




*Cross-Industry
Working Team*


Customer View of Internet Service Performance: Measurement Methodology and Metrics

October 1998



*Intel Telecom
American Management Systems
Ameritech
AT&T
Bell Atlantic
Bellcore
BellSouth
Cisco
Citicorp
Compaq
Corning
CyberCash
EarthLink Network
Ericsson
GTE Laboratories
Hewlett-Packard
Houston Associates
Hughes Network Systems
IBM
Intel
InterTrust
Inverse Network Technologies
Kaiser Permanente
Lucent Technologies
MCI Communications
Motorola
National Institute for Standards and Technology
New York Times
Nortel (Northern Telecom)
Novell
Philips Research Briarcliff
Pitney Bowes
QuantumLink
Real Networks
Science Applications International Corporation
SBC Technology Resources
Sun Microsystems
Texas Instruments
USWest
West Group*





*Cross-Industry
Working Team*

***Customer View of Internet
Service Performance:
Measurement Methodology
and Metrics***



Contents

- Abstract*
- 1.0 Introduction***
- 2.0 Measurement Methodology***
- 3.0 Metric Definitions***
- 4.0 Summary***
- 5.0 Glossary***
- 6.0 References***



Abstract

The Internet is increasingly critical for conducting business and engaging in commercial transactions. Consequently, attention is being focused on reliability and performance issues, and customers are being driven to negotiate service guarantees with their Internet service providers (ISPs). Such negotiations can be complex and time consuming: they are complicated by a lack of common definitions for metrics and a similar lack of agreed-upon methodologies for measuring and monitoring compliance. Furthermore, addressing these issues on an ad hoc basis between individual customers and ISPs is highly duplicative and could result in a large increase in Internet measurement traffic.

Through the promulgation of a common set of metrics and a common measurement methodology to assess, monitor, negotiate, and test compliance for service quality, a significant reduction in time and associated costs can be achieved and numerous benefits realized—for both ISPs and their customers.

A measurement architecture, measurement methodology, and common set of metrics are proposed in this paper. By limiting the scope to metrics and scenarios that are the most meaningful, a feasible architecture and methodology are derived; examples are given showing how they can be applied to assessing, monitoring, and testing compliance of service quality. It is expected that the application of this methodology will lead to improved Internet performance and foster greater cooperation between customers and ISPs.

1.0 Introduction

For large corporations and small businesses, government and private organizations, educational institutions, consumers, and individuals, the Internet has become an integral means for conducting business, engaging in commercial activities, and simply communicating. Increasingly, the Internet is being viewed as critical infrastructure (PCCIP 1997), and its contribution to economic output is growing.

As reliance on the Internet increases, expectations for reliability also increase. These expectations are driving customers to negotiate with their Internet service providers (ISPs) for guarantees that will meet customer requirements for specific quality-of-service levels.¹ This, however, poses a number of problems. First, users' perception of service quality can extend "end-to-end"; that is, remote networks that extend beyond the responsibility of the customer's ISP can dictate application-level service quality. Second, reaching agreement can be a complex and time-consuming task, encumbered by the myriad possible metrics that define service quality and the lack of any common definitions for these metrics. Third, there are no agreed-upon methodologies in place for measuring and monitoring negotiated metrics for compliance.

¹ Throughout this paper, we distinguish between an *end-user* and an *ISP customer*. The end-user is an individual who uses the ISP's services to access the Internet (browse the Web, for example); an ISP customer is a larger entity, such as a corporation, that pays the ISP to provide Internet services to its end-users (employees, for example).

If these problems are addressed on an ad hoc basis between each customer and its ISP, then the evaluation of performance and service quality will lead to a great duplication of effort and a large increase in measurement traffic. A common set of metric definitions and a common measurement methodology would significantly reduce these multiple independent efforts (and their associated costs) while facilitating the specification of service level requirements between a customer and an ISP, as well as the dissemination by an ISP (to its customers) of data on service quality.

This paper is intended as a first step in promulgating a common set of metrics and a common measurement methodology that can be used to assess, monitor, negotiate, and test compliance of service quality. The intended audience is ISPs and customers of ISPs. The hope is that the application of the metrics and methodology will lead to improved Internet performance and foster greater cooperation between customers and service providers.²

1.1 Benefits

Adaptation of common metrics and methodologies has specific advantages for both an ISP and its customers. Benefits to an ISP include the following:

- faster identification of performance problems, leading to better service for all customers;
- a means for providing evidence that service degradations are beyond the ISP's boundaries and thus the information needed to work with other ISPs to help resolve problems;
- a common language to facilitate the diagnosis and resolution of service-related problems;
- a means for differentiating service offerings from other ISPs; and
- a reduction in measurement traffic on the network—achieved by collecting common measurement data sets that can be shared—and consequent reduction in the need for independent measurements by customers.

For the customer, the benefits include:

- the opportunity to shift some of the task of monitoring and measuring service performance to the ISP,
- mechanisms for auditing the performance of a service provider,
- methods for quickly troubleshooting network problems and a means to better cooperate with ISPs to resolve these problems, and
- the ability to compare ISP service quality and make price-performance tradeoffs.

Additionally, both ISPs and their customers will benefit from the fact that less time will be needed to reach agreement on negotiated levels of service.

1.2 Scope

The metrics defined in this paper can be measured and monitored to determine service quality levels and, if necessary, used for auditing compliance with a negotiated and mutually agreed-upon contract. While such a model of cooperation is expected to be mutually beneficial, the metrics and methodology presented here are not to be limited to such a scenario:

² Measurements based on these metrics and this methodology are being collected by XIWT; results will be made available at a later date.

they can also be used to monitor intranets or service quality outside the domain of the customer's ISP. This paper focuses on metrics that relate to the network; application-level metrics are deferred to a future publication.

Several fundamental assumptions about these metrics are needed to ensure their relevance and to limit the scope of this effort to a practical level:

- The metrics are indicative of the quality of service perceived by the end-user; the metric values (measurements) are directly or indirectly affected by the quality of service offered by the ISP.
- Cooperation by an ISP is not a prerequisite to obtaining relevant and meaningful measurement data.
- The customer has direct dedicated access to a service provider (via a premises router). This consequently excludes dialup, ADSL, and cable modem users, for example. Nondedicated access will be considered in a future study.
- Service quality can mean different things to different users. The metrics must be broad and include both performance- and reliability-oriented measures.

Initially, active testing and measurement—in which traffic is introduced in order to collect data—will be the basis for assessing service quality. Passive measurement—in which measurement data are derived from normal Internet traffic—will be the subject of a future study and may be incorporated after careful validation and after appropriate monitoring tools are developed.

1.3 Related Work

Several other efforts are also under way that seek to define metrics for Internet quality of service. It is our intent to complement and build on these efforts, and to collaborate with them as appropriate, in order to meet the goals of improving Internet performance and encouraging greater cooperation between ISPs and their customers.

1.3.1 Internet Engineering Task Force (IETF)

IETF's Internet Protocol Performance Metrics Working Group (IPPM WG) is developing a set of standard metrics that can be applied to the quality, performance, and reliability of Internet data delivery services. While the working group will define specific metrics, actual implementations and applications are beyond its scope; rather, IPPM WG will promote the sharing of effective tools and procedures for measuring these metrics. The metrics will be defined so that the tools can be used by network operators, end-users, and/or independent testing groups.

IPPM WG has work in progress that defines a general framework for developing specific IP-level metrics (Paxson et al. 1998). To the greatest extent possible, our metrics and methodology will build on this framework.

1.3.2 Automotive Industry Action Group (AIAG)

AIAG's Automotive Network Exchange (ANX) effort aims to provide a large virtual private network for interconnecting automotive industry trading partners. This network of certified and monitored providers has been designed to reduce the data communication costs of the trading partners while meeting quality of service requirements in the areas of performance,

reliability, security, and administration and management (AIAG 1997). Service quality is characterized by requirements for mission-critical business-to-business communications.

AIAG has developed metrics and stringent requirements (AIAG 1997) focused on the specific needs of the trading partners. Specialized measurement tools (hardware and software) are needed to collect data for the metrics defined. Our effort, in contrast, is intended to be more generic and applicable to the general customer. The use of specialized hardware is specifically avoided.

1.3.3 T1A1

The T1A1 Network Survivability Performance Working Group, T1A1.2, studies network survivability performance by establishing a framework for measuring service outages and a framework for classifying network survivability techniques and measures. T1A1.2 is addressing the challenges posed by the Internet in maintaining a highly available, reliable, and survivable public switched telephone network (PSTN) (T1A1.2 1998a). T1A1.2 is also considering the impacts of the interaction between the PSTN and the Internet and addressing these impacts from an end-user's perspective. As part of its work, T1A1.2 will describe parameters for quantifying network failures and methods for measuring these parameters (T1A1.2 1998b).

This work is potentially related to the present efforts through its specification of and measurement methodology for reliability objectives.

Another T1A1 working group, the Performance of Digital Networks and Services Working Group (T1A1.3), is developing standards and technical reports describing performance and data services, and their multimedia integration within U.S. telecommunications networks (T1A1.3 1998). A current T1A1.3 project (T1A1-14) addresses the specification and allocation of Internet service performance, but no output documents will be available before the third quarter of 1998.

1.3.4 ITU-T Study Group 13

ITU-T Study Group 13 has developed Recommendation I.380, Internet Protocol Data Communication Service—IP Packet Transfer and Availability Performance Parameters. This recommendation is currently in its final form and is targeted for ITU-T approval in February 1999. Ongoing work related to Internet service performance will harmonize with Recommendation I.380 as appropriate.

1.3.5 Other Efforts

The **National Internet Measurement Infrastructure (NIMI)** effort (Mahdavi et al. 1997) uses software methodologies and tools as a vehicle for testing measurement strategies and as the basis for a ubiquitous infrastructure (i.e., one that is readily available and easy to deploy). The challenges NIMI seeks to address include the following:

Scaling—How to deal with millions of nodes and with costs for end-to-end measurements that increase as the square of the number of nodes.

Mapping—How to discover the topology needed for path decomposition and the statistical distribution of end-to-end tests.

Optimization—How to minimize the amount of measurement traffic while maximizing the value of the data.

Result distribution—What techniques, such as caching, should be employed for scalable data movement.

NIMI is especially aimed at solving the broader, longer term issues that will arise when many entities are involved in measuring and characterizing Internet performance. Until a ubiquitous infrastructure is available, focused efforts such as ours will be needed. By concentrating on the customer-ISP relationship, our initiative will help define the types of measurements NIMI will need to address.

Surveyor is a joint initiative of Advanced Network & Services, Inc. (ANS), and the Common Solutions Group, a collaboration of 23 universities jointly addressing challenges to networking (ANS 1997). Surveyor is focused on an early implementation of metrics developed within IPPM WG. Measurement devices (Surveyor tools) will be deployed at each university, and several traffic measurements will be taken from these sites, including estimates of one-way delay and packet loss. The measurement devices are PCs (running FreeBSD) equipped with GPS antennae to provide accurate (± 50 microsecond) time stamps. A database system is being developed to analyze data and permit Web-based access by participants. This initiative parallels our own measurement effort, with the major differences being the nature of the participants (XIWT comprises corporations and government organizations); the types of tools used (ours comprise tools developed by the high-energy physics community); and, to some extent, the types of traffic measurements (XIWT will focus initially on round trip delay and packet loss).

1.4 Organization of This Paper

The remainder of this paper is organized as follows. Section 2.0 defines the XIWT measurement architecture and methodology. These are motivated by a set of ideal requirements, an assumed generic topology, and a set of usage scenarios representative of how users access data. Examples of how the measurement methodology can be applied to typical usage scenarios are also provided. In section 3.0, definitions of the relevant performance and reliability metrics—measurable using the methodology described in section 2.0—are provided. Guidelines for aggregating measurements to provide useful estimates of long-term performance and details on the techniques for measuring and computing the metrics are also given. Section 4.0 provides a summary and an outline of future work.

2.0 Measurement Methodology

2.1 Requirements

Ideally, a measurement methodology should meet the following requirements:

Isolate sources of problems—Since one of the motivating factors for defining the metrics is to isolate sources of problems in end-to-end service, the measurement methodology should help identify administrative domains responsible for the problem, thus ensuring accountability. For example, measurements should isolate whether the problem is within an ISP's domain, and, if so, should indicate components that may be the cause of the problem.

Provide meaningful results—The results must facilitate communication between customers and operators. Thus, for example, while a “router queue length” metric may be very meaningful to an operator, it is unlikely to be used by an ISP customer. The customer may be more interested in a “packet loss” or “response time” metric.

