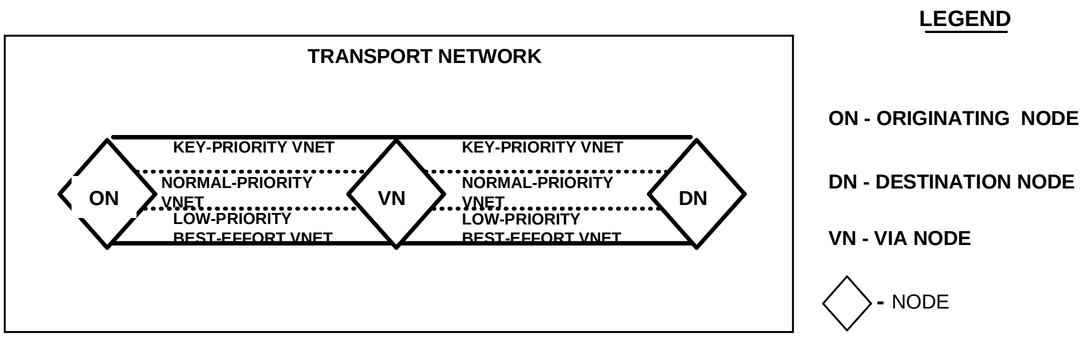

- distributed method applied on a per-virtual-network basis
- ON allocates bandwidth to each virtual-network (VNET) based on demand
- for VNET bandwidth increase
  - ON decides link-bandwidth-modification threshold (Pi) based on
    - bandwidth-in-progress (BWIP)
    - routing priority (key, normal, best-effort)
    - bandwidth allocation BWavg
    - first/alternate choice path
  - ON launches a CRLDP label request message with explicit route, modify-flag, traffic parameters, & threshold Pi (carried in setup priority)
- VNs keep local link state of idle link bandwidth (ILBW), including lightly loaded (LL), heavily loaded (HL), reserved (R), & busy (B)
- VNs compare link state to Pi threshold
- VNs send crankback/bandwidth-not-available notification message to ILSR if Pi threshold not met

Figure 1.4. Use MPLS/DiffServ/Virtual-Network-based QoS Resource Management with Dynamic Bandwidth Reservation & Priority Queuing Methods

In the models the per-VNET method compares favorably with the per-flow method, which is all feasible within the current MPLS protocol specification and is therefore recommended for the TE admission control and dynamic QoS routing methods.

Furthermore, we find that a distributed event-dependent-routing (EDR)/STT method of LSP management works just as well or better than the state-dependent-routing (SDR) with flooding. An example of the EDR/STT method:

Figure 1.5 illustrates the recommended STT EDR path selection method and the use of a sparse, single-area topology.

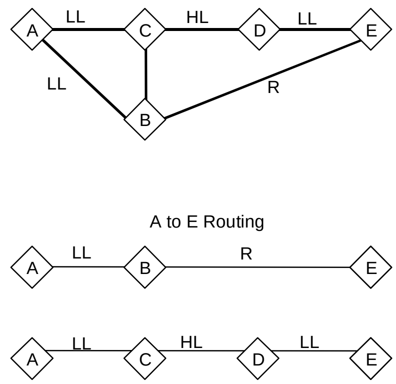

Example EDR TE method (success-to-the-top (STT) EDR):

1. if LSR-A to LSR-E bandwidth needs to be modified (say increased by delta-BW) primary LSP-p (e.g., LSP A-B-E) is tried first
2. available bandwidth tested locally on each link in LSP-p, if bandwidth not available (e.g., setup priority is heavily-loaded HL state and link BE is in reserved R state), crankback to LSR-A
3. if delta-BW is not available on one or more links of LSP-p, then the currently successful LSP-s (e.g., LSP A-C-D-E) is tried next
4. if delta-BW is not available on one or more links of LSP-s, then a new LSP is searched by trying additional candidate paths until a new successful LSP-n is found or the candidate paths are exhausted
5. LSP-n is then marked as the currently successful path for the next time bandwidth needs to be modified

Figure 1.5. Use Success-to-the-Top (STT) Event-Dependent-Routing (EDR) TE Path Selection Methods in a Sparse, Single-Area Topology

The EDR/STT method is fully distributed, reduces flooding, and a larger perhaps even a single backbone area could be used as a result. Edge-router (ER) to backbone-router (BR) hierarchy is also modeled. We modeled an MPLS/DiffServ ER-BR resource management, although it is sometimes claimed that DiffServ alone would suffice on the ER-BR links. The problem there is what happens when bandwidth is exhausted for the connection-oriented voice, ISDN, IP-telephony, etc. services versus the best-effort services. One needs a TE admission control mechanism to reject connection requests when need be. In the ER/BR hierarchy modeled, there is a mesh of LSPs in the backbone, but separate LSPs ("big pipes") for each ER to the backbone BRs, that is, for each ER-BR area (i.e., there is no ER-ER LSP mesh in this case).

Some example VNET definitions are given in Figure 1.6 along with example Service Identity components, as well traffic allocation characteristics such as service priority and bandwidth characteristics.

10.4 Recommended Transport Routing Capabilities

As discussed in ANNEX 5, the following recommendations are made for transport routing:

- dynamic transport routing for better performance & design efficiency
- traffic and transport restoration level design, which allows for link diversity to ensure a minimum level of performance under failure

Virtual Network Name	Service Identity Examples	Virtual Network Traffic Priority & Traffic Characteristics
1. BUSINESS VOICE	VPN, DIRECT CONNECT 800, 800 SERVICE, 900 SERVICE	NORMAL PRIORITY; 64 KBPS CBR
2. CONSUMER VOICE	LONG DISTANCE SERVICE (LDS)	NORMAL PRIORITY; 64 KBPS CBR
3. INTL VOICE OUTBOUND	INTL LDS OUTBOUND, INTL 800 OUTBOUND, GLOBAL VPN OUTBOUND, INTL TRANSIT	NORMAL PRIORITY; 64 KBPS CBR
4. INTL VOICE INBOUND	INTL LDS INBOUND, INTL 800 INBOUND, GLOBAL VPN INBOUND, INTL TRANSIT INBOUND	KEY PRIORITY; 64 KBPS CBR
5. 800-GOLD	DIRECT CONNECT 800 GOLD, 800 GOLD, VPN-KEY	KEY PRIORITY; 64 KBPS CBR
6. 64 KBPS ISDN	64 KBPS SDS, 64 KBPS SWITCHED DIGITAL INTL (SDI)	NORMAL PRIORITY; 64 KBPS CBR
7. 64 KBPS ISDN	64 KBPS SDS & SDI (KEY)	KEY PRIORITY; 64 KBPS CBR
8. 384 KBPS ISDN	384 KBPS SDS, 384 KBPS SDI	NORMAL PRIORITY; 384 KBPS CBR
9. IP TELEPHONY	IP TELEPHONY, COMPRESSED VOICE	NORMAL PRIORITY; VARIABLE RATE, INTERACTIVE & DELAY SENSITIVE; VBR-RT: 10% OF VN1+VN2+VN3+VN4+VN5 TRAFFIC LOAD, CALL DATA RATE VARIES FROM 6.4 KBPS TO 51.2 KBPS (25.6 KBPS MEAN)
10. IP MULTIMEDIA	IP MULTIMEDIA, WWW, CREDIT CARD CHECK	NORMAL PRIORITY; VARIABLE RATE, NON-INTERACTIVE & NOT DELAY SENSITIVE; VBR-NRT: 30% OF VN2 TRAFFIC LOAD, CALL DATA RATE VARIES FROM 38.4 KBPS TO 64 KBPS (51.2 KBPS MEAN)
11. UBR BEST EFFORT	VOICE MAIL, EMAIL, FILE TRANSFER	BEST-EFFORT PRIORITY; VARIABLE RATE, NON-INTERACTIVE & NOT DELAY SENSITIVE; UBR: 30% OF VN1 TRAFFIC LOAD, CALL DATA RATE VARIES FROM 6.4 KBPS TO 3072 KBPS (1536 KBPS MEAN)

Figure 1.6. Use Virtual-Network Traffic Allocation for Multiservice Network Differentiate Key Service, Normal Service, & Best Effort Service 1

10.5 Recommended Network Operations Capabilities

As discussed in ANNEXES 5 and 6, the following recommendations are made for network operations and design:

- monitor traffic & performance data for traffic management & capacity management

Figure 1.1 illustrates the monitoring of network traffic and performance data to support traffic management and capacity management functions.

- traffic management methods to provide monitoring of network performance and implement traffic management controls such as code blocks, connection request gapping, and reroute controls

Figure 1.7 illustrates the recommended traffic management functions.

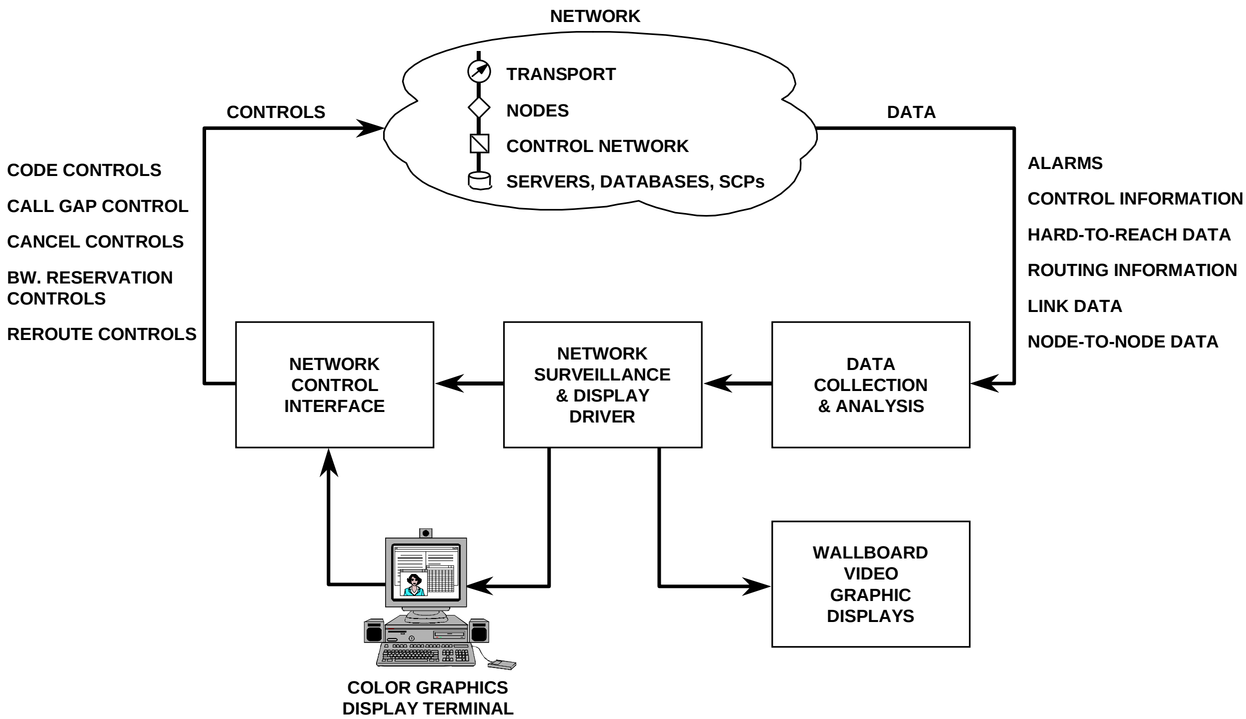

Figure 1.7. Employ Traffic Management Methods to Provide Monitoring of Network Performance and Implement Traffic Management Controls (such as code blocks, connection request gapping, and reroute controls)

- capacity management methods to include capacity forecasting, daily and weekly performance monitoring, and short-term network adjustment

Figure 1.8 illustrates the recommended capacity management functions.

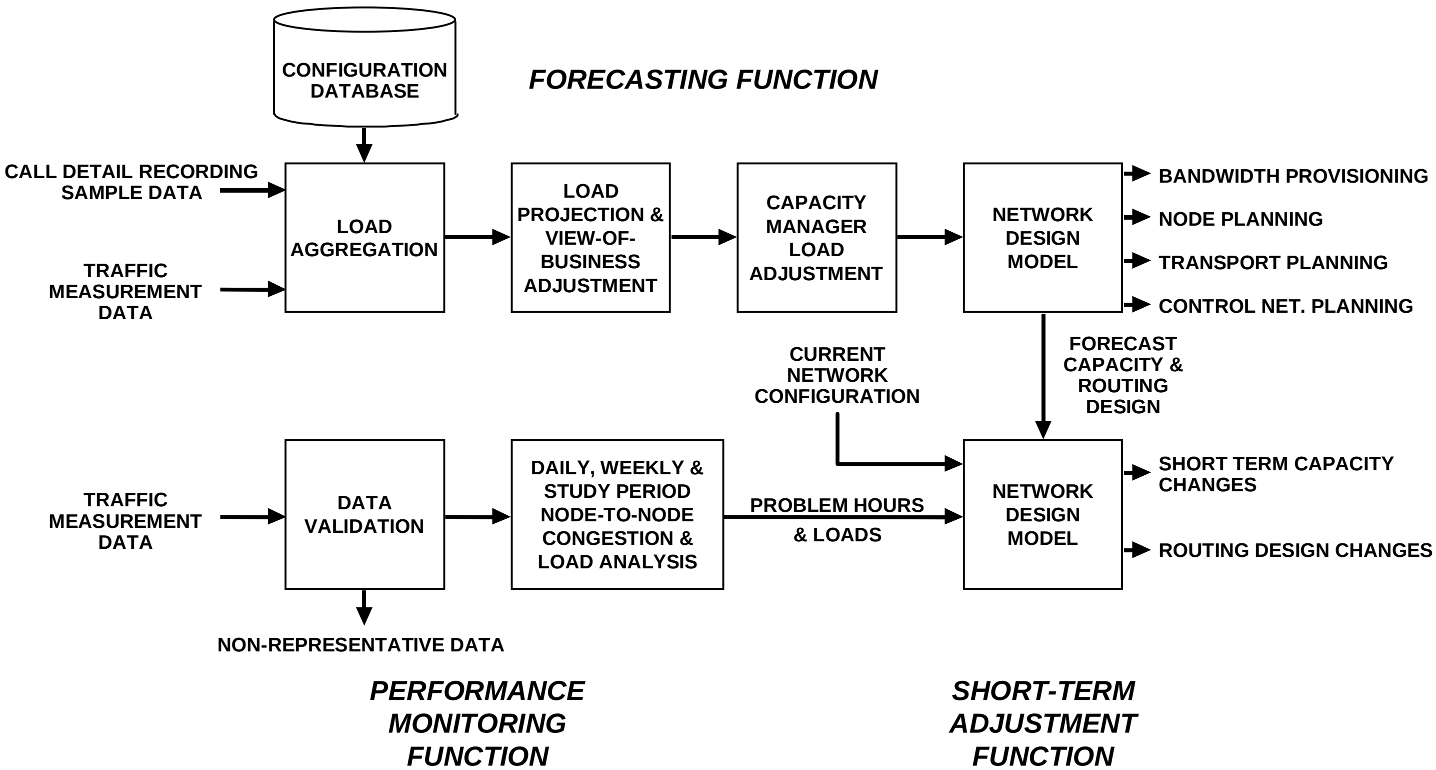

Figure 1.8. Employ Capacity Management Methods to Include Capacity Forecasting, Daily and Weekly Performance Monitoring, and Short-Term Network Adjustment

- discrete event flow optimization (DEFO) design models to capture complex routing behavior and design multiservice TE networks

Figure 1.9 illustrates the recommended DEFO design models. The greatest advantage of the DEFO model is its ability to capture very complex routing behavior through the equivalent of a simulation model provided in software in the routing design module. By this means, very complex routing networks have been designed by the model, which include all of the routing methods discussed in ANNEX 2, TDR, SDR, and EDR methods, and the multiservice QoS resource allocation models discussed in ANNEX 3. Complex traffic processes, such as self-similar traffic, can also be modeled with DEFO methods.

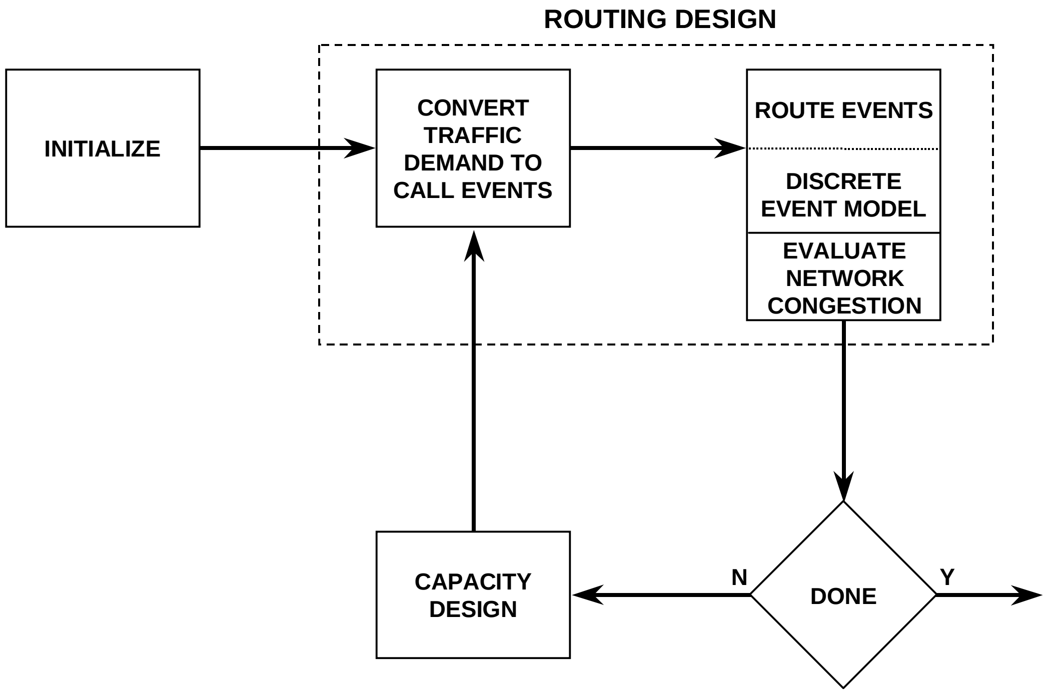

Figure 1.9. Use Discrete Event Flow Optimization (DEFO) Design Models to Capture Complex Routing Behavior & Design Multiservice TE Networks

10.6 Benefits of Recommended TE/QoS Methods for Multiservice Integrated Networks

The benefits of recommended TE/QoS Methods for IP-based multiservice integrated network are as follows:

- IP-network-based service creation (Parlay API, CPL/CGI, SIP-IN)
- lower operations & capital cost
- improved performance
- simplified network management

The IP-network-based service creation capabilities are discussed in ANNEX 4, the operations and capital cost impacts in ANNEXES 2 and 6, and improved performance impacts in ANNEXES 2 and 3.

Simplified network management comes about because of the following impacts of the recommended TE admission control and dynamic QoS routing methods:

- distributed control, as discussed in ANNEX 2
- eliminate available-link-bandwidth flooding, as discussed in ANNEX 4

- larger/fewer areas, as discussed in ANNEX 4
- automatic provisioning of topology database, as discussed in ANNEX 3
- fewer links/sparse network to provision, as discussed in ANNEX 2

11. Security Considerations

This document does not introduce new security issues beyond those inherent in MPLS and may use the same mechanisms proposed for this technology. It is, however, specifically important that manipulation of administratively configurable parameters be executed in a secure manner by authorized entities.

12. Acknowledgements

The author is indebted to many people for much help and encouragement in the course of developing this work. In the IETF I'd like to especially thank Dan Awduche of Movaz, Jim Boyle of Level 3 Communications, Angela Chiu of Celion Networks, and Tom Scott of Vedatel for all their helpful comments, assistance, and encouragement. Within AT&T I'd like to especially thank Chuck Dvorak, Bur Goode, Wai Sum Lai, and Jennifer Rexford for the excellent support and on-going discussions. In the ITU I'd like to particularly thank Mikael Ahman of Skanova Networks, Anne Elvidge of British Telecom, Davide Grillo of Alcatel Italia, Tommy Petersen of Ericsson, Bruce Pettitt of Nortel, Jim Roberts of France Telecom, Michael Tuxen of Siemens, and Manolo Villen of Telefonica for the valuable help throughout the course of this work. I'd also like to thank Professor Lorne Mason of INRS/University of Quebec for his many insights, discussions, and comments in the course of this work.

12. Authors' Addresses

Gerald R. Ash
AT&T Labs
Room MT D5-2A01
200 Laurel Avenue
Middletown, NJ 07748
Phone: 732-420-4578
Fax: 732-368-8659
Email: gash@att.com

13. Full Copyright Statement

INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.