An overview on the Quality of Service Development Group (QSDG) activity, and the computer telephony integration and the reliability issues

LUIS de SOUSA CARDOSO

QSDG CHAIRMAN
C.P.RADIO MARCONI
Av. Alvaro Pais, 2
1699 LISBON PORTUGAL

ABSTRACT

The past few years witnessed a rapid growth of interest in network reliability. This is because telecommunications services are now an integral part of businesses, national security, and public health and safety. Also, recent technological advances such as fiber optics, high-capacity digital switches and the increasing concentration of capacity in the telecommunications infrastructure have made networks more vulnerable to single failures. Because of this capacity concentration and the increasing reliance on telecommunications, the potential impact of a single outage has never been greater on business, public perception, and geographic scope. On QSDG work, attention has been paid to the network reliability and network outages. In particular, the need for developing a methodology for quantifying the customers' impact of a network outage has been recognised. Network reliability is one of the major future items of the QSDG.

The Quality of Service Development Group (QSDG), which is in nature an operational group including a field trial group, was created in the Study Group 2 organisation in 1984. The original idea came from Mr. C. McCauley from AT&T who felt that, despite their usefulness, the writing of Recommendations as such or drafting a handbook was not enough to improve the quality of service. Telecommunications people need to come together, communicate and help each other, to achieve the best actions for improvements and to have a free exchange of views and ideas in a multilateral environment. The support was overwhelming and ITU Study Group 2 has approved the establishment of this Group mainly because many Administrations and ROA's felt that practical approaches to improving quality of service should be implemented and this could best be done by a Group not burdened by ITU-T (then CCITT) procedures, where "formal agreements" must be reached on Recommendation.

The first meeting of the QSDG was held in Hague, The Netherlands in 1984 and very soon the participation increased as can be seen in the following table

YEAR	LOCATION	CONTRIBUTIONS	PARTICIPANTS	COUNTRIES	COMPANIES
1984	The Hague	23	32	16	20
1985	Orlando	23	29	16	18
1986	Hong-Kong	31	39	21	25
1987	Manila	31	51	19	26
1988	Albufeira	41	54	25	31
1989	Bern	36	50	26	34
1990	Padova	43	57	25	33
1991	Blois	56	71	26	39
1992	S. Francisco	60	88	30	42
1993	Rotorua	65	95	35	45
1994	Chester	68	98	36	50
1995	João Pessoa	51	114	36	54

From the formal point of view, QSDG is directly associated with Study Group 2 Question 8/2. Not only is the QSDG an integral part of the Question with which it is linked, but also in practice the development of Recommendations under the Question follows directly from the experiences shared and contributions submitted at meetings of the QSDG.

The primary aim of the QSDG is to improve the quality of the international telecommunications services, to the mutual benefit of both the customers and the Administrations/Roe's. The QSDG in its Work Program will continue to study all aspects of Quality of Service from the point of view of the customer who uses the telecommunications network. The Group is conscious that in addition to the voice use of the PSTN, there is also a need for it to cover non-voice uses, e.g. fax, ISDN, B-ISDN, packet, video and other services. The work includes general studies on how customer service can be improved, including the interpretation and use of customer surveys, how information can be exchanged between Administrations/ROA's, the effect of network digitisation, appropriate network measurements and their correlation with customer input and multilateral benefits of improving quality of service and network performance.

Considerable work has been accomplished, and several achievements has been reached. An example is the "IDD Completion Rates World-wide" database which is updated yearly and used extensively by the participants. From this database, it has been concluded that the number of destinations which could be considered as offering poor call completion is decreasing. Following the criteria defined in the Recommendation E.426, the following table provides insight to the ASR (Answer Seizure Ratio) trends

ASR TIER	1991	1992	1993	1994
> 60%	9.8%	13.5%	14.6%	18.9%
> 30% < 60%	61.4%	63.4%	67.1%	61.6%
< 30%	28.8%	27.7%	18.3%	19.5%

Considering the information in this database and taking into account that QSDG has met 6 times in Europe, 3 times in Asia-Pacific Region, twice in USA, and once in South America, it was decided that future meetings should be held in regions of the world where QSDG had not met before. It has been felt that if QSDG meeting took place in these areas, the number of delegations participating would increase, and more people would be involved in Quality matters. Consequently it is expected that the quality and performance of some destinations in these areas will be improved. This objective has been achieved and the next meetings of QSDG are planned to be held in Africa, Middle East and in Oceania by the kind invitations of South Africa, Bahrain and Fiji Islands, respectively.

It is anticipated that as long as the participating entities are benefiting from it, the QSDG will continue to exist and will meet around the world regularly. All costs involved are borne by the

Administrations/- themselves, and this implies that for the ITU-TS there is no significant costs involved, other than the publication of the Report of the Meetings as a White Contribution.

Currently the QSDG is Chaired by Mr. Luis Sousa Cardoso from CPRM-MARCONI, Portugal, who was appointed to this position at the Rotorua meeting in 1993. It must be noted that all information and work presented at the QSDG meetings are in the public domain and can be obtained through the Chairman.

◆*The QSDG and the Computer Telephony Integration*◆

Things used to be simple. There was only one type of light bulb. Gasoline was all leaded. Mustard was only yellow. Now incandescent lights are being replaced with fluorescent and halogens, gasoline has at least three octane rating, and mustard fills three shelves at the supermarket. It's the same with telephone and computers. When they didn't need to talk to each other much, a modem was more than sufficient. But now, the business advantages of computer-based call control are forcing its natural bond.

No one denies that computing and telecommunications are converging, but the vision of the converged world has changed. Until recently, it was still possible to imagine that the future would look pretty much like today's telecommunications environment, but with fancier tricks, like video built into the familiar telephone, and maybe even telephone companies selling combined PC-phones. It is now becoming clear, however, that the reality will consist of more powerful and flexible computers that are capable of taking over most, or maybe even all, of the traditional telecommunications functions. Convergence increasingly means a computing take-over. This may still be controversial among some of the larger private branch exchange (PBX) vendors and telecommunications service providers, but confidence in the computing camp is growing to such an extent that some are even handing out diagrams showing the information technology business environment of the near future, and there is no room in the diagrams for a PBX.

Today's CTI (computer telephony integration) systems generally fall into one of four different architectures or configurations, based on their approach to making the actual connections and making calls:

- Phone-centric systems
- Server-centric systems
- Voice-server systems
- PC-centric systems

Phone-centric systems are the easiest to implement; they only require a direct link from the phone to an external adapter that connects to the PC's serial or parallel port. They don't require extensive changes to an existing phone systems. Users can have direct control over call routing (known as first-part call control). To transfer a call, for example, the user just clicks on an icon, and the PC sends a message to the switch that emulates a command from the phone requesting the switch to transfer the call. Many PBX vendors offer adapters that give that kind of control to the PC. Unfortunately, these adapters don't provide a connection to the phone line and can't be used to connect data or fax lines to the PC.

Server-centric systems connect your telephone switch to a server on your LAN. Here, the phone system becomes another part of the computer network, and you don't need a physical connection between the phone and individual desktop PCs. The LAN server, not the switch or the user, is responsible for routing calls (thus, it's termed third-party call control). To transfer a call, the user clicks on the transfer icon, which sends a message to the server requesting that it transfer the call. The third part (the server) sends a message telling the switch where to route the call. The server's processing power lets it screen and route incoming calls. For example, caller-ID information may help route the call to the proper person. Third-party call control is particularly helpful in workgroups and call centers. But the server-centric model manages only call control. The switch-to-server link is for status and request only. It doesn't carry the voice path and in no way physically connects a phone line to the server. For a server to send and receive faxes and data, a physical phone line would have to be connected to a fax-modem board in the server.

Voice-server systems are a variation on the service-centric model. Where server-centric systems deliver call-control links but not the calls themselves, voice-server systems deliver the call directly, but not a separate control-and-status link. In a voice-server model, phone lines from the switch connect to a board in the voice server. Depending on the board's capabilities, the lines can be analogue, ISDN, or proprietary digital. The board can do anything that a phone it replaces can do; for example, it can issue a flash hook to transfer, conference, hold, call park, call forward, initiate call pickup from another office, and so forth. Digital phones usually have other features, such as speaker-phone control and caller-ID display. With the phone line going into a voice server, you get the media - that is, the voice path, or data path for faxes and modem calls - but you don't get all the information and control that's available on the server-centric model. For example, a phone line can't force the switch to take control of another call; it can only control a call that it has received or placed. It can't tell the switch to forward a call from the next office to another phone.
In a server-centric call-center application, the server receives the caller-ID and tells the switch where to send the call. In the voice-server model, on the other hand, the call is sent to the server, which must then answer it and transfer the call. But this is just too slow for a call center.

PC-centric systems have the telephone line and the telephone itself connected directly to an add-in board in the PC. The telephony board emulates the type of telephone that the switch is designed to support, whether analogue or proprietary digital. When we have isochronous Ethernet or ATM (asynchronous transfer mode) data pipes going directly into our PCs. which looks to be the long-term prospect, we'll use PC-centric telephony systems. For the short term, however, we'll see fax-modem boards with telephony features that will provide an interim solution.

Just as a crisis hit the mainframe computing world when PCs first burst onto the scene, PBXs are seen as the inevitable victims of smarter PCs that can understand human speech, play real-time video and still have enough spare random memory to run telecommunications functions in software. Equip such a PC with some slot-in cards that can send and receive telephone signals, and computing has effectively swallowed telephony.

There will still be space for the PBX vendors, just as there is still, even today, a place for mainframes in the corporate computing world. In the office backbone network, and in the wide area networks that links corporate offices, convergence is taking the form of a new communications protocol that can handle data, voice and video with more or less equal ease. Called asynchronous transfer mode (ATM), this protocol promises to be the basis of new kinds of networks and may eventually - though this still controversial - extend all the way from the international networks to the desktop-based local area networks.

ATM is still fertile ground for the PBX vendors, whose detailed under-environment makes them prime partners for the computing and data networking vendors. The larger PBX vendors are also manufacturers of public telephone exchanges, and they have been among the front-runners in developing ATM technology .

The requirements of the public networks in the converged world are difficult to foresee, however. A lot will depend on what happens at the individual customer sites, both in home and in the office. Much as the public telecommunications carriers would like to forecast the demand in the next decade, or even to shape that demand to suit themselves, the chop-and-change nature of the computing world is bound to make that wish impossible to realise. What is hardest to predict is the extent to which computer power will infiltrate the public exchanges sites themselves. It is already possible to build a public telephone exchange out of standard computer components. Some believe the resulting system cannot possibly be resilient enough to provide the service availability that telecommunications customers expect, but the reality is that computing is becoming increasingly more reliable. Indeed, the problem of 100 percent availability has already been solved, time and time again, by computer systems designed for banks, finance

FOUR CTI ARCHITECTURES

PHONE-CENTRIC SYSTEM

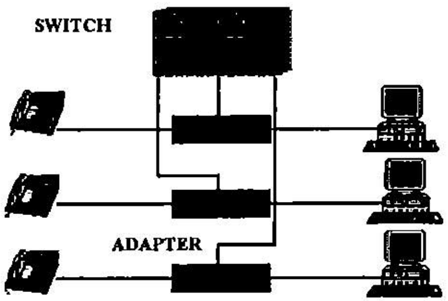

SERVER-CENTRIC SYSTEM

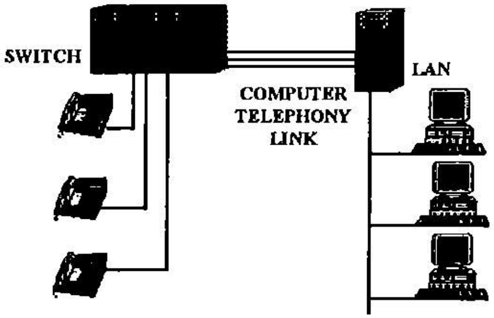

VOICE-SERVER SYSTEM

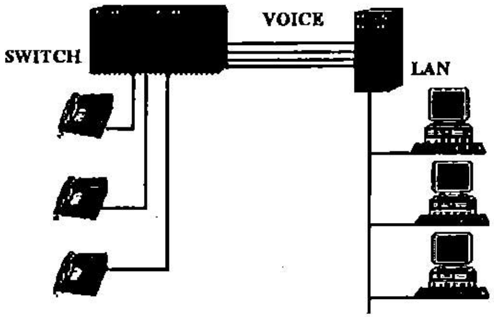

PC-CENTRIC SYSTEM

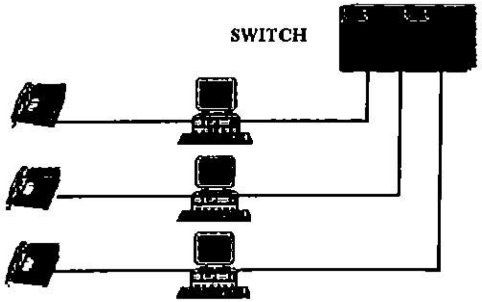

companies and the military. It used to be that all telecom systems required hardware and software more sophisticated that what was on offer for standard business computing solutions. This is increasingly no longer the case. Standard computers are now powerful enough to handle the human voice and full-motion video in real time. The operating systems running on today's servers are sophisticated enough to cope with the logic required of a digital telecommunications switch.

In fact, the insides of the telecommunication boxes, which used to look sophisticated compared with corporate computer systems, now look out of date. PBXs and telephone switches actually contain their own computers, operating systems and applications software. Typically, the computers are non-standard, the operating systems are proprietary, and applications are not flexible.

There is no real need for all that computing to be inside the telecommunications box at all. Transferring the guts of the PBX to a UNIX or Windows NT server suddenly opens the market up to price competitive hardware, to operating systems that integrate properly with corporate information technology and to software that can be modified by the customer. In the long term, the same will probably be true of the public telephone exchange.

What this means for the world-wide compatibility of systems and services is an important question. We are growing accustomed to being able to pick up a telephone in one part of the world and get through to any other telephone on earth. The more complicated and varied the systems get - and more they come to resemble computers - the harder it will be to maintain its uniformity.

In 10 years, convergence will have changed the look and feel of telecommunications and computer technology, and also the shape of the information technology industry. Just as the advent of distributed systems revolutionised the industry in the 1980s, convergence will do the same between now and the millennium. The previous revolution forced some companies to merge, especially the larger system providers, while some once obscure companies rose to international prominence. The convergence revolution is bound to make a few reputation, but it may also break a few.

♦ *The OSDG and the reliability issue* ♦

The use of International Telecommunications services is dependent upon the interconnection of numerous items and media which are controlled, often independently, by various entities and agencies. This is increasingly true also for national telecommunications services as deregulation proliferates. Whilst individual administrations, agencies, or network operator may be subject to desirable Quality of Service measurements and improvement programs, it is the overall (end-to-end) performance which customers perceive as important.

When credit-card, long-distance telephone, and retail banking services users are asked to rate (on a scale 1 to 10) the importance of the different dimensions of the Quality of Service, reliability clearly emerges as the most important dimension, regardless of the service being studied. A typical example can be seen in the next table.

Service Quality Dimension	Credit Card		Long-Distance Telephone		Bank Services	
	Average Rate	Most important	Average Rate	Most important	Average Rate	Most important
Tangibles	7.43	0.6	7.14	0.6	8.56	1.1
Reliability	9.45	48.6	9.67	60.6	9.44	42.1
Responsiveness	9.37	19.8	9.57	16.0	9.34	18.0
Assurance	9.25	17.5	9.29	12.6	9.18	13.6
Empathy	9.09	13.6	9.25	10.3	9.30	25.1

The customer's message to the service provider is clear: Be responsive, be reassuring, be empathetic, and most of all be reliable. Complementary there is another message: Human performance plays a major role in customers' perception of service quality. Three of five dimensions - responsiveness, assurance, and empathy - result directly from human performance. Moreover, reliability often depends largely on human performance. Concerning this particular point, it has been felt that QSDG, in their working methodology, is playing a significant contribute.

The reliability performance of a system is its ability to perform a require function under given conditions for a given interval. Important factors influencing reliability performance are:

The quality of the constituent components
Environment factors
The network structure
Security against secondary failures

Telecommunications networks are vulnerable to many threats: natural disasters, intentional sabotage, accidents, and faulty hardware and software. But technical solutions exist and can be deployed to minimise both failures in network elements and the impact of disruptions on telecommunications infrastructures, thereby meeting the requirements of telecommunications users.

Telecommunications services have traditionally been characterised by high level of service dependability. Dependability is a characteristic that includes reliability, availability, maintainability and survivability. In briefly, dependability means that a service works the way a customer wants it to work when the customer wants it. Three different considerations are making service dependability a topic of great interest in the industry today. First, as the use of telecommunications services has become a more integral part of many businesses, customer expectations for dependability have increased. Second, revolutionary changes are being made or proposed in the network. While each of these changes brings with it the opportunity to provide customers with additional services that meet their needs in a timely fashion, each also provides a challenge in assuring the continued dependability of new services. Third, some unfortunate events in the recent past have disrupted services for a large number of customers for an extended period of time and have caused concern among customers.

It is recognised that a three-pronged attack is needed to meet the challenge at hand:

- Making the network elements less vulnerable to hardware failures, software errors, and procedural errors, and therefore minimising network elements outages.

- Making the network less vulnerable to network element outages, for example, through use protection switching, diverse routing for protection, self-healing rings, dual homing, and other means of minimising network outages or the impact of the outages.

- Improving and automating the techniques used to restore service after a network outage.

Reliable network elements, survivable network architectures, and efficient restoration strategies will be the keys to assuring dependable telecommunications services in the 1990s.

Networks have traditionally not been designed to provide 100% service continuity in face of catastrophic failures because the cost was thought to be too large. However, with the increasing dependence of customers on telecommunications for their business survival, customers expect a

high level of service restoration after a catastrophic failure. More and more, customers require and demand service approaching 100% availability on an end-to-end basis. Even if customers settle for less than 100% availability during a failure, their minimal requirements may be 100% availability for critical services, with some minimum level of performance for all services to all locations. Network providers must balance meeting requirements of specific customer and service with providing a high-reliability network with service assurance to all customers.
The ultimate reliability goal is to make all failures imperceptible to users. An interim goal may be to reduce the impact of a failure so that calls in progress are not cut-off and data sessions are not prematurely terminated.

The past few years witnessed a rapid growth of interest in network reliability. This is because telecommunications services are now an integral part of businesses, national security, and public health and safety. In addition because telecommunications services have traditionally been so reliable, even through national disasters, public expectations are very high. Also, recent technological advances such as fiber optics, high-capacity digital switches and the increasing concentration of capacity in the telecommunications infrastructure have made networks more vulnerable to single failures. Because of this capacity concentration and the increasing reliance on telecommunications, the potential impact of a single outage has never been greater on business, public perception, and geographic scope.

Therefore, there has been a concerted effort by end-users, service providers, standards bodies, and government and regulatory agencies to undertake multiple activities addressing issues related to network reliability and service outages, over the past few years.

This will lead to a scientific assessment and enhancement of network reliability. A major question is how to quantify a network outage.

For years, network reliability planning has concentrated on establishing the reliability of each network component. The individual component reliability are then aggregated to establish, through modelling, a target network reliability. This resulting reliability was then examined for adequacy.

How does a network planner determine what is adequate? The current view of quality suggest the planning start with customers expectations and needs for their services. Those service reliability needs can then dictate the demands on the network, and in turn the individual network component requirements.

With this current methodology, then, the understanding of reliability levels delivered to customers must also begin with an end-to-end network assessment of what quality of reliability is being delivered to the customer. Thus the process now begins with the customer.

Reliability as defined by AGREE (1957) and ANSI/ASQC is the probability that an item will perform without failure a required function under stated conditions for a stated period of time.

To define network reliability, QSDG has accepted to describe the six elements of the reliability definition for a telecommunication network. In network applications the six elements of the reliability definition are:

⇒ Probability - Network reliability is quantified in terms of probability.

⇒ Item - The item to be considered here is a telecommunication network.
However, a variety of distinct, but interconnected, networks exist. These networks could be owned and operated by different companies. For instance, there are multiple Local Exchange Carriers (LECs) and Interexchange Carriers (ICs). A single end-to-end usage of the network may utilise several of these distinct interconnected networks. In network reliability studies the network under investigation needs to be precisely defined and its boundaries specified.

⇒Required function - A telecommunication network is required to perform several functions. From the users' perspective and the network reliability point of view, the first and the most common function of a network, is the ability to communicate from a source to a target. This is the ability to establish a new connection and/or maintain the established connection.

⇒Failure - A network fails to perform the required function, from a source to a target, when the failure probability exceeds a pre-specified failure threshold value (the maximum designed value). The failure threshold and/or customer impacting values need to be determined for different networks, applications, and services.

⇒Stated conditions - A telecommunication network consists of many distinct elements. Networks are designed to perform the connection establishment function under certain conditions. These conditions depend on the network, service, and application. QSDG is considering that a failure due to CPE (Customer Premises Equipment) should not be counted as a network failure.

⇒Stated Period of Time - A time period must be specified in order for a probability measure to be meaningful. It is desirable to know what is the probability a network will perform the connection establishment function during the next time interval.

It is seen, therefore, that a statement about network reliability requires an explicit, clearly defined and formulated description of the network under study, the network function, network failure, and allowed environment and conditions.

Finally, a network outage needs to be clearly defined in terms of customer perceivable and measurable quantities. The customer's perspective of a network outage can be classified into three major components of Intensity, Duration, and Extent. The (I,D,E)-triple, provides a framework for measuring network outage. This framework can be applied to different networks and services. Each network and/or service needs its own specific formulae and computational procedures. Quantitative scales measuring end-to-end customers' impact of an outage may be developed through appropriate integration of I, D, and E, as well as customers' related parameters.

The Federal Communication Commission (FCC) of the USA issued a report (February 27, 1992) defining "service outage" to constitute loss of service to more than 50000 customers (later reduced to 30000) for a period of 30 or more minutes.

Many international organisations, fearful of a damaging loss of customer confidence, refuse publicly to admit ever having suffered a major network outage, and are reticent even to acknowledge the possibility. New Jersey-based Ascom Timeplex, is one of the few companies to have a successfully commissioned independent research. Its survey of 100 of the top companies in the UK shows that 50% of those interviewed predict a loss of business if a communications network is down for more than an hour. Almost nine out of ten firms (87%) said a network failure of less than 60 minutes would have a measurable impact on efficiency and throughput. A similar survey indicates that 73% of respondents experienced LAN failures, 23% of which lasted for more than 10 hours. Some companies logged as many as 20 outages a year. Failure was particularly high on networks supporting between 100 and 299 users, with router and bridge interconnection technology being the most likely to go down. Furthermore, a high degree of the recurrence of specific problems was reported. Token ring networks apparently failed more often than Ethernet, but disabled fewer users when it was out of action. Research specific to the US only shows that a typical LAN there is likely to be disabled twice as often for an average of five hours. These findings are bolstered by a European-wide survey of 100 top IT professionals, conducted by the Yankee Group. One financial institution is on record as saying that it could easily lose US$10000 in direct costs for every minute that its network is down. Indirect costs, such as loss of credibility and market share, breach of statutory requirements, cashflow

problems and collapse of share price make the overall expense of major outage, literally, incalculable.

Within QSDG, work on network reliability started on 1993. Network reliability is one of the major future work items of the QSDG. In particular, the need for developing a methodology for quantifying the customers' impact of a network outage has been recognised by regulatory agencies, industry leaders and end-users. A network reliability policy begins with identification and definition of quantitative scales to measure an outage and assess its impact on end-users.

REFERENCES

Contribution QSDG - 12/95-41 - Reliability of the US Telecommunication Services
M. Daneshmand, C. Savolaine - AT&T

Contribution QSDG - 10/93-18 - A Framework for Measuring Telecommunications Switched
Network Outages — M. Daneshmand, C. Savolaine - AT&T

"What sort of convergence is coming? " - Bob Whitehouse -International Herald Tribune, October 9, 1995.

"Standard issue" - James Burton - CEO of C-T LINK, Inc, Boston, Massachussets, September 1995