Not require new infrastructure—It should be possible to make the measurements using existing infrastructure. Although there may be cases where operators or customers may install measurement instruments for improved efficiency or ease of administration, the methodology itself should not require this.

Avoid unnecessary or duplicate measurements/traffic—Measurements that are required to meet multiple customer/ISP needs should not be replicated; instead, the results should be shared as appropriate. This ensures that nonvalue-added traffic is minimized over the network.

Be auditable—Since the intent of the measurements is to communicate service quality metrics between suppliers and consumers, the architecture should support independent measurement and validation.

Be robust—Because a large part of the reason for measurement is to detect and correct problems, the methodology must be robust in the presence of faults and should contain checks to ensure measurement accuracy.

2.2 Assumed Topology

Figure 1 presents a general topology of an ISP network

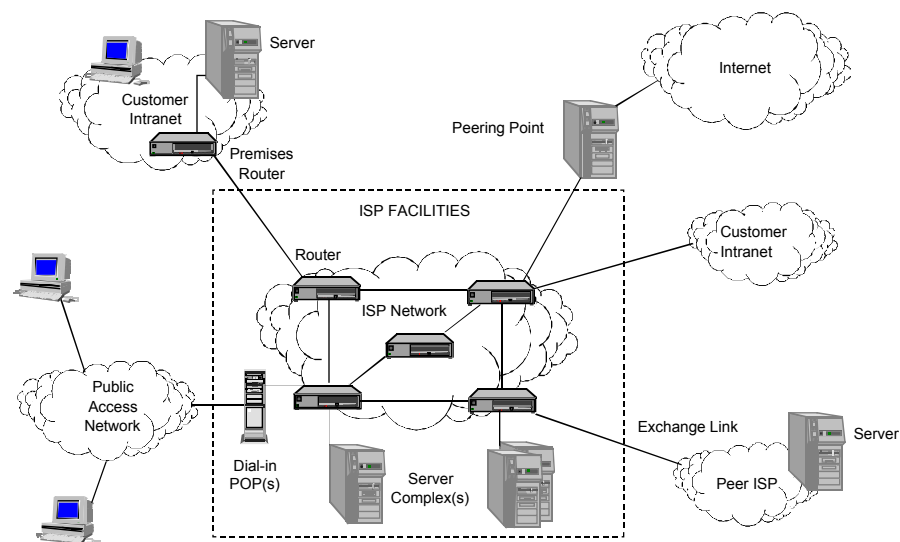


FIGURE 1. ISP Network Topology

In the figure, we assume that the ISP network consists of a number of routers communicating over links (an IP “cloud”). The ISP has servers providing services (DNS, e-mail, Web hosting) on hosts inside server complexes that are connected to the IP cloud and one or more POPs for dialup connections from the public access network. The IP cloud controlled by the ISP is connected to the other parts of the Internet either through peering points (e.g., public NAPs) or through exchange links based on private peering arrangements. The ISP also provides IP services to customers through dedicated links. As mentioned above, nondedicated access is not considered in this paper; therefore, measurements of the public access network are not discussed. We consider that the ISP boundary is demarcated at the premises routers located at the customer’s facilities and at the ISP border routers at public or private peering points.³ The methodology and metrics defined below can easily accommodate other demarcation points.

2.3 Usage Scenarios

Since any performance metrics are intended to determine the experience of the end-users (as opposed to the customers), the different ways in which services affect end-users must be taken into account when defining these metrics. Table 1 shows the scenarios that occur when end-users access data owned by the ISP customer. Note that, in the table, we assume that end-users *within* the customer intranet are corporate employees; those *outside* the customer intranet may be corporate employees, clients, or other unrelated end-users. While not comprehensive, this table helps identify the various conditions that can be encountered by end-users.

TABLE 1: Usage Scenarios for Internet Services

		Data location		
		At customer site	At ISP facility	Beyond ISP boundary
End-user location	At customer site	N/A	2. Corporate end-users accessing data hosted by ISP	4. Corporate end-users accessing data on the Internet or at remote sites
	Directly connected to ISP facility	1. End-users accessing data located at customer premises from remote sites or the Internet	3. End-users accessing data hosted by ISP	5. End-users accessing data hosted by third party
	Beyond ISP Boundry			N/A

In each of the scenarios, the customer is interested in both the *reliability* as well as the *performance* of the service as seen by the end-users. Note that in certain cases (e.g., case 3),

³ The ISP boundary is usually at a router interface. This could be the “inside” (customer-side) interface or the “outside” (ISP-side) interface, depending on whether the ISP or the customer owns the router. Boundaries between different ISPs are defined similarly at router interfaces depending on who owns the router. We here assume that an appropriate demarcation point can be identified at or near these boundaries.

there is no direct way for the customer to measure service from its site—it has to rely on measurements made by the ISP or some third party, or rely on users to notify the ISP that the data are not accessible.

These scenarios demonstrate that end-to-end service quality depends on the ISP facilities, customer facilities, and—in many instances—on elements that are outside the control of either party.

2.4 Measurement Architecture

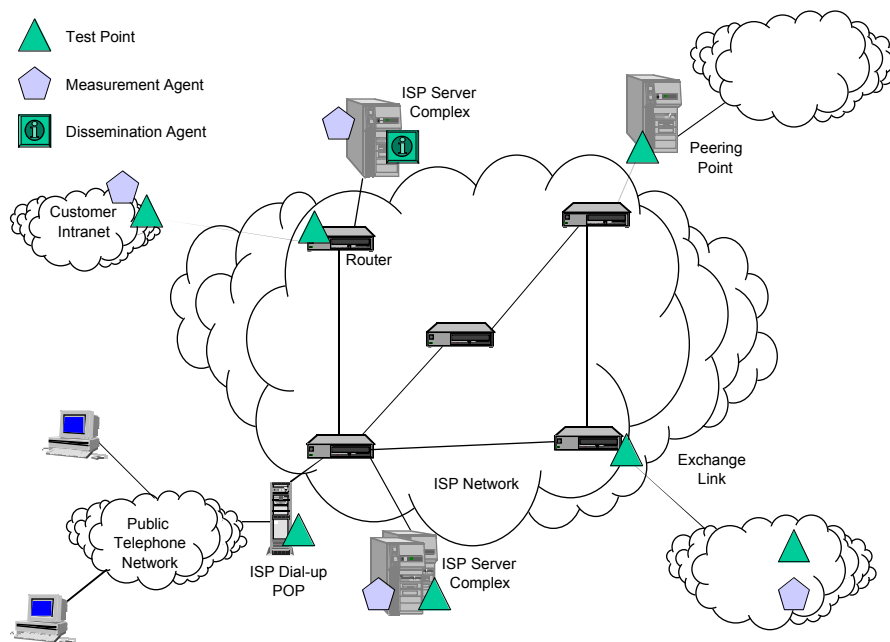


FIGURE 2. Measurement Architecture

Our measurement architecture, which is shown in figure 2—and which is very similar to that defined by AIAG (1997) for the ANX—uses a “black box” approach, where the metrics and measurements are defined in terms of externally visible properties. The architecture is defined in terms of the following logical components:

Test point—Test points are hosts that either collect performance data or have been configured to respond to measurement queries (Paxson et al. 1998). Most existing equipment is already designed to respond to measurement queries such as those using SNMP or ICMP. It is assumed that other software may be added to test points to enable measurements if desired. Providers may decide to install separate hosts as test points to reduce load on critical network elements such as routers. Our architecture assumes that appropriately configured test points have been identified at various places in the network and on customer sites.

Measurement agent—A measurement agent is software that runs on a host and actually initiates the various measurements. Measurement agents communicate with test points to conduct the measurements or collect data (such as SNMP data) present at the test points. Measurement agents may also examine log files and simulate service usage to measure performance.

Dissemination agent—A dissemination agent provides, to interested parties, the results of the measurements that have been collected. The actual form/format of, and mechanisms for, this information dissemination are left open. Based on need, the results from various agents may be combined/correlated before dissemination or may be left in a “raw” form. The process by which data are transferred from the measurement agent to the dissemination agent is also left to the implementation.

We assume that appropriately configured test points can be identified (or located) as needed to enable measurements. Measurement agents may run on both the customer site as well as on the ISP network and server complexes. In many instances, both the measurement agent and the test point may be co-located on the same host.

2.5 Methodology

For each measure of interest, *the location of both the measurement agent as well as the test points used are specified* along with the metric. This defines exactly what is being measured and within which domain(s) the responsibility for maintaining performance lies. Proper identification of measurement agents and test points allows different parts of the network or service to be tested independently, and thus allows problems to be isolated to specific segments of the network.⁴ In particular, *if test points exist at administrative boundaries, it may be possible to isolate problems to either inside or outside those boundaries*. This lets providers offer service guarantees to customers for their part of the service, even when the end-to-end service spans multiple providers.

As part of its operations, the ISP’s staff needs to monitor the health of the network and other services provided by the ISP. Thus, we assume that the ISP can run measurement agents on various network hosts and configure various network and service elements as test points. By making some of these measurements available through a dissemination agent (e.g., by providing the measurement summary data at a Web site), the ISP can provide the results of these tests to its customers and peers.

Even if customers rely on the ISP for most measurements, they may conduct their own measurements (1) to validate (audit) the results provided by the ISP, and (2) to measure items of interest to them that are not being provided by the ISP. Customers may also choose to have third parties make such measurements on their behalf. To do so, a customer may ask the ISP to make selected test points available to it. The customer may also make some of its own test points “public” to enable other interested parties to test overall performance in communicating with the customer’s site.

2.6 Examples

This section shows how the methodology described above can be applied in the various usage scenarios from section 2.3 to monitor service reliability and performance.

2.6.1 Service Level Agreement (SLA) Monitoring

In this example, we illustrate how an ISP can provide SLA reporting capabilities to its customers. Consider a customer that has signed up with the ISP for Internet access and has

⁴ Strictly speaking, this requires that the end-to-end measurement be decomposed into independent measurements over the individual segments, which may not always be possible. However, for the purpose of coarse-grained fault isolation, a segment that shows problems has a high likelihood of being the cause of the end-to-end problem in the absence of other detectable faults.

outsourced its e-mail and Web site to the ISP for hosting (figure 3). The ISP has offered an SLA to the customer with guaranteed availability for network access, as well as availability guarantees on the mail and Web services. The ISP has peering arrangements at two points on its network and runs a server complex where services are provided. Let us assume that the ISP provides the customer with a premises router which it continues to own and manage.

The ISP identifies the two border routers and the customer premises router as the test points for checking customer access to the Internet; it also identifies the Web and e-mail servers as test points to meet its Web and e-mail service guarantees. The ISP locates measurement agents near the customer premises router and at the data center. It periodically tests (A) the customer premises router as well as (B) the border routers to ensure that they are up, and makes network measurements (C) to ensure that the customer traffic can reach the border routers. The ISP also makes periodic measurements (D) to check that the customer’s e-mail and Web hosting services are up. The results are aggregated and made available (E) to the customer using a customer service Web site, where the customer has access (F) to the most current data.

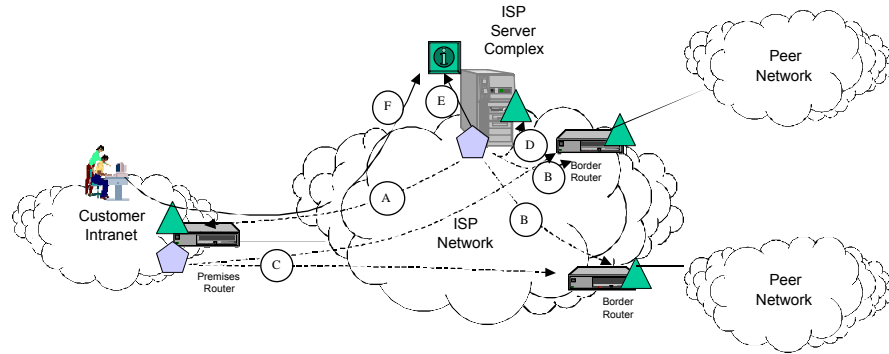


FIGURE 3. SLA Monitoring for Internet Services

Figure 4 expands this example to show how measurement agents and test points can be identified for each of the usage scenarios shown in table 1. We assume now that the ISP has located measurement agents at each of its border routers, and that the customer has located a measurement agent at its premises router to make independent measurements. The border routers (and the customer premises routers) are enabled to act as test points (or the ISP has identified hosts connected to the routers for this purpose). We also assume that the measurement agents can use the data servers as test points.




FIGURE 4. Corporate Data Access Performance Measurement

Case 1: External end-users accessing data located at a customer-owned server at the customer site—In this case, the customer is interested in end-users accessing data on the customer’s server from outside its intranet. The user experience in this case requires instrumentation of the user’s client-machines. However, measurements taken from measurement agents A and B to test point 1 can provide client experience *insofar as it is affected by the ISP*. As an alternative, if the ISP provides only test points 4 and 5, the customer can obtain an approximation of client performance by making measurements using measurement agent C to test points 1, 4, and 5. Assuming that test point 6 is located as close to the premises router as possible, it can be used to determine if the problem is on the ISP network or on the customer intranet.

Case 2: Corporate end-users accessing data hosted by the ISP—In this case, measurement agent C and test point 2 located at the ISP-hosted server would be used. Note that in this case the customer can directly make these measurements or can choose to outsource them to the ISP.

Case 3: External end-users accessing data hosted by the ISP—As in case 1, the ISP could use measurement agents A and B with test point 2 to measure performance. In this case, the customer cannot directly measure performance from inside its intranet without relying on third-party measurements.

Case 4: Corporate end-users accessing data on the Internet—The customer can make direct measurements using measurement agent C. If the ISP provides test points 4 and 5, it is also possible to decide if problems in accessing the Internet are within the ISP’s boundary or beyond it. Test point 6 can be used to determine if the problem is on the customer’s intranet.

Case 5: External users accessing third-party hosted data—This is similar to case 3, with measurement agents A and B using test point 3 to make the measurements. As in case 3, it is difficult for the customer to make this measurement directly.

2.6.2 ISP Performance Comparison

In this example, we assume that a multi-homed customer wishes to compare how its ISPs are meeting the needs of its employees who are interested in reaching a number of sites on the Internet.




FIGURE 5. Performance Comparison Between Two ISPs for Accessing Sites on the Internet

Figure 5 shows this scenario. The customer selects the set of sites (as test points) and uses one or more agents on its intranet to access those sites through both ISPs (paths A and B).

Because the same destinations are used in both cases, this is a meaningful comparison. If path A provides better performance most of the time (as measured to the test points of interest), the customer may consider the service offered by ISP1 to be better than that provided by ISP2, even if the problem is in fact on the Internet beyond ISP2’s boundary.

2.6.3 Network Monitoring and Domain Isolation

As mentioned earlier, multiple network providers are usually involved in providing end-to-end service. In the SLA monitoring example, we assumed that the ISP provided guarantees only to its boundaries. This is likely to be unacceptable to most customers. Clearly, to guarantee end-to-end performance, multiple ISPs need to cooperate. In this example, we show how the methodology can be used to isolate problems in end-to-end service by cooperating ISPs.

Consider the scenario in figure 6. In this case, a customer uses ISP1 and is interested in reaching sites A to D. Clearly, since site A is on ISP1’s network, ISP1 can measure (and guarantee) performance to A. Let us assume that ISP1 has agreements with ISP2 to share performance data and to permit access to appropriate test points. ISP2 has similar agreements with ISP3. By having access to measurement data from all three domains, ISP1 can provide end-to-end performance measurements to the customer and isolate which ISP domain is causing a customer-visible problem. Additionally, if the SLAs between the ISPs provide service guarantees, ISP1 would be able to extend these guarantees for access to sites B and C to the customer.

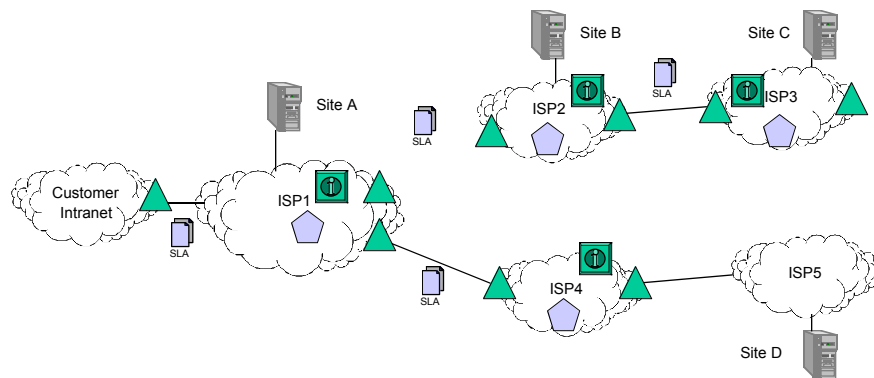


FIGURE 6. Network Monitoring and Domain Isolation

What happens if some ISP does not have such an agreement? This is shown for site D. In this case, ISP4 does not provide any guarantees about site D to ISP1, which in turn cannot provide any guarantee to the customer about site D. However, because the benefit of such agreements is mutual, ISP4 can also deny ISP5 access to any performance data on its network.

2.6.4 Virtual Private Networking

In this example, the customer has used the Internet to create a virtual private network among multiple sites, as shown in figure 7.




FIGURE 7. Virtual Private Networking

The customer locates measurement agents (and test points) as close to the premises routers as possible at each of its sites and uses them to measure performance. If a customer site is located on the network beyond the ISP's boundaries, the ISP can locate measurement agents near its boundaries (as in the SLA monitoring example) to assure the customer that the problem is not within its domain.

3.0 Metric Definitions

As should be obvious from the above examples, the architecture requires the following:

- The metrics of interest must be specified in terms of quantities (such as round trip delay) that can be measured by appropriately configured measurement agents and test points.
- The location of measurement agents as well as test points must be specified to allow a clear understanding of what part of the network or service is being measured. We suggest that test points be located at administrative domain boundaries to ensure accountability when multiple service providers cooperate to provide end-to-end service.

This section discusses some metrics of particular interest to customers and service providers. While metrics such as these are widely used, it is our goal in this section to define them more precisely using the methodology described above; they can then be more efficiently and effectively used by customers and service providers when making “apples-to-apples” comparisons and communicating service quality concepts.

3.1 Performance Metrics

Performance metrics quantify end-user visible perceptions of service performance. For network services, commonly available tools such as “ping”—which measures round trip time using ICMP—can be used to determine quantities such as *packet loss* and *round trip delay*. As throughput measurement tools become more widely available, *throughput* could also be used as a performance metric. And, with the spread of real-time multimedia applications, it is expected that *packet jitter* (Demichelis 1998) and *one-way delay* (Almes et al. 1998) may also become important. However, because few widely available tools exist to measure jitter and one-way delay, we do not use them in our definitions at the present time.⁵

⁵ While we believe that *throughput* and *response time* measured at the network level are important metrics, they should also be measured in an application context. We defer their discussion to a later paper.

We recommend that (at a minimum) network performance be measured using the following metrics.

3.1.1 Packet Loss

Packet loss is defined as the fraction of packets sent from a measurement agent to a test point for which the measurement agent does not receive an acknowledgment from the test point. This includes packets that are not received by the test point as well as acknowledgments that are lost before returning to the measurement agent. Acknowledgments that do not arrive within a predefined round trip delay (see below) at the measurement agent are also considered lost.

3.1.2 Round Trip Delay

Round trip delay is defined as the interval between the time a measurement agent application sends a packet to a test point and the time it receives acknowledgment that the packet was received by the test point. Round trip delay includes any queuing delays at the end-points or the intermediate hosts, but does not include any DNS lookup times by the measurement application.⁶

3.2 Reliability Metrics

System reliability is particularly important in business-critical applications. Because of the nature of the Internet, it is difficult for ISPs and customers to specify service levels based directly on reliability (which typically indicates only if a system is operational), and terms such as “availability” are used instead. However, because no common definitions and computational methodology exist for such terms, it is often difficult to negotiate reliability guarantees—and harder still to compare quality of service based on quoted values. We believe that from a customer perspective, there are three components to service reliability:

- Can the customer *reach* the service?
- If so, is the service *available*?
- If not, how *frequently* and for how *long* do the outages last?

We capture these components in the following metrics.

3.2.1 Reachability

A test point is considered *reachable* from a measurement agent if the agent can send packets to the test point and, within a short predefined time interval, receive acknowledgment from the test point that the packet was received. In most instances, we can consider the ping test (described in section 3.1) as a sufficient metric of reachability.

Thus, if each measurement sample consists of multiple pings, the test point is considered reachable from the measurement agent if the latter receives at least one acknowledgment from the test point.

⁶ Note that we define round trip delay to include the queuing delay at the end hosts including the one running the measurement agent. The measurement agent can estimate the delay at the local host by using a local “loopback” to ensure that the measurement host is not contributing significantly to the delay measured value.

This definition can be extended for various topologies to derive more general reachability metrics. For example, an ISP may consider a host on the Internet reachable from a customer site if:

- a measurement agent placed just before the customer premises router can send packets to a test point placed just beyond the last router on the ISP's network that has a common initial subpath with the site being accessed; and
- within a short time interval, receive acknowledgment that the test point received the packets.

Note that unless the site is on the ISP's network, this does not imply that the customer can actually reach the site in question. Rather, it implies only that the problem, if any, is not in the ISP's network and that the ISP is, in fact, meeting its commitment to provide access to sites on the Internet.

3.2.2 Network Service Availability

The network between a measurement agent and a test point is considered available at a given time t if, during a specified time interval Δ around t , the measured packet loss rate and the round trip delays are both below predefined thresholds.

Network service availability is defined as the fraction of time the network is available from a specified group (one or more) of measurement agents to a specified group of test points.

Again, this definition depends on network topology. Thus, for example, a customer may consider an ISP's network service to be available if, over a predefined interval,

- test points located just before peering points and exchange links on the ISP's network are reachable from a measurement agent located near the customer's premises router, and
- the packet loss rate and round trip delays to the test points from the measurement agent are all below predefined thresholds.

3.2.3 Duration of Outage

The duration of an outage is defined as the difference between the time a service becomes unavailable and the time it is restored. Note that because of the statistical nature of Internet traffic, the duration over which service is measured to be unavailable should exceed some minimum threshold before it is declared an outage. Similarly, when service is restored after an outage, it should stay available for some minimum duration before the outage is declared over.

3.2.4 Time Between Outages

Time between outages is defined as the difference between the start times of two consecutive outages.

3.3 Ancillary Metrics

Ancillary metrics are often needed to interpret the results obtained from direct measurement of performance or availability. Specifically, most performance metrics depend on the utilization of underlying resources (defined below) and of network services. For example, DNS, a critical network service, is now used by almost all applications to resolve host names to the corresponding IP addresses. A performance problem with DNS results in a large variety of symptoms including poor response time, or even failure, of applications.

3.3.1 Network Resource Utilization

Network resource utilization is the percentage of a particular part of the network infrastructure used during a given time interval. It is a dimensionless number calculated by dividing the amount of the particular resource used during a given time interval by the total theoretically available amount of that particular resource during that same interval. Measuring resource utilization is especially important for key resources that include links and routers. Both utilization peaks and percentiles must be monitored.

3.3.2 DNS Performance

DNS has become an increasingly important part of the Internet because almost all applications now use it to resolve host names to IP addresses. As a result, application-level response times can appear slow if DNS performance is bad. We define DNS performance measures using two metrics—*DNS query loss* and *DNS response time*.

DNS query loss is defined as the fraction of DNS queries made by a measurement agent for which the measurement agent does not receive a response from the DNS server within a predetermined time. This definition is analogous to the packet loss definition cited earlier.

DNS response time is defined as the interval between the time a measurement agent application sends a DNS query to a DNS server and the time it receives a response from the server providing the result of the query. This is analogous to the round trip delay metric.

3.4 Tradeoffs Between Metrics

In general, the metrics described above are related, but provide different types of information to the customer. Good results with one metric may be balanced by a poor showing in another. If time between outages is high, that is good. But if the duration of the outage is very long when it does occur, the overall picture is not good, despite the infrequency of outages. Similarly, if an ISP network demonstrates good performance metrics for points within its domain, but poor peering arrangements lead to poor reachability of other sites on the Internet, the ISP's service may not be acceptable to customers.

Actions taken to improve one metric may have a negative effect on others. For example, round trip delay can be improved by reducing router queue lengths. While this may have the positive effect of reducing delay, it may also have the negative impact of increasing packet loss.

Finally, most metrics depend on the utilization of resources in the network. If a customer uses low-capacity links to connect to the ISP and traffic patterns show that the link runs near capacity most of the time, obtaining higher levels of performance requires subscription to a higher speed access link. Metrics should thus be evaluated as a group to provide a complete picture of the service and to examine the tradeoffs of improving one metric at the expense of another.

3.5 Aggregation of Measurements

Because the Internet is based on best-effort protocols, instantaneous performance does not relate directly to long-term performance. Although individual measurements can be used to detect operational problems in near real time, metrics of interest usually need to be aggregated over time to obtain valid estimates of performance. Due to system complexity, it is difficult to predict the performance that can be achieved a priori; it is thus necessary to compute baselines that can be used for setting quality-of-service targets and for comparing performance.

Statistical aggregates such as means or standard deviations are not appropriate to quantify performance of data networks because the underlying primary metrics have “heavy-tailed” distributions that are not represented well by those aggregates. These metrics can be more appropriately represented by *percentiles* (Bickel and Doksum 1977) and *order statistics* such as the median.

We define each measurement aggregate A as follows:

Measurement values V_m —Each measurement sample M results in a value V_m . Note that in most cases, V_m will itself be computed from a set of measurements. Thus, V_m could be the fraction of responses received when a given host is pinged 10 times at one-second intervals, or it could be the median round trip delay computed from the returned responses. As another example, V_m could represent the median packet loss measured between a group of sites in North America and a group in Europe.

Measurement interval I_m — I_m is the interval between measurement samples. If measurements occur at random times, then I_m is the expected value of the interval associated with the measurement. Thus, for example, a measurement may be taken every five minutes (periodic) or at intervals that are Poisson distributed with expected time of arrival equal to five minutes. Note that I_m defines the temporal resolution of the measurements—i.e., events that are shorter than I_m in duration are likely to be missed by the measurements.

Baseline value B_m —A baseline value represents the expected value of M . Baselines may be static (i.e., time invariant) or dynamic (time dependent). Baselines may also be dependent on service load and other system parameters. It is expected that under normal circumstances, baselines will be computed from historical records of measurement samples. As noted above, the sample mean is a poor baseline value, and the median is probably a better baseline.⁷

Baseline spread S_m —The baseline spread is a measure of the normal variability associated with M . As with baselines, the baseline spread may be static or dynamic. Once again, we believe that spread should be computed using quartiles or percentiles rather than the standard deviation. A measurement is considered to be within baseline if

$$|V_m - B_m| / S_m \leq T_m$$

where T_m is a threshold.⁸ If the underlying measurement distributions are significantly asymmetric, the baseline spread may be specified in terms of an upper specification limit U_m and a lower specification limit L_m .

⁷ Some XIWT members are currently making measurements over the Internet by periodically pinging one another’s sites. We expect to use these data to try various means of baselining and will report on our findings in a subsequent paper.

Aggregation interval I_a — I_a is the time over which multiple measurement samples are aggregated to create the metric A which is representative of system behavior over that time. Typically, aggregation may be provided over hourly, daily, weekly, monthly, quarterly or annual intervals.

For example, an SLA may contain multiple aggregation intervals over which performance is measured. Note that aggregation intervals may be disjoint, i.e., aggregation may occur only at peak times or during business hours.

Aggregate value F —The aggregate value F_a is the fraction of measurements that are within baseline over the aggregation interval. Thus, if N measurements are taken over the aggregation interval, and N_b are within baseline, the aggregate value is

$$F_a = N_b/N$$

Bounds may be placed on F_a to specify “acceptable” service behavior. Measurements that return illegal or unknown values (e.g., if all packets are lost in a round trip delay measurement) should normally be considered out of baseline for computing the aggregate value.⁹

Note that while aggregation intervals used to compute F_a are likely to be large for monitoring and planning purposes, alarms may be generated by as few as one or two sequential measurements if they are sufficiently out of baseline.

Also note that we resist the temptation to arrive at an aggregate value by averaging the measurement values. In our opinion, because of the long time intervals involved in aggregation, such values do not provide meaningful conclusions about service quality. The baseline value(s), however, can be used for historical comparisons if they are dynamic.

3.6 Computation of Metrics

In this section, we use the SLA monitoring example (section 2.6.1) and the corporate data access scenarios (table 1) to show how to measure and compute the metrics. For each scenario, we should measure the following:

- packet loss,
- round trip delay,
- network service availability,
- time between outages, and
- length of outages.

Packet loss and delay estimate the performance that users expect to see when the network is “operational.” Network service availability, time between outages, and length of outages quantify service availability to the customer’s Internet users.

⁸ Division by S_m allows T_m to be specified as a dimensionless quantity. This is useful for specifications such as “take action when measurement falls outside the 99 percentile value.” The threshold could also be specified in the same units as V_m , in which case normalization need not be done.

⁹ Strictly speaking, such measurements should be considered out of baseline only if the loss of data is caused by the system and not due to problems in the measurement process itself. Determining the cause of data loss is difficult, and significant effort is required in the implementation to decide how to deal with this.

Case 1: External users accessing data located at customer premises—In this case, the customer has some resources that it wishes to make available for users on the Internet. This case encompasses virtual private networking where the external users are employees of the customer connecting through the Internet to resources within the company. Ideally, the customer should instrument a test point as close as possible to the demarcation between its intranet and its ISP. The ISP should provide measurement agents as close to its peering points with other ISPs as possible. These measurement agents can then be applied to estimate the performance for users of those ISPs, aid in problem isolation, and give the customer some idea of how well its ISP peers with other ISPs.

The measurement agents will sample at some given interval—5-, 10-, 15- or 30-minute sampling intervals are typical. A measurement agent can ping each test point and then use the results to measure round trip delay, packet loss, and reachability. These measurements can then be used to determine availability during the sampling interval. When a test point is not reachable, the measurement agent records the time that the outage occurred. When the outage has been repaired, duration of the outage can be calculated with these data. When another outage occurs, time between outages can be recorded.

Measurement first needs to take place for a few weeks (a minimum of two) to establish baseline values and baseline spreads for all metrics. The objective is to establish norms of network behavior. Since some of the measurement agents are not on the customer's ISP, there may be different levels of performance and availability depending on the measurement agents' location. In this particular case, aggregation is an important issue. Aggregation by ISP or geographic region will help isolate peering problems between ISPs or within local regions.

Once the baseline values and spreads have been determined, operational measurements can begin. Measurements that fall sufficiently far outside the baseline can be flagged as problems. Comparisons between the measurements made by agents on the customer's ISP and agents on other ISPs can be used to decide if the problem is due to the customer's ISP or beyond it.

There are two difficult aspects to this particular case. The first aspect is gathering data from agents on other ISPs. This can be done by cooperative associations (e.g., XIWT's Internet Performance Working Team) or by third-party vendors (e.g., Keynote or Inverse Networks). The second difficulty is what to do when problems are found. ISPs have little direct control to fix problems on other ISPs' networks. This is an issue that may be addressed by bodies such as Internet Operators (IOPS.ORG 1998).

Case 2: Corporate users accessing data hosted by the ISP—In this case, the customer's users access data hosted by the customer's ISP. A test point should be located at the ISP's hosting center. A measurement agent should be located as close as possible to the customer's demarcation point with the ISP. After a baseline is established (as described in the previous case), operational measurements can begin and deviations from established baseline spreads flagged for ISP notification. All the data can be aggregated together. Monitoring performance and correcting problems in this case is much easier because only one ISP is involved.

Case 3: External users access data hosted by the ISP—Here, users on the Internet access the customer's data hosted by the customer's ISP. This case is similar to case 1, except that the test point should move to the ISP hosting facilities. Measurement agents should be stationed within the customer's ISP and on other ISPs. As with case 1, a key point of interest is the ability of the customer's ISP to peer with other ISPs.

Data should, as in case 1, be aggregated by ISP or geographic region where the measurement agents are located. And again, as in case 1, issues arise with regard to fixing problems when they occur at an ISP other than the customer's ISP.

Case 4: Customer's employees accessing data on the Internet—This is the inverse of case 1. A measuring agent should be located as close to the ISP's point of demarcation at the customer's premises as possible. Test points should be distributed at the customer's ISP and other ISPs. This is the exact opposite placement of test points and measurement agents as in case 1. Instead of aggregating data by ISP or geographic region of the measurement agents, data should be aggregated by ISP or geographic region of the test points.

As in cases 1 and 3, difficulties exist in fixing problems when these occur at an ISP other than the customer's own.

Case 5: External users accessing customer's data hosted by third party—This is basically case 3, but the data are hosted at a location not on the customer's ISP. This scenario often occurs when a customer wants to outsource handling of its content. The test point then moves to the third party's data center. As in case 3, aggregation is done by ISP and geography. And again as in case 3, there may be difficulties in debugging problems that occur at ISPs other than the customer's.

4.0 Summary

To simplify the process of negotiating service quality agreements between customers and their ISPs, a measurement architecture and methodology, and a common set of metrics that can be quantified using this methodology, have been described. Motivated by ideal requirements and common usage scenarios, and scoped by metrics most meaningful to both customers and ISPs, the methodology is both practical and applicable to a wide range of performance and reliability issues. Guidelines for aggregating measurements to provide useful estimates of long-term performance and details on the techniques for measuring and computing the metrics have also been given.

Ongoing effort in this area by XIWT is expected to address the following outstanding issues:

- The feasibility of the proposed architecture and methodology needs to be demonstrated with data from an example implementation. XIWT members are currently running such an experiment, and it is expected that initial results will be published in the near future.
- The measurement architecture and methodology have been limited to customers with dedicated access to their ISP (via a premises router). For other customers with dedicated access (such as small business customers with xDSL connections or individual customers with cable-modems, for example), the architecture and methodology ought to be applicable. The subject of dialup customers also needs to be addressed.
- The methodology has focused on active measurements. The use of passive measurements should be investigated, since these can significantly reduce the amount of overhead traffic.
- The metrics definitions should be expanded to include those relevant to specific applications (e.g., TCP- and HTTP-level measurements, jitter for voiceover IP applications, etc.).

5.0 Glossary

ADSL	asymmetric digital subscriber loop
AIAG	Automotive Industry Action Group
ANX	Automotive Network Exchange
DNS	Domain Name System
DSL	digital subscriber loop
GPS	global positioning system
HTTP	Hypertext Transfer Protocol
ICMP	Internet Control Message Protocol
IETF	Internet Engineering Task Force
IP	Internet protocol
IPPM WG	Internet Protocol Performance Metrics Working Group of the IETF
ISP	Internet service provider
NAP	network access point
NIMI	National Internet Measurement Infrastructure
POP	point-of-presence
PSTN	public switched telephone network
SLA	service level agreement
SNMP	Simple Network Management Protocol
TCP	Transmission Control Protocol

6.0 References

Advanced Network & Services, Inc. (ANS). 1997. Surveyor home page.
<<<http://www.advanced.org/csg-ippm>>>

Almes, G., S. Kalidindi, and M. Zekauskas. 1998. "A One-Way Delay Metric for IPPM."
Internet Engineering Task Force, Network Working Group, Internet Draft.
<<<http://www.ietf.org/internet-drafts/draft-ietf-ippm-delay-03.txt>>>

Automotive Industry Action Group (AIAG). 1997. *ANX Release 1 Draft Document Publication*. TEL 2. Southfield, MI. << <http://www.aiag.org/pub/>>>

Bickel, P., and K. Doksum. 1977. *Mathematical Statistics: Basic Ideas and Selected Topics*. San Francisco: Holden-Day.

Demichelis, C. 1998. "Instantaneous Packet Delay Variation Metric for IPPM." Internet Engineering Task Force, Network Working Group, Internet Draft.
<<<http://www.ietf.cnri.reston.va.us/internet-drafts/draft-ietf-ippm-ipdv-01.txt>>>

Internet Operators (IOPS.ORG). 1998. IOPS.ORG home page. <<<http://www.iops.org/>>>

Mahdavi, J., M. Mathis, and V. Paxson. 1997. "Creating a National Measurement Infrastructure." Presentation at the Internet Statistics and Metrics Analysis Workshop, May 1997. <<<http://www.psc.edu/networking/nimi/slides/welcome.html>>>

Paxson, V., G. Almes, J. Mahdavi, and M. Mathis. 1998. "Framework for IP Performance Metrics." The Internet Society RFC 2330. <<<ftp://ftp.isi.edu/in-notes/rfc2330.txt>>>

President's Commission on Critical Infrastructure Protection (PCCIP). 1997. *Critical Foundations: Protecting America's Infrastructures*.
<<http://www.pccip.gov/report_index.html>>

T1A1.2. 1998a. Alliance for Telecommunications Industry Solutions Network Survivability Performance Working Group. <<http://www.t1.org/t1a1/_a12-hom.htm>>

T1A1.2. 1998b. Technical Subcommittee T1A1 Project Tracking Report.
<<<ftp://ftp.t1.org/pub/t1a1/t1a1.0/8a100081.txt>>>

T1A1.3. 1998. Alliance for Telecommunications Industry Solutions Performance of Digital Networks and Services Working Group. <<http://www.t1.org/t1a1/_a13-hom.htm>>

XIWT

Corporation for National Research Initiatives
1895 Preston White Drive #100 Reston, VA 22091