



Managing Quality of Service

Strategic White Paper

Managing Quality of Service

A Strategic White Paper

March 1, 1997

by Richard C. Sturm,
Principal
Enterprise Management Professional Services, Inc.

Table of Contents

I. Introduction	1
II. Divergent Views	3
III. Defining Quality of Service	5
A. Availability	5
B. Performance	6
C. Accuracy	6
D. Affordable	6

I. Introduction

Over the last 10-20 years, as global competition has increased, managers have increasingly sought ways to improve their organizations' competitiveness. Often, as a new idea appeared, it was adopted rapidly, and some would say unquestioningly. All too often, these new ideas have ended up merely being short-lived fads that have produced few tangible benefits. Many of these ideas have been held up to ridicule in the popular cartoon, "Dilbert."

Managers of Information Technology organizations (IT) have proven to be susceptible to this same phenomenon. The frequent misuse of popular technology has been rampant. Currently there appears to be the beginning of a groundswell of interest among IT managers in quality of service management (or as some refer to it as service level management). The question must be asked. "Is quality of service management another management fad, or is it a legitimate business focus?" Many cynics would be quick to say "yes." However, a review of today's situation quickly refutes such skepticism.

Traditionally, IT managers have measured the effectiveness of their organization by looking at the various hardware and software components for which they are responsible. The units of measure that were applied have been varied. The most common have been such things as availability, utilization, performance, etc. This has given IT a fragmented, misleading perspective of its effectiveness. It is a perspective that is distorted and misleading. It has led IT managers to believe that they have been doing an excellent job at meeting the needs of their clients for IT services.

In recent years, while IT managers have been congratulating each other on the ever improving quality of service that they have been providing, their clients have had a radically different perspective. They have been complaining about the inadequacy of the service and the unresponsiveness of the IT organization to client requirements. In recent years clients of the IT organization have become more and more dependent upon the services provided by IT. The very existence of many businesses depends upon the delivery of those services. In addition, their clients (users) have become more technically sophisticated. This has given them a better understanding of what they can reasonably expect IT to deliver. Together, these factors have led IT's clients to demand a much higher quality of service from IT. In some cases, it has also led the users to stage a coup, toppling the IT organization from its pedestal and replacing it with an out-sourcer's organization. The clients of IT hold the corporate power and control the corporate purse. One way or another, they are determined to receive the level of IT services that they require.

Today's interest in managing the IT organization in terms of the quality of service it delivers has come to the forefront due to the convergence of two major factors. Basically it is the convergence of supply and demand. The confidence of IT managers has been shaken by the complaints of disgruntled clients and also by continuing reports of outsourcing deals. IT, and their clients, have a need to accurately measure and report on the quality of the service being delivered by IT. This is the demand side of the equation. True end to end service level measurement has only recently become practical. There are now tools emerging to support the measurement of the quality of the service being delivered by IT.

It must be remembered that not only is IT a service provider, it is also a major consumer of services. If those providing services to IT do not deliver a consistent level of service, IT will not be able to be consistent in the quality of services that it delivers to its clients. Those providing services to IT include groups internal to the company (e.g., data base administrators, application developers, etc.) as well as some that are external (e.g., telephone company). However, it must be

realized that managing the quality of service delivered is not limited just to IT organizations. Certainly, for most organizations the IT function is a critical service provider, but there are many others. For example, the telephone company that provides voice and data links to other companies, is a major service provider. It is only reasonable to demand a guarantee of the quality of service being provided.

Is the current desire to establish effective measurement and reporting of the quality service being delivered a fad? Absolutely not! As will be seen in this document, quality of service management (service level management) is absolutely essential as a survival strategy for the IT organization.

II. Divergent Views

What has gone wrong? Why have the views of IT and their clients about the quality of service being provided by IT become so divergent? The bottom line answer is that the two groups do not have a common understanding of what is an adequate level of service or even about what the key indicators are of the level of service being provided.

Until now, most IT organizations have not been measuring the quality of the service being delivered to their clients. Instead of measuring the quality of the service being delivered, IT has been measuring various attributes of the hardware and software used to deliver the service to its clients. Certainly this data is valuable for evaluating the performance of the IT employees or planning for adequate equipment capacity in the future. Very technical metrics are also essential in assessing the quality of the service being delivered by the external network service providers (i.e., telephone companies). However, this very technical data is not indicative of the level of service being delivered by IT to its clients.

There have been several flaws in the strategy employed by IT. First, it has not been possible to capture equivalent information for every component. That is, in examining the end-to-end quality of service being delivered, there are numerous components that impact that service. There are various sources of information about those components and about some portion of the service being delivered. These different data sources do not capture the same types of information, nor is their data capture synchronized. The result is that the data available has been fragmented, scattered and disjointed.

Second, for some components, only limited data has been available. In some cases, no data relevant to quality of service has been available. The third flaw has been the type of information captured. Certainly, CPU uptime, dropped packets, network congestion, etc. will each impact the quality of the service that is delivered. However, these factors can only indirectly suggest something about the quality of service being delivered. Finally, IT has been struggling under the assumption that the whole (of quality of service management) is equal to the sum of all of the parts. In mathematics, this assumption is valid. In the management of IT services, it is totally false. Consider an simple example in which a typical transaction requires that each of the following pieces of equipment and software listed be available.

Table 1

Component	Average Availability	Minutes of Downtime
LAN	99.97%	0.32
Local Server	99.95%	0.54
Building Hub	99.96%	0.43
Intranet router	99.99%	11.88
Remote Host	99.99%	1.08
Order Entry applic.	99.91%	9.72
Customer data base	99.92%	8.64
Inventory data base	99.96%	4.32
Average/Total	99.66%	36.94

Let's assume that the user group performing order entry demands 99.90% availability for the order entry "function." (Note, the term function is being used to indicate the desired activity, such as entering a new order.) Expressed another way, the client expects to be able to have the function available 24 hours per day, 7 days per week, with no more than 10 minutes of downtime in a week. Looking at the table above, most IT organizations will congratulate themselves on having met the client's requirements in all but one of the categories, and even that was very close to the objective. However, this is probably not correct. Unless there was the coincidence of all of the problems occurring at the same time, then the total time that the function was unavailable to the order entry department was much greater than the acceptable level. If the outages each occurred at a different time, then the total amount of time that the function will be unavailable to the client was almost 37 minutes! Also, the availability can be calculated by multiplying together the availability of each of the components. The result is 99.66% - far less than the target of 99.99%. In truth, there may be some overlap between the various events. This could reduce the total amount of time that the function was not available to the order entry clerks. However, overlaps could not conceivably reduce the minutes of outage to the target level.

From the perspective of the end user of the IT services, the availability of individual components simply does not reflect the quality of the service being provided. The one merit of this particular metric is that it does reflect a concept that the clients can relate to, even if it is from a different perspective. If data for individual components is all that is available, then what is needed is a more sophisticated correlation and analysis of the data to translate it into meaningful information.

In terms of managing the quality of service to client organizations (that is, end users), there is absolutely no value in reporting such arcane data points as packets dropped, network latency, repair interval for network outages, etc. With a few rare exceptions, the clients simply do not understand their significance. Even if understood, these pieces of data are not measurements of the quality of the service being provided by the IT organization. It is true that such information may be valuable measurements for internal IT assessments. In fact, some groups develop what are termed "Technical Service Level Agreements." These are basically agreements that are used to look at how well the groups within IT are fulfilling their responsibilities.

Very well, technical measurements are valuable for special purpose assessments of quality, but in their raw form do not reflect the quality of service that is being received by end users. However, if device availability, as well as other detailed statistics, are not valid indicators of the quality of service being delivered to end users, what is? Before this question can be answered it is necessary to deal with the even more fundamental issue of defining what is meant by the term "quality of service."

III. Defining Quality of Service

To begin with, in considering the meaning of the term “quality of service,” it is necessary to assume the perspective of the client - the consumer of the service. The client’s perspective is the only perspective that matters in this definition. This is true regardless of the industry. It is as valid in the fast food business as it is in the world of IT. The client is the one utilizing the service. The client’s perception of the quality of the service being delivered will determine whether or not they are satisfied with that service and whether or not they will want to continue to work with that service provider.

There may be a temptation to consider the adoption of a client’s perspective to be biased and possibly even one-sided. However, IT managers must remember that they are both providers and consumers of services. They are the client when dealing with service providers such as the local telephone company. They are also consumers of services from their own organization. The IT department uses the services of various groups within IT to provide the total package of services that they offer to a client. For example, a typical IT department will have a group responsible for operating the corporate Intranet. Another group may be charged with the operation of the cen-

There are also less obvious aspects of availability. Consider the case in which a company has a system that allows sales representatives to dial in using their laptop computer to check product availability while they are at a customer's office. In such a case, if a modem port is not available when a sales representative attempts to call, then for that sales representative the service is not available.

B. Performance

Words can hold vastly different meanings for different people. This is certainly true of the word performance. For IT personnel performance will suggest such things as packets per second. However, for the client, the meaning is simpler - do the IT services function at an acceptable speed? The question of speed may be the question of response time in an online transaction system. In another case it might be the time that it takes to move a copy of a file from one office to another, or the time required to load an application to a desktop system from a server.

C. Accuracy

From the perspective of the user of a service, there is a third part to the question of quality. That third part is accuracy. This component is concerned with the question of whether or not the service performs accurately. An example of this is the question of whether or not e-mail is delivered to the correct recipient. Similarly, in the case of applying transactions to a data base, it is essential that the change is applied to the proper version of the data base. If services do not accurately perform their functions, high availability and high performance are worthless.

Having loosely defined quality of service, let us next look at the question of how to approach the management of that service and the level of the service that is provided. We will defer the question of which specific indicators or measurements are appropriate to use until later in this document. First we must look at how they will be used.

D. Affordable

While availability, performance and accuracy, are strictly related to the quality of service, the cost of the service cannot be ignored. There is a saying in application development, "Fast, cheap, good - you can have two out of three." This expression is true when it comes to delivering any service. There is always a tradeoff between cost of the service and the quality of the finished product, as well as the timeliness of the delivery. It may be possible for a service provider to raise the quality of the service that it is providing dramatically. However, the cost of such an improvement might be so great, that it, if implemented, that it would cause the bankruptcy of the corporation. Therefore, while cost is not directly part of the quality of the service being delivered, it is a limiting factor, and cannot be ignored.

IV. Approaching the Problem

Managing the level of service that is provided is a little bit like the physics problem of creating a vacuum. Assume that you have a vacuum pump attached to a sealed chamber and in a cycle lasting one hour, the pump is able to remove 50% of the air in the chamber. It is almost certain will never get to a total vacuum (that is, no air molecules left in the chamber), at least not in a reasonable amount of time. Likewise, service providers and their clients must accept the fact that they will never achieve perfect service, at least not at a cost that any business can afford. Perfect service, that is always available with excellent performance and 100% accuracy, is simply not possible.

If perfect service is not possible, then it is necessary to determine what level of imperfection is acceptable. This determination is not something that can be done by the service provider alone, nor is it possible for the client to do this task. The service provider lacks some of the necessary information. The service provider does not have adequate information about the value of the service to the client organization. The clients do understand their own need for the service and its value to the organization. However, the client does not have a good understanding of the issues facing the service provider. The client may have enough technical expertise to realize that it is possible to substantially increase availability, but they will probably not know the financial and other implications of doing so.

The answer for determining the appropriate service level lies in approaching the problem as a partnership between the service provider and the client organization. (It must be remembered that although IT is most often thought of as the service provider, they are also major clients for many external organizations.) The ideal is for the two groups to come together to work out what could be best for the company, in a “win-win” approach to the problem.

Unfortunately, not all groups or individuals are willing or able to take such an approach. Some insist upon a win-lose approach to these discussions. Others come to the discussions with absolutely no ability or willingness to make concessions or adjustments of any kind. This latter case is best seen when an IT manager attempts to negotiate a service level agreement with the local telephone company. The IT manager might as well try to negotiate with the government’s tax collectors! Even if one party is not a willing participant, it is important that the dialog takes place.

When discussions take place, everyone should be prepared to share as much factual information as possible about the costs and impacts of various service levels. This information can then make it possible to drive decisions based upon business impacts, rather than merely emotions.

What is the expected result from the dialogue regarding the quality of the service to be provided? Very simply, the dialogue should produce an understanding of what is desirable, possible and acceptable to all parties.

It must be remembered that when IT provides a service to a client or uses the services of an outside provider, it is still entering into a business relationship. It is easier for most people to recognize the existence of a formal business relationship when dealing with an external supplier. However, the business relationship is just as real when it is between two internal entities (i.e., IT and a client group).

Whenever a business relationship is established, expectations are set. In the case of internal IT services, the client expects that service will be delivered at an acceptable level of quality. IT funding may not directly depend upon the level of quality of the services that it provides. However, IT certainly has the right to expect that if it delivers the services requested at an agreed upon level of quality, the client will be satisfied with those services.

Time clouds even the very best of human memories. Any understanding that is reached needs to be documented. This requirement is not limited to just the issue of quality of service. It applies to nearly every business situation. Documenting agreements and understandings is simply a prudent business practice.

V. Documenting the Understanding

A. Need for documentation

When a company hires an external supplier to provide goods or services, the terms of that relationship are usually documented in a contract. It would be unthinkable to enter into a multi-million dollar business deal without a contract defining prices, payment terms, responsibilities, deliverables and delivery dates. Yet that, in effect, is what IT organizations and their clients are doing every day.

Contracts are a valuable tool for businesses and they should be applied to internal commitments, as well as to ones with outside suppliers. They can be a vehicle for avoiding misunderstandings and also for ensuring that there is a common, clear understanding of what is expected. In the absence of some documentation of the service commitments, problems will arise.

The first problem that can arise is a lack of a common understanding. Both groups have agreed to the same set of words. Yet, in the absence of documentation, those common terms can carry vastly different meanings to the two groups. The result is that they can go away from the discussions with significantly different interpretations of the level of service that will be provided.

Let's consider the hypothetical case of an IT organization as the service provider and a department in the same company is the client. At the beginning of the year, representatives of the two organizations sit down together and discuss the client's service requirements and the ability of the IT organization to deliver those services. After much discussion, the representatives agree upon the level of service to be provided by IT during the coming year. Being typical of many organizations, the service levels are ill-defined and the agreements are not documented.

Assume that the IT organization performs superbly and during the first quarter of the year, every one of the service level commitments are met. What will be the result? Will the client organization's management host a dinner to honor the IT organization? Never! Instead, in the absence of written documentation, IT will become a victim of "expectation creep." That is, in spite of the agreement reached at the beginning of the year, the clients will start to wish for even better service. Soon, the client will have (unconsciously) translated those wishes into expectations. Therefore, during the remainder of the year, even if IT continues to meet all of the commitments that it made at the beginning of the year, in the client's opinion, IT will have failed. The failure will be due to the fact that IT has not met the higher expectations that have developed on the part of the client. It should be remembered that in this example, IT has never agreed to the higher service levels. Yet in spite of this, the client will be dissatisfied with the service provided by IT.

B. Creation process

The documentation for services is a contract. If the service provider is an independent company, the agreement will probably be incorporated as part of the contract between the two companies. However, if the service provider is part of the same company as the client, the agreement will probably be less formal. Regardless of the legal form of the documentation, the information contained in it will be essentially the same, as will be the process for creating it.

The essence of the process for documenting the agreement, is one of careful negotiations. There are several points that should be remembered when entering into the negotiations. First, the two parties must be entering the negotiations as equals. That is, negotiations are not possible when

there is a great difference in the authority or power between the two parties. When a significant imbalance exists in a negotiation situation, it is possible for it to degenerate to the point where the more powerful group dictates the terms of the agreement and the other party acquiesces. Certainly there will be situations in which the service provider is at a significant disadvantage in the negotiations in terms of power or leverage. However, for the discussions to produce meaningful results, the group with the advantage, must be willing to resist the temptation to wield that power during the negotiations. If the negotiations degenerate to the point that one party is dictating terms to the other, then the resulting agreement is of dubious value. The one-sided process may be an agreement that is binding. In the case of an external service provider, the agreement may even be legally enforceable. However, that does not ensure that the service provider will be able to meet the terms of the agreement. On the other hand, it does ensure that the client group will have the expectation of the commitments being met, even if that is not possible.

The next point that must be remembered when beginning the negotiation process is that the objective should be what is termed a “win/win” result. That is, both parties must be able to feel that the resulting agreement is fair and reasonable. In order to achieve this, both groups must be flexible in their positions in the negotiations. One characteristic of the successful negotiator is the ability and willingness to understand and respect the objectives, requirements and limitations of the other party in the negotiations.

A third requirement for successful negotiations relates to the individuals who actually engage in the negotiations. The negotiators need to have sufficient authority to represent their organization in the discussions. Some managers like to use the negotiating ploy of sending a representative who does not actually have the authority to make a commitment and conclude the negotiations. This is an unethical tactic and should be avoided.

In addition to having sufficient authority, the negotiators need to understand their organization's issues and objectives in the negotiations. There is no value in trying to negotiate an agreement using personal opinion and speculation. Obviously, the client organization's negotiator must know what levels of service the organization desires. Perhaps less obvious is the fact that the negotiator must know the value of a given service level versus a different level of service. On the other side of the table, the service provider's negotiator must know whether it is possible to provide a desired level of service and also what will be the incremental cost of providing that level of service. Essentially, the heart of the negotiations should be based upon fact, and as much as possible, the discussions should be driven by a cost benefit analysis. In the case of negotiations involving an external service provider, the client will not normally be as candid with financial data as might be the case with an in-house supplier. Even if the data is not shared between an external service provider and the client, it is important for their negotiators to have the information, to provide an objective, factual basis for their bargaining.

C. Documentation characteristics

When an agreement is reached it is important that it contain certain fundamental pieces of information and that the agreement have certain characteristics. Those key characteristics are:

- Attainable
- Measurable
- Meaningful
- Controllable

- Mutually acceptable
- Cost Effective

1. Attainable—When an agreement is arrived at for the level of service to be provided, one of the most fundamental requirements is that the objectives set forth in the agreement must be attainable. There is no value in setting objectives for the level of service, that the service provider knows cannot be met. For example, an agreement might call for no packets to be dropped in a corporate Intranet. While theoretically possible, over any reasonable period of time, this objective could not be met. Other examples might be found at the Help Desk. An agreement might specify that technical support would be available 24 hours per day, 7 days per week. Yet, there could be a union contract that specifies that the Help Desk personnel are not permitted to work on Sunday.

The items above are examples of terms in the agreement that would be impossible to meet. There are also conditions that might be specified in the agreement that would be theoretically possible to meet, but impractical or unaffordable. Such items are as irrelevant and inappropriate to a good agreement as are the ones that are completely impossible. Consider the case where a client insists upon having extremely high availability for key applications. In order for IT to provide this it will be necessary to acquire \$1 million in additional hardware and software. If IT cannot obtain the funding necessary to make those acquisitions, then this item in the agreement has to be considered impossible. That, in turn, makes it irrelevant for any agreement regarding the service to be provided by IT.

It is reasonable to ask why any service provider would ever make a commitment to provide a level of service if it was not possible to meet that commitment. First, it should be noted that this normally happens when the service provider is an internal group, rather than an outside provider. Such a commitment can be the result of the negotiator for the service provider not realizing that the objective could not be met. However, this is extremely rare. The most common reason for making impossible commitments is an inequity in the power of the groups represented in the agreement. If the client group is sufficiently powerful, they be able to coerce an IT organization into making unrealistic commitments. The fourth reason that can lead a service provider to commit to provide a level of service that cannot be met, is if the service provider does not take agreement seriously. That is, if the service provider perceives that there will be no consequences if the commitment is not met. Regardless of the reason that commitments are made that cannot be met, they benefit neither the client nor the service provider. Such commitments have no place in any agreement regarding the level or quality of service to be provided.

2. Measurable—In order to assess the performance of a service provider, there must be yardsticks in place. Developing an agreement between the service provider and the client organization is one of the steps in that process. That agreement must define the quality of service provided in terms of components, or indicators, of that service quality. Since “quality” is an imprecise and subjective concept. It is not possible to directly measure the quality of a service. Instead, clients and service providers identify factors that are indicative of the quality of service. For the factors that are selected, data must be available to support to the objective measurement of the factors. If data is not available, or if the data is not perceived by both parties as being objective, then there is little difference between this and an unachievable factor. While the service provider may be able to provide excellent service, if the factor can’t be measured, it will not be possible to determine whether or not the target has been met.

Measurement of a quality indicator may mean the use of data from performance tools (such as, Expert Sniffer, etc.). Alternatively, trouble tickets (such as Remedy, Vantive, etc.), platforms (such as OpenView, Spectrum etc.), error logs, etc. can be the source of data used for measuring

an indicator of the quality of service. Whatever the source, the data used must be reliable. Consider the case of an internal telecommunications service provider. The group's annual bonus was based upon network availability. The availability was calculated with a complicated formula using data from an internally developed network management system. At the end of each month, the service group would review the outages. Any outages that were exceptionally long were declared to be "errors" and the records were adjusted to reflect a more acceptable figure. Not surprisingly, that organization was able to receive a generous bonus during the first year of the program. However, management and client groups soon realized that the basis for the measurement was not reliable and changes were made.

It is important to note that very often it is necessary to use multiple data sources to generate a reliable measurement. The multiple sources may be necessary to actually produce a measurement of a single factor, or it may be required in order to provide a way to validate a measurement. It is in this analysis, correlation and synthesis of data that highly technical measurements can be useful. These data points can be processed to become factors in producing reports that contain information that is meaningful to the client organizations. Consider the question of end-to-end response time. Unless an application is instrumented, the response time can only be calculated by drawing upon data from each of the elements that service the transactions (that is, LAN, WAN, local server, mainframe, application software, etc.). This quickly can become a data intensive, time consuming, complex effort. Organizations need to be certain that the generation of the requested information is technically feasible and economically affordable.

3. Meaningful—There are many factors that could be suggested as indicators of some aspect of the quality of the service being provided. In order to be useful, an indicator must be meaningful to both parties. There are many indicators that could be proposed, but that would not be meaningful to a client group. (In fact, much of the data collected by IT organizations is not meaningful to their clients, as indicators of quality of service. Much of the data provides a fragmented view. That fragment may be useful to IT for managing sets of resources, but it does not reflect the clients' perspective of the service.) For example, few clients will be able to relate to dropped packets as an indication of the quality of the service that is being provided to them. Similarly, the LAN or WAN network component of response time is not meaningful to the client. The client's perception is based upon the whole and not a part.

The client's perspective of the quality of the service being provided tends to be technically unsophisticated and at the same time difficult to measure. The typical client is used to looking at assessing the quality of service from the perspective of a user sitting at a terminal. Therefore, the best indicators are ones that can reflect that perspective. Unfortunately, these are the most difficult to measure. Even if you consider an IT organization as a service consumer (client), they are susceptible to the same problem. It simply is not possible for anyone to be an expert in everything. While operations group in IT may be more knowledgeable technically about many things, as a client of a telephone company, or of the internal data base administration group, they clearly are not as expert as their service providers.

4. Controllable—The factors included in an agreement between a service provider and the client must be controllable. That is, the factors must be ones that the service provider can control. There may be factors that are valid indicators of the quality of service, at least as perceived by the user. However, the factors may be beyond the control of the immediate service provider. An agreement might include availability of the wide area network circuits. However, if the agreement is with an internal service provider, and if the links consist of circuits leased from an external telecommunications carrier, the internal group does not exercise any direct control over the availability of the links. Therefore using the availability of the links as an indicator of the quality of service provided by the internal group is not valid.

5. Mutually Acceptable—As noted previously, when a client and a service provider set out to document their understanding of the factors that indicate the quality of service being provided, it is a negotiation process. The result of a negotiation process should be a document that both parties find acceptable. While neither party may care for all aspects of the agreement, they must at least find them acceptable. They must consider every element incorporated to be legitimate, even if onerous or cumbersome. If there are portions in the document that a party finds unacceptable, then they should not agree to the inclusion of those parts. Incorporating components in an agreement that are unacceptable to one of the parties is not unlike including objectives that cannot be met. It negates much, if not all, of the value of the agreement.

6. Cost Effective—As a basic principle, the cost of generating the quality of service reports should not exceed the benefit derived from those reports. While it is difficult to quantify the value received from the reports and associated processes, organizations will be wise to give at least some thought to this question. On the cost side of the equation, some managers tend to dismiss the costs of service level reporting as trivial, but they are not. There are real, incremental costs in collecting, analyzing and the data used for the generation of reports. Tools must be purchased, installed and maintained. In some cases, additional hardware is required. The manipulation of the data is normally either a labor intensive process, or dependent upon internally developed programs, with their associated costs of development and ongoing maintenance. In either case, the cost can quickly become significant. In some cases, this cost has made the production of quality of service reports more costly than an organization was willing to accept. Therefore, an effective strategy for quality of service reporting needs to address how the costs will be minimized. The most effective way to do this is, through the automation of as much of the process as possible.

D. Service Level Agreements

What has been discussed in this document is a service level agreement (SLA). The term has been avoided to this point because so many managers (both client and service provider) have had unfavorable experiences with them in the past and are reticent about considering them as viable options today. This raises two questions. First, why have SLAs been so notoriously unsuccessful in the past? Second, why should managers be willing to use them at this time?

1. Problems of the past—In the past there have been numerous shortcomings in the process of creating and implementing SLAs. These shortcomings have, in many cases, led to the ineffectiveness of the SLA process and its subsequent abandonment. The shortcomings of the past can be divided into two very broad categories. The first, encompasses those errors described earlier in this section. Clients and service providers have approached the process of developing an SLA as a contest - a test of the power of their respective organizations; the indicators chosen have not been meaningful, etc. Managers today are still susceptible to those problems. However, these are problems that have been identified. If two groups will agree to approach the SLA creation process using the guidelines in this document, the problems in this category can be minimized or avoided entirely.

The second problem that dogged SLAs in the past was the lack of effective measurement tools. Much of the data available was either piecemeal, or created manually. If the data was piecemeal, then its relevance to the SLA was usually marginal, or it forced the SLA to be written to accommodate what could be measured, even if it was not particularly meaningful. Data that was created manually was subject to another problem. Often it was not reliable due to human error, or due to deliberate manipulation of the input or the output (i.e. reports) to achieve the desired result. In some cases, individuals have been known to falsify reports. In recent years there has been a tremendous proliferation of hardware and software products (tools) for the monitoring and control of networks and systems. These management tools have vastly increased the amount and

the quality of the data available for use in SLAs. Furthermore, these advances in tool technology and product availability have also led to an increase in the reliability of the data that is generated. Finally, even when tools have been available to provide the means to capture the data required, they have provided a very narrow time perspective. That is, the tools do not generally serve as an archive of data. This has made it even more difficult to generate reports covering an appropriate span of time.

2. Value of SLAs—The service level agreement can be a powerful tool for both service providers and their clients. As was noted previously in this document (“Divergent Views”), the clients and the service providers have significantly different perspectives regarding the services being provided. It is, in fact, accurate to characterize these two communities as speaking different languages. These are languages that are not understood by the other group. This leads to frustration on the part of both groups. There are the repeated complaints such as “They don’t understand.” Or, “They just don’t listen.” The service level agreement offers a way to bridge this communications gap. By creating a document through process of negotiation, it should contain objectives that are realistic and, also, service level indicators that are meaningful. Furthermore, the ongoing process of administering the SLA requires regular meetings and dialog between the service provider and client organizations. These meetings foster improved communications. Over time, the increased amount of communication will help elevate the quality of the communications that take place. Another interesting result of this process can be increased trust, as the two groups interact more frequently and communicate more successfully.

SLAs offer a particular advantage for service providers. They offer a tool that the service provider can use to manage client expectations. It forces the client to think about what it means to have good service, and how much they are willing to pay for it. Once documented, the SLA also insulates the service provider against the gradual escalation of client expectations that can occur in the case of undocumented service level commitments.

VI. Building the SLA

The service level agreement is a valuable tool for service providers and also for their clients. However, an SLA does not appear magically overnight. The creation of an SLA that will offer value to the parties involved, takes serious commitment and effort and some advance preparation. There are several key steps involved in the creation of the SLA. Those steps are:

- Assemble a team
- Document the agreement
- Define the ongoing process
- Define responsibilities
- Define corrective action

Of course, before these steps can be taken, a prerequisite must first be satisfied. That is, the management of the service provider and of the client organization must agree that an SLA should be created. Particularly in light of the past shortcomings of SLAs in some organizations, agreement to create an SLA may not be immediate or automatic. However, once there is agreement to create an SLA, it is possible to move on to the actual creation process.

A. Assemble a team

The very first step consists of assembling a team of people from the client and service provider organizations. This team will have responsibility for developing the agreement. As noted in the previous section (Documenting the Agreement), the team must consist of people who have the authority to negotiate for the organization that they represent. They must also be knowledgeable, from their organization's perspective, about the services that will be covered by the agreement.

1. Team size—The appropriate number of members for an SLA team will vary significantly from one organization to another. Obviously, the absolute minimum number of team members is two, one from each organization. However, some companies can't seem to conceive of anything worthwhile being done by a group of less than at least 6 people, and even then only the more minor projects will have a team that small. For those companies, important projects will tend to be staffed with a team of 20 or more people. There are other, more streamlined companies that consider any group of more than 4 people to be a flagrant waste of time and resources. The size of the SLA team will also vary depending upon the scope and complexity of the agreement that is contemplated.

There is not a single "right" answer regarding the number of members required for the SLA team. However, each organization should have approximately the same number of people on the team. Some groups will be prone to send everyone except the janitor. Others, particularly those groups that are reluctantly participating in the process, will be inclined to send as few as possible. Whatever the numbers, they must be kept roughly equal. Exact equality is not essential. There should not be more than a 20% difference in the number of team members from each group. The objective of these points is to maintain a sense of equality in the course of the negotiations. Just as it is important that there not be a great disparity in the organizational power represented by the two groups, likewise, there should be equality within the team. This should help to generate more productive and fruitful negotiations.

One last point to remember when staffing the SLA team is that the team should be no larger than is thought to be absolutely necessary. People should only be added to the team when it is clear that they are required and there are specific functions that can be assigned to them.

2. Team membership—The composition of the team will vary depending upon the availability of individuals, the importance attached to the SLA, etc. Ideally, the SLA would be negotiated by the head of each organization. At least that would be ideal, in terms of their responsibility and accountability relative to the services. However, these would probably not be the ideal people in terms of their knowledge of the service and the underlying technology. Also, these are people who tend to not have time in their busy schedules to take on such a task. Therefore, the responsibility for negotiating a service level agreement will almost always be delegated downward in the organization. If the agreement is to be taken seriously, it is important that the primary responsibility not be delegated any lower in the organization than is absolutely necessary. From the service provider's organization, the person with primary responsibility for negotiating the SLA should be the person with the greatest responsibility for the service (provided that they can make time for full participation). If a manager does not have all of the technical knowledge that the negotiations may require, it will be better for them to add subordinates to the team as supporting resources, than to delegate the responsibility further down in the organization. Similarly, in the client organization, the primary responsibility needs to be assigned to the manager whose group is the largest user of the service(s) to be covered by the agreement.

Another point to keep in mind when assembling the team is that many people and organizations tend to be "level conscious." That is, they are very concerned with a person's level in the organization, or in military terms, their rank. In organizations where this is the case, steps must be taken to ensure that the team should be staffed with peers from the two organizations. That is not to say that everyone on the team should be from the same level. Rather, each person on the team should have a counterpart from the other organization. That counterpart should be from approximately the same level in the company.

There is a simple answer to the question "Who should be on the SLA team?" The answer is that it should be led and staffed by key stakeholders in providing and using the service(s) that will be covered by the agreement.

3. Team Structure—A team that is formed to create a service level agreement is not very different from teams established for other projects. One aspect of the SLA that is different from most other teams is that its members are often polarized into two opposing camps. This makes it difficult to name a single leader for a team. Yet, it is imperative that the team have leadership. The two most widely used approaches dealing with this problem are: shared leadership and arbitrary selection. Using the approach of shared leadership can be awkward, but it can be made to work. The key to making the structure of co-leaders work is to divide the responsibilities so that each has clearly defined functions rather than have them be totally redundant. The term "arbitrary selection" sounds harsher than this solution is in reality. Under this approach, one of the two leaders will essentially be designated as being the "first among equals." The actual selection is most likely to be made based upon skills, interest, or time available to devote to the function. The decision to use this approach should be made by the two leaders and not by anyone outside of the team.

4. Functioning of the team—The SLA team needs to be run in much the same way as any other team working on a project. Responsibilities of the team members need to be defined as clearly as possible. A schedule for the project needs to be developed with milestones and deliverables specified. Most important, the project needs to have a deadline for completion. All too often, members of an SLA team seem to almost view this as an exercise that should go on for the

remainder of their career. While it may not be literally true, this is the appearance that is given. If an SLA is to be created, it needs to be made a priority for the people working on it.

5. Preparation—Before the actual negotiation of an SLA between an internal service provider and a client organization, can take place, each of the teams must do some preparation. In particular, the service provider needs to secure service level guarantees from its service providers, both internal and external. For example, in the case of an IT organization that uses public network services, it must obtain service level guarantees from the telephone company. Next, the service level provider examine another question. Of the many possible service level indicators, which ones can actually be measured with a reasonable expenditure of time and money? Likewise, the client team must know what are its actual requirements. The team must resist the temptation to create a list of “requirements” that is actually a description of what they wish they could have, rather than a list of what is essential to the successful functioning of their organization.

B. Document the agreement

Once the SLA team has been assembled and organized, they can begin the process of actually creating the agreement. This process is discussed at length in the earlier section “Documenting the Agreement.”

C. Define the ongoing process

The creation of a service level agreement is a project. That is, it is a finite activity. It should have a clear beginning and a clear end. However, once the SLA has been created, there is a shift from project to process. The SLA is much more than a document to be stuck on a shelf and forgotten. The SLA document reflects the creation of an ongoing process focused upon managing the quality of the service delivered. Therefore, the SLA must not just list what indicators of quality will be used at the objective levels for those indicators. It must define the processes around those service level indicators.

D. Define responsibilities

Closely related to process definition, when the SLA is created, it must include of definition of responsibilities associated with the SLA process. For each of the functions defined, it is necessary to identify the person (by position, rather than by name) who will be responsible for that ensuring that the function is successfully performed.

E. Define corrective action

This section of the SLA defines what steps are to be taken in the event that a particular service level objective is not met. This not defined to the extent of such things as which IT technician will make changes to a router configuration. Rather, the section on corrective action should contain specifications of which group will “own” the problem of a particular service deficiency, as well as such things as escalation over time, etc.

While corrective action might correctly be classified as one of the SLA processes, it is singled out here because it is a special case. Ideally corrective action should be a process that is only performed on an exception basis. However, that is not what makes this unique from other processes. It is the fact that is omitted from most SLAs. In fact, in far too many cases business contracts omit any provision for corrective action or consequences in the event that one party fails to meet their obligations under the contract. It is quite common to see even high priced attorneys only specify in a contract what each party must do, but nothing about the consequences of failing to meet those obligations. The result of this omission is that in the event that one party fails to perform as the contract specifies, the only legal recourse is litigation. That may be great for ensuring employment for attorneys, but it does not ensure a speedy, predicable, or satisfactory resolution to the problem.

A far better strategy for constructing an agreement is to specify what each party is required to do, and the consequences of non-performance. For example, consider an example of an out-sourcer that is obligated to staff a data center facility 7 days a week and 24 hours per day. The agreement needs to acknowledge that it is possible that at some time in the future, the vendor will make a decision not to honor that part of the agreement. If this is aspect of the service contract is important to the client, then there needs to be some consequence if that commitment is not met. If you like, it can be thought of as a penalty clause. However, many people find the connotation of a penalty objectionable. The real objective of the consequence, is not to punish a service provider for failing to meet a commitment, but rather to give them an incentive to try harder to meet their obligation. Creating and negotiating inclusion of such a clause can be challenging, but in the long run it can very beneficial to a client company.

An alternative to the penalty clause that is sometimes used with effectiveness is a reward clause. That is, the parties agree that if the quality of service exceeds certain thresholds, the service provider will be rewarded. Typically, the reward will be in the form of an extra payment. However, here as elsewhere in structuring the agreement, the negotiators should try to be creative when contemplating their options.

Consider the case of a company buying software. This company insisted that in its contract for technical support for software a clause be included that covered the consequences of poor support. Basically, they insisted that a software supplier agree to “address” all problems within a specified amount of time. (Note, it was only a requirement to address, not to resolve the problem within the specified time period.) If the vendor providing the support service failed to meet this nominal obligation, they were then obligated to issue credits to the customer. This clause was rarely invoked. The amount of the credit could never represent a large amount of money. Instead, the incentive provided by this clause was that it would force the vendor to do something for which they had no process. As a consequence, they would be forced to raise the visibility of the problem in order to secure the necessary approvals required to issue the credits. The managers responsible for providing the support were usually willing to turn themselves inside out in order to avoid exposing themselves and their organization to the executive scrutiny that was almost certain to result from such an escalation.

As can be seen from the above example, the dollar value of a penalty clause does not have to be large. More important is that it will cause discomfort for the organization or individuals in the organization. For example, a client of an out-sourcer might demand that if quality of the service drops below certain levels, that the CEO of the out-sourcer must meet with the client at the client’s site. The effectiveness of this particular clause would depend upon the out-sourcer. The point is that consequence of poor quality service does not need to be financially onerous.

Creating consequences for SLAs with internal service providers can be more challenging. Seldom will cash payments be viable. Yet, a client might negotiate a provision that would have a similar effect. For example, if key systems were unavailable during prime hours of usage, then the client might be authorized to charge the cost of resulting overtime against the budget of the service provider. This is only an example. The point is to be creative and to come up with ways to apply pressure, or incentives for the service provider to meet its commitments. A very common consequence that is included in internal SLAs is a provision for the escalation of a problem under certain conditions. That is, progressively higher levels of management are notified when service commitments are not met.

VII. Elements of an SLA

As has been noted throughout this document, a service level agreement is a contract between service provider and a client. In some cases it will be a formal legal agreement (i.e., contract). In others, it will be an internal document that records an understanding the commitments of two groups within the same company. In either case the contents of the SLA are fundamentally the same.

A. Parties to the agreement

Although it may seem to be stating the obvious, the first piece of information that a service level must establish is who are the parties to agreement. This should list each of the groups who are part of the contract. That is, it should list each organization that, in the context of this agreement, can be considered to be a service provider or a client of the service provider.

In most of this document, the discussion of SLAs has presumed that there would be only a single service provider and a single client group. In reality, there may be multiple service providers, and even more likely is that there will be multiple client groups. Each of these organizations needs to be identified in this section of the agreement.

B. Term of the agreement

As with any contract, an SLA should be for a specified period of time. Typically the term will be 1-2 years. Normally, terms of less than a year are seen only in special circumstances. In most situations, there is too much volatility in technology and organizations for it to be wise to try to have a term greater than two years. In this section of the agreement, the beginning and end dates of the agreement need to be specified.

C. Services Included

It is very common for a service level agreement to include multiple services. In this section, each of the services covered by the agreement is identified. In addition, there should be a brief definition of each of the services.

1. Service level indicators—For each service covered, the service level indicators need to be identified and clearly defined. Not only must the indicator be defined, but also the SLA needs to include a description of how that indicator will be measured and who will be responsible for performing each measurement.

There is an important note of caution with regard to establishing the service level indicators. One of the most significant service level indicators for the client is the end to end response time for online transactions. All service providers need to be extremely cautious before agreeing to incorporate this into the agreement. While it is certainly an important factor, it is one that is extremely difficult to measure accurately. Before agreeing to include end to end response time as a service level indicators, the service provider needs to be absolutely certain that it will be able to accurately measure that item. Also, it is important that the measurement taken by the service provider be one that is somewhat indicative of what the client sees. If it is not, the client will dismiss it as being invalid. In truth very few service providers have the appropriate tools to allow them to be able to measure true end to end response time. The greatest obstacle to being able to measure end to end response time is the lack of instrumentation in the application software. However, this is gradually changing, and as it does, tools are emerging to take advantage of the new capability.

It is appropriate to consider a few examples of service level indicators. In the case of a web site, the service level indicators might include such things as availability (that is, how much of the scheduled time could it actually be accessed, not just server up time), the number of “hits” (connections) made to the site, the time required to download or display data to a user. In the case of a telephone company, offering a frame relay service, one currently offers service level guarantees to its customers in terms of availability, percent of frames delivered, response time (250 milliseconds) and time to repair problems (maximum of 4 hours).

There is an extremely important point that must be remembered when defining service level indicators. That is those working on this need to be creative. It is true that service level indicators must be measurable. In turn, that means that there must be a data source for that information. However, it is not always possible to find a single source that will provide the exact data that is required. In such a case, alternate sources may provide the same result. For example, if the service level indicator is availability of an application (that is, access to the functionality of the application), there may not be any management software capable of providing this information. Yet, it may be possible to obtain an accurate view of the overall availability either by correlating data about the multiple components extracted from various management systems, or, alternatively, it may be derived by extracting and analyzing data from the trouble ticket (help desk) system.

2. Non-performance—In this section, the SLA will define what is to be done when the service level indicators do not meet the levels specified in the agreement. Also, consideration needs to be given to the amount of deviation from the objective that is required before the actions specified in this section are triggered. Usually it is unwise to invoke the provisions of this section for every deviation, no matter how slight or how brief. It can result in normal “jitter” triggering alarms. This is not desirable. This section should be structured so that it will be invoked only in cases of material deviations from the objective. For example, consider an agreement that includes a service level indicator that the response time an order entry application will not be exceed 2 seconds. If on a particularly busy day, there are five transactions with a response time of 3-7 seconds, declaring that the agreement is not being met and invoking the provisions of this section would be foolish. Non-performance, should be triggered only by material deviations from the specifications of the agreement. Considering the response time example again, if 25% of the transactions processed each day had a response time of 2.1-7.0 seconds, and the mean response time for all transactions was 2.1 seconds, then a discussion with the service provider would probably be appropriate. Also, consider the case of a telephone company. If they guarantee that they will repair any outages within 3 hours, and there is an outage that lasts 8 hours, then invoking this section would probably be appropriate.

The key words to remember when working with this section of the agreement are “meaningful” and “reasonable.” If the deviations from the intended level of service are significant then the situation may warrant further consideration. When working with this section, managers should consider what it is reasonable to expect in the context of a situation. If there is a massive earthquake, even if a telephone company’s facilities escape undamaged, they will probably be overloaded. Is it reasonable to expect them to meet the service level agreement even in such a situation? That is a question that only the manager of the client organization can answer.

3. Fees—This section should document any fees that are payable for each service. In the case of an in-house service provider, there will typically not be any fees specified in this section. Although, some companies have charge-back schemes for the use of internal services and some even operate service provider groups as profit centers. In such cases as these there will be fees associated with the service level agreement. In addition, there may some fees associated with optional services from an internal service provider. In the case of an external service provider, obviously fees will be an essential part of the agreement.

D. Optional services

While the agreement is intended to be comprehensive, there may be some services that the service provider is willing to supply, but they are not included within the scope of the current agreement and/or funding arrangements. Those optional services should be listed in this section, along with all of the detailed information that is provided for each of the included services. In addition, it will be necessary to include information about how the client should request the service.

E. Exclusions

This section should be included to guard against misunderstandings. Basically, this section should clarify aspects of services that are intended to be included in either the Included Services section or in the Optional Services.

F. Reporting and reviews

It is not sufficient to just define the level of services to be provided and the service level indicators. It is also necessary for the service provider to regularly provide reports to the client about the service provider's performance. Therefore regular reports need to be produced that will compare the actual performance with the objective. The frequency of the reporting can be a subject for negotiation. However, frequent reporting seems to be the most productive. Some companies have reported that they have particularly good results from providing some real time information about key quality of service indicators. Clients are able to access this information at any time they wish. Organizations contemplating new service level agreements are strongly advised to consider implementing such a structure of limited real time reporting, coupled with periodic reports (daily, weekly and monthly), and exception reports (triggered when thresholds are crossed).

The reports that are generated must be consistent, regardless of the source(s) for the data that they are based upon. It is at this point that consistency is mandatory. In spite of the disparate sources, when the reports are generated, the data must have been analyzed and interpreted to present a semantically consistent, synchronized view of the service delivered. Also, in order to be most relevant, the reports need to be able to reflect the location of the use of the service, both geographically and organizationally. That is, in order for the information contained in the reports to be relevant, it must be able to be mapped to the groups that are using the service. Sometimes, those groups will be defined in terms of where they are physically located. In other cases, groups will be defined in terms of the organization in which they are located. Flexibility is important in this area. In part, this is important because of frequent organization changes. Also, the flexibility will support the inevitable need to summarize or abstract the data to higher or lower organizational levels.

From a practical perspective, is often advisable to present the information about the quality of service in both tabular and graphical form. Tabular reports are more difficult to assimilate, but can provide much the information more precisely. Graphical reports offer the advantage of being more quickly viewed and making trends more apparent. Also, by providing data in both forms, each user can choose to view the information in the format that they prefer.

Once the reports have been generated, it is essential that representatives of the client organization meet with representatives of the service provider to review the information contained in the reports. Like the reports, greater frequency is generally more valuable than less frequent meetings. Some companies have reported excellent results by having daily meetings between the IT organization and their client groups. The SLA should specify the frequency of these meetings and who should participate in them.

G. Other provisions

This section is not always necessary for a service level agreement. However, it is mentioned here to provide a place for any unique provisions that may need to be incorporated into the agreement. These will normally be things that reflect some unique aspect of the company's business environment, regulatory restrictions, or corporate culture.

H. Modifications

No matter how well designed an agreement may be, it is impossible to create one that can be guaranteed not to be impacted by changes going on around it. Also there is the possibility that what appeared to be good provisions in the agreement (e.g., objectives for service level indicators) are not appropriate. Yet another possible reason for modifications to be made to the agreement is the initiation of new services or the discontinuation of old ones (e.g., when a new version of an application is introduced with major changes in function and performance.). Regardless of the reason, it may be necessary to modify the agreement before it has expired. This section of the agreement needs to briefly describe the process for making changes to the SLA. It needs to describe how the process is initiated. The process for actually negotiating and approving the revised agreement should parallel that of creating the initial agreement.

I. Refinements

No agreement is ever perfect, nor are the environments covered by the agreement ever static for very long. New technology is inevitably introduced that impacts the environment covered by an SLA. For example, the introduction of a lower cost, higher speed network service by the local telephone company (e.g., frame relay) could lead to a shift in the wide area network strategy. This would change what is possible (at an affordable cost) in terms of reliability and/or response time. This would then warrant reviewing, and possibly redefining the some of the service level commitments. Also, organizations are notorious for restructuring. Often these organizational changes will also have implications for the SLA.

Therefore, any agreement must include a provision for periodic reviews and revisions. The parties need to meet periodically (e.g., quarterly) to examine whether or not the agreement is still pertinent to the current environment. If not, they must negotiate the changes necessary to restore it to relevancy. The change process will need to follow the same steps as the original process for creating the agreement.

J. Tracking Changes

Just as technology may change the level of service that can be delivered, so can changes in the client environment. For example, a significant increase in the size of the organization being served, can place unexpected traffic on a network, resulting in poorer response time. Similarly, relocating personnel, or even the IT organization can drastically effect the quality of service that can be delivered. The introduction of a new application which vastly increases the amount of memory allocated per user will also impact the quality of service, or the cost of delivering it. Changes, whether they are initiated by a client or by a service provider must be documented. When it comes time for a reviewing the quality of service being delivered, it is important to have this information available. It permits the parties to take into account the impacts of these changes, otherwise they might be forgotten or ignored.

K. Approvals

The final section should be the approvals. This is merely a section for signatures of the key stakeholders in the agreement. Normally, this should be the senior management responsible for providing and using the services covered by the agreement. This is included to document their commitment to the objectives contained in the agreement.

VIII. Administering an SLA

Once a service level agreement has been negotiated and approved, it transitions from being a project to being a process. There are three key components of the ongoing process of managing the quality of service with the SLA. Those components are actually processes themselves. They are:

- Data gathering
- Report generation
- Service level review

Since each of these administrative processes have been discussed earlier in this document, they will not be elaborated upon here. However, it is essential that these administrative processes be given just as much attention and priority as was given to the initial process of creating the SLA.

IX. Conclusion

The service level agreement is a powerful tool for both service providers and their clients. Certainly it is possible for an SLA to fail to realize its potential. However, if the guidelines of this document are followed closely, the companies can be very successful with their SLAs. Managers in service provider organizations can find that the SLA, along with the reports that are produced for it, are very useful tools for demonstrating the effectiveness of their organization, defending it from criticism and also in justifying requests for funding.

In order to be successful in developing and managing an SLA, it is essential that before starting into the process of creating one that managers first address the prerequisites. That is, they must know what data is available (that is, know what they can measure) and they must have SLAs with each of their service providers. They must define the agreements in terms of service indicators that are: attainable, measurable, meaningful, controllable and mutually acceptable. They must document their performance with accurate, comprehensive reports that go beyond providing just data and instead deliver information that is derived from a variety of data sources. Finally, information must be presented in a manner that is relevant to the client organizations.

The technology of network and systems management is finally reaching a stage where meaningful service level management is finally possible. At the same time, client groups are becoming increasingly vocal in their demands for higher quality service, and a better return on the money they are spending for a service. This is true of all service providers, whether an IT organization, telephone company, or other. The demand for meaningful service level management and reporting is going to continue to escalate. Service providers will be wise to prepare now to meet those demands.

Author

Rick Sturm is a principal with the Enterprise Management Institute, Inc. He is responsible for the program for the Enterprise Management Summit. He has done strategic planning, for the management of networks and at a Regional Bell Operating Company (RBOC). In that position he was also responsible for the review and evaluation of tools for enterprise management. Prior to that he held a variety of responsibilities, working in system software, computer operations, application development and expert systems for network management. He has managed projects that developed an expert system for network management and a project management system. Throughout all of these assignments the management of the quality of services has been an overriding concern.

Rick is involved in a variety of industry activities. He was one of the founders of OpenView Forum, the OpenView users group. He served two terms as President of that organization and is currently a member of its Board of Directors. Rick is the co-chair of the Internet Engineering Task Force (IETF) working group developing an SNMP MIB for the management of application software. Rick speaks regularly at industry conferences. He writes frequently on the subject network and systems management including a monthly column for Communications Week.

He has taught various information systems courses at the University of Colorado and at the Community College of Denver. He also serves on two advisory boards for the Community College of Denver. Rick holds a BA in Economics from Allegheny College and an MBA in Information Systems from Bowling Green University.

EUROPE

**Worldwide Headquarters
InfoVista S.A.**

6, rue de la Terre de Feu
91952 Courtaboeuf Cedex
Les Ulis
France
Tel: +33 (0)1 64 86 79 00
Fax: +33 (0)1 64 86 79 79
E-mail: sales@infovista.com

**NES Operations
Northern Europe & Scandinavia
InfoVista UK Limited**

64 Queen Street
Prince Rupert House
London EC4R 1AD
England
Tel: +44 (0) 207 246 88 50
Fax: +44 (0) 207 246 88 55
E-mail: sales@infovista.com

**SEMEA Operations
Southern Europe,
Middle East & Africa
InfoVista S.A.**

6, rue de la Terre de Feu
91952 Courtaboeuf Cedex
Les Ulis
France
Tel: +33 (0)1 64 86 79 00
Fax: +33 (0)1 64 86 79 79
E-mail: sales@infovista.com

**CEE Operations
Central & Eastern Europe
InfoVista GmbH**

Landsberger Str. 302-312
80687 München
Germany
Tel: +49 895 898 86 0
Fax: +49 895 898 86 79
E-mail: sales@infovista.com

**BNL and Nordic Operations
Benelux and Nordic Region
InfoVista BNL nv**

Drève Richelle 161, Building N
B-1410 Waterloo
Office : +32 2 352 87 19
Fax : +32 2 352 88 88
Mob : +33 671 585 003
E-mail: sales@infovista.com

NORTH AMERICA

**US Headquarters
InfoVista Corporation**

30 Corporate Center
10440 Little Patuxent Parkway,
4th Floor
Columbia, MD 21044
USA
Tel: +1 410 997 4470
Fax: +1 410 997 4607
E-mail: sales-usa@infovista.com

**North American Operations
InfoVista Corporation**

521 Fifth Avenue
Suite 2101
New York, NY 10175
USA
Tel: +1 212 972 5305
Fax: +1 212 972 5279
E-mail: sales-usa@infovista.com

**West Coast Sales
InfoVista Corporation**

1250 Oakmead Parkway
Suite 316
Sunnyvale, CA 94086
USA
Tel: +1 408 774 2960
Fax: +1 408 774 2966
E-mail: sales-usa@infovista.com

**InfoVista Solution Center
InfoVista Corporation**

12950 Worldgate Drive
Suite 250
Herndon, VA 20170
USA
Tel: +1 703 435 2435
Fax: +1 703 435 5122
E-mail: marketing@infovista.com

Asia Pacific

**Asia Pacific Headquarters
InfoVista (Asia-Pacific) Pte Ltd.**

Block 750C, #03-16/17,
Chai Chee TechnoPark
Chai Chee Road
Singapore 469003
Tel: +65 44 9 7641
Fax: +65 449 3054
E-mail: sales-ap@infovista.com