

NEGOTIATING A MORE PERFECT SLA

READING AND NEGOTIATING AN SLA IS ABOUT AS INTERESTING AS COPYEDITING THE NEW YORK CITY PHONE BOOK. HERE'S PRACTICAL ADVICE ON NEGOTIATING WITH SERVICE PROVIDERS TO GET THE MOST FROM YOUR SLA.



No part of a communications service agreement is harder to negotiate than the SLA. One reason for this is that customers and carriers have divergent expectations and perceptions that create tensions between the two parties.

On the one hand, customers demand reliable service. No IT executive wants a key data center, call center, or Web portal going offline for minutes—let alone hours—because of a service outage. For customers, SLAs ensure that quality service is received, and that tough remedies are meted out when service fails.

Carriers, on the other hand, are reluctant to negotiate SLAs, for practical and historical reasons. Carriers know that outages are inevitable: Backhoes cut fiber, router cards burn out, and software goes buggy. As negotiations begin, the carrier touts “great service” and “best-of-breed” capabilities. But when pushed to translate rhetoric into SLA metrics and remedies, it quickly shifts to expectation management mode, where “the network is the network” and networks go down.

By Mark Johnston + Justin Castillo

A legacy of limited liability and common carriage regulations, which demand that carriers serve all comers on equal terms, also influences a carrier's willingness to customize SLAs. Clinging to this tradition, carriers often insist on a "one-size-fits-all" approach, citing obsolete requirements to justify inflexibility, though most of them are now free to negotiate customer-specific arrangements.

Then there's the grim state of carrier economics. Custom SLAs mean nonstandard processes for tracking and reporting performance, which creates additional expense. They also mean more rigorous performance standards and stiffer consequences for service failures than standard SLAs.

Despite these obstacles, customers can get a fair shake. With a little care, persistence, and planning, it's possible to negotiate reasonable SLAs that promote solid performance and provide useful remedies when service flags.

THE BATTLE OF THE FORMS

The threshold issue in every negotiation is whether to start with the carrier's form SLAs. These are replete with loopholes, traps, gotchas, and excuses. Nevertheless, the pragmatic choice is to start with the form. Unless a large volume of business is at stake (\$10 million or more in annual purchases), starting from a customer's draft SLA creates conflict and delay without yielding substantial benefits. In general, SLA negotiations proceed more smoothly when the carrier's form is used, and customers focus on revising the substance.

This isn't always easy, however. Appearance and reality often diverge in form SLAs. A favorite tactic of carriers is to offer a tough-looking "target," such as 99.95 percent end-to-end availability, with remedies that start only when uptime slips below 99.8 percent. Another tactic is to exclude from SLA measures access or tail circuits leased from local exchange carriers. Still another is to count outage time for a particular service interruption only after the outage continues for some minimum period, such as 15 minutes or even an hour. As with any statistics, SLA metrics are only as useful as the fine print defining them.

Negotiating a good SLA means understanding the carrier's form, identifying the gotchas and exceptions, then neutralizing as many of them as possible. A good first step is to identify when and how the carrier's commitments apply, and when remedies start. How are outages defined and tracked? Does the customer have to report SLA violations to get credits? When do credits start accruing? When do escalation commitments start? The only way to know the value of an SLA is to determine what real guarantees it offers.

OUTAGE? WHAT OUTAGE?

Not surprisingly, carriers draft their SLAs to minimize their risks, obligations, and financial exposure. One way they do this is through the definition and use of terms. Outage and downtime illustrate the importance of terminology. Carriers often define "outage" as a complete loss of service, and "downtime" as the period during which a service experiences an outage, meaning the system is completely down. This effectively excludes any situation where there's nominal connectivity, but the service is operationally useless.

Carriers may resist revising the definition of an outage or the calculation of downtime, but it's worth pushing

Remember that the goal of SLAs isn't outage credits for the customer, but good performance and prompt resