

# WORKSHOP



## SERVICE-LEVEL AGREEMENTS

### How SLAs Are Used

By Jon Saperia

In today's business environment, IT managers have eyes only for technologies and approaches that improve their enterprises' profitability and competitive position, demonstrate the value of the internal IT organization, or otherwise contribute to paying salaries. Increasingly, service-level agreements are being used to help meet these objectives.

The Open Group's Quality of Service task force shared with us its recent research study on current practices related to services and SLAs in enterprises (the Open Group is an international, vendor-neutral consortium that brings together IT buyers and suppliers to lower the time, cost and risks associated with integrating new technology across the enterprise; for more information go to [www.opengroup.org](http://www.opengroup.org)). U.S. enterprises represented in the survey came from a variety of vertical markets. There were a total of 151 respondents to a Web-based survey and a number of telephone follow-up calls.

#### Not Everything Is Rosy

Many enterprises' biggest complaint with their service providers is that they see their SLAs as being difficult to enforce and document, and the agreements often don't have sharp enough teeth (see "SLA Enforcement and Business Issues," right). When deploying or renegotiating SLAs you need to go in with your eyes open—just knowing where pitfalls lie is half the battle.

Don't give up on SLA negotiations, even if you can't get the teeth that you would like—more than half of enterprises surveyed say that SLAs are useful even if there are no penalties involved. For example, assuming you choose good metrics, you can get a sense of where you stand and use this information to leverage future agreements with that provider or obtain better agreements with other service providers. In that vein, make part of the negotiation process delivery of good metrics to show whether the SLA has been met, and ensure that the performance documentation is in a form that is understandable and useful to your organization. Clearly, given the problems and costs associated with SLAs, businesses must have a compelling reason for taking

the leap, and indeed, SLAs can support business goals (see "How SLAs Support Business Goals," page 72).

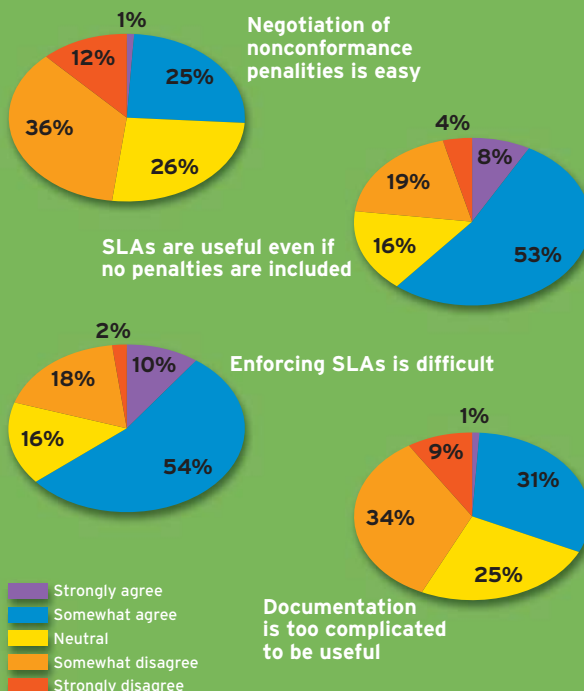
SLAs are so important that a majority of respondents would subscribe to some services *only* if an SLA were available for it. They also say that SLAs help them compete more effectively: If it's worth the expense to deploy a new network service, it's worth the expense to know if the service is supporting the business goals that drove its deployment. In fact, the majority of enterprises surveyed deploy SLAs with a wide variety of services, ranging from low-level infrastructure and transmission services, such as frame relay, to high-level services, like videoconferencing.

So whether the services you plan on offering are complex or simple, SLAs will help maximize customer (end user) satisfaction. SLAs also will help you be more competitive by letting you know exactly where your dollars are going and if you are getting what you paid for.

#### Details, Details

Many of those surveyed place great value on SLAs (see "SLA Benefits," page 72). However, SLAs grow more valuable the more precisely they are written and measured. That's where marking and measurement come in.

#### SLA Enforcement and Business Issues



Source: The Open Group

# WORKSHOP

## SERVICE-LEVEL AGREEMENTS

Marking in SLA parlance means the identification of specific services through the network. In this context, measurement means monitoring the marked services.

If an agreement is not written so that work performed as part of the service can be identified and counted, how can you know if you're getting what you paid for? Precise measurements let you confirm to your users and upper management that you have delivered value for the money spent. The more precise the marking and measurement, the more likely you are to know if the service covered by the SLA is meeting your business objectives.

How can you achieve a precise accounting? By marking and measuring the service by traffic or application type. There is a consensus that when you can identify and control services by application type you are in a far better position than when you cannot. There are two basic approaches employed by enterprises.

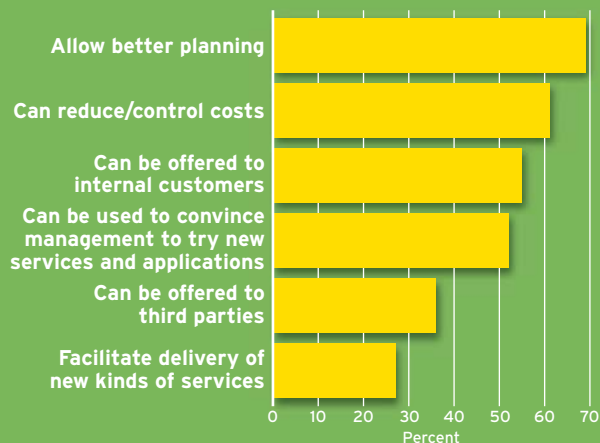
One is to employ technologies that identify different types of network traffic. In "How Services Are or Will Be Marked" (page 74), we see that for those that do mark traffic using the technologies listed, the most common service marked is Web services via URL identification. Note that only 23 percent of the respondents use the techniques shown in this figure.

The problem here is twofold: Technologies for identification of services on a network are usually complex to configure and monitor. And unfortunately, there are some services that are not easily pigeonholed or even distinguished on the network. For example, would all voice over IP traffic be important or only that traffic between specific places? The configuration and monitoring associated with this level of specificity could be overwhelming.

Also, many of those surveyed said the separation of services and their metering and monitoring are problems. While companies value information based on application type, for example, they have not for the most part invested in technologies like DiffServ (Differentiated Services) that can provide that data. Instead, they tend to provide dedicated network components and application servers for mission-critical applications.

This brings us to the second, tried and true

### How SLAs Support Business Goals



Source: The Open Group

approach: over-provisioning. Put important services on dedicated hardware with more than enough capacity to do the job. In some cases traffic is run over lines that do not share any other service traffic, which gives you even more control. This approach is also commonly used to guarantee that important services remain up.

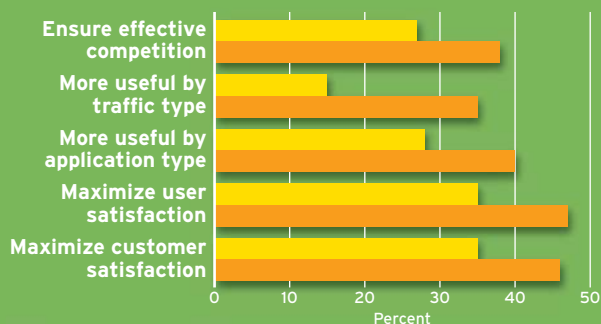
While over-provisioning vastly simplifies service-level monitoring, there is a tension between these two approaches. The first model, which typically employs complex technology such as DiffServ, can be expensive to configure, monitor and manage, while the second approach requires a hefty investment in hardware and costly dedicated bandwidth.

Making the choice is difficult, but not impossible; you need to break down costs so you're comparing apples to apples. Cost factors when considering over-provisioning are cut-and-dried: cost of the hardware and incremental bandwidth, and configuration cost of additional resources.

On the other hand, using one or more technologies for controlling traffic on the wire can save bandwidth, server and other network expenses and ensure that your network resources are indeed used for their intended purposes. The drawbacks to this approach include more complex configuration and ongoing data collection to ensure your SLAs, and personnel needed to make this more complex environment work.

Another way to look at this is to evaluate the types of monitoring and measurement necessary to support your SLAs. You'll end up with a cost analysis of over-provisioning versus "high-tech" solutions. Over-provisioning usually wins thanks to the complexity of new technologies, even though the over-provisioning approach may not scale as more and more enterprises use their networks for mission-critical applications that have SLAs associated with them. It is also true, however, that DiffServ and its ilk, while found in some new devices, such as Cisco routers, has a way to go before it can be effective.

### SLA Benefits



Source: The Open Group

tively managed. It's just too complicated for many. Two examples put this into context: In implementing QoS (Quality of Service) within a campus setting, over-provisioning, especially with the coming of 10 Gigabit Ethernet, is a valid strategy. However, when the resource in question is scarce and costly, like WAN connectivity, traffic management may be worth the overhead.

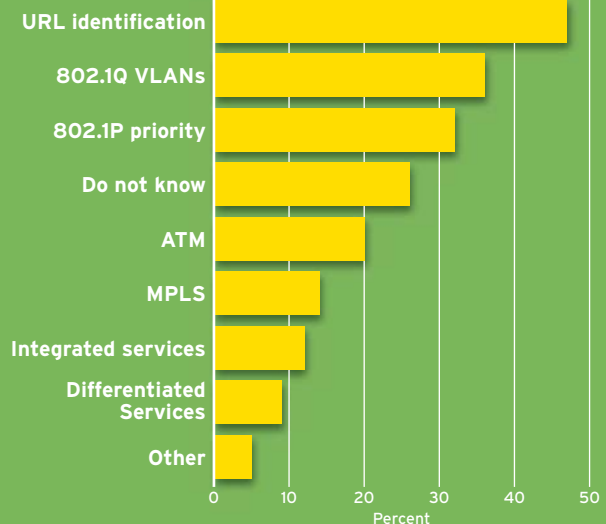
### Watching the Watchers

**"Most Valuable Monitors and Measurements"** (page 75) shows what enterprises value in regard to monitoring. Check to see if your approach provides you with these types of data. All these measures, including service-usage reporting, real-time threshold monitoring and real-time monitoring by application type, help validate that the service has, in a concrete way, measured up to the SLA.

Many enterprises do not yet take these kinds of measurements, whether they use over-provisioning to guarantee good service or one of the many technologies that control services on shared resources. It is still common for internal IT organizations to get much of their information from non-real-time user feedback. While this qualitative input is critical, augmenting that user information with real-time quantitative data will put you in a much better position to defend your budget and will provide real leverage with your service providers. Make sure you ask them for this type of data when you negotiate your next SLA.

Quantitative input adds additional value: It can be used to more intelligently perform capacity planning, resource allocation, and server and network redesign,

### How Services Are or Will Be Marked



Source: The Open Group

all of which can be costly and without significant benefit if attempted without real data.

### Get the 411

**What are the building blocks of successful SLA management?** If you want to report service outages and monitor real-time traffic, you must get, for most services, some specific information from the network to back up claims of conformance or non-conformance. There are no global standards as yet that we can apply to say that most of the high-level services are running or have run properly. For the most part, these are defined locally.

Some of the metrics are latency, packet loss, mean time to repair, mean time between failures, uptime, sometimes called availability, and transaction response time. By combining these in different ways, it may be possible for you to write into your SLAs enough metrics so that you'll know when agreements have not been met. But even though these metrics can be of value, many people are still not able to get them from their service providers. And many internal IT departments cannot provide them on a per-service basis.

"Management SLA Information Wanted and Provided" (page 75) shows information that can be used to verify the performance of a service and what percentage of users can actually get this from their service providers versus those still just wishing.

### Integrated Service Management

**So if all this quantitative data is so useful** for so many important tasks, and given the value we see in good measures, why is it that so few of us have it? Only slightly more than one third of those responding to our

## Survey By The Numbers

#### OF 151 POLL RESPONDENTS:

- 36** are in companies with 500 to 999 employees
- 54** are in companies with 1,000 to 4,999
- 61** are in companies with 5,000 employees or more.

#### THESE COMPANIES WERE IN A NUMBER OF DIFFERENT INDUSTRIES:

- 30** in manufacturing
- 26** in banking/insurance and finance
- 20** in government
- 25** in health and education
- 50** in other

#### WE ASKED HOW MANY SITES WERE SUPPORTED:

- 25** had five or fewer sites
- 43** had six to 20 sites
- 43** had 21 to 100 sites
- 40** had 100+ sites

survey got even the most basic measure of packet loss, while 25 percent lacked uptime stats.

There are a number of potential explanations. The most cogent seems to be that we do not have an integrated set of tools available in our networks that can provision and then monitor services and the agreements that we make to provide these services. What we have is a patchwork of individual tools for configuration; performance monitoring; fault management; root-cause analysis; security, for example, intrusion detection; and performance reporting. As we found in “Network Management on \$1.19 per Day” ([www.nwc.com/1402/1402f1.html](http://www.nwc.com/1402/1402f1.html)), the low-hanging fruit and performance fruit are about all you can expect without a six-figure investment. This non-integrated array of systems is a core obstacle to effective network and services management.

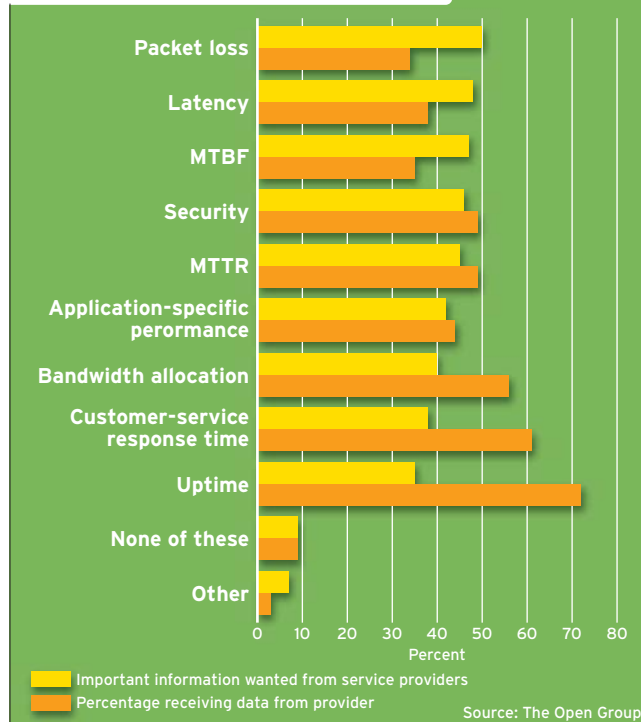
If integrated tools are so important to service management, why don't we have them? It's not that the user community has not complained about their absence. Because internal IT departments are under increasing pressure to reduce costs and deliver high-quality, reliable services, providers are being driven by their customers to provide effective service-level guarantees—and the wherewithal to back them up.

### Living With SLAs

**Don't despair—things are not as bleak** as they might sound. There are things that we can do to maximize the SLAs we make with our providers as well as those we make with our internal clients.

- » Don't hesitate to put SLAs in place, despite their shortcomings. Just be sure that the SLA is written with enough precision so that you can get an accurate reading of the state of the SLA at any point in time, even if the snapshot is not as focused as you would like.
- » Keep SLAs and your network as simple as possible—just because you *can* gather a lot of data doesn't mean that you *should*. There's a real cost associated with col-

### Management SLA Information Wanted and Provided



lecting and analyzing network data. If you have a high-value service that must stay up and offer good performance, do what many others have done: Rather than introducing new technology, devote hardware and bandwidth resources to it. Then, as management software improves, migrate to network infrastructures that allow services to share resources while at the same time maintaining the reliability and performance characteristics they need.

- » Write your SLAs with specific metrics, and define how they are to be measured. Even if you express performance as latency experienced at point A or B, that may not be precise enough. You want to be sure you understand the measurement approach and are getting a closer feel for what the users are experiencing.
- » Pressure your service providers and vendors for better management tools. Ultimately, a new generation of management software is needed, but for now, turning up the heat can bring incremental improvements. **NWC**

*Jon Saperia, co-chair of the IETF SNMP Configuration Working Group, is co-author of several recent Internet drafts in the area of policy and configuration management. He is also founder of JDS Consulting and author of a recent book on services and service management, SNMP at the Edge: Building Effective Service Management Systems. Write to him at [saperia@jdscons.com](mailto:saperia@jdscons.com). Post a comment or question on this story at [www.nwc.com/go/ask.html](http://www.nwc.com/go/ask.html).*

### Most Valuable Monitors and Measurements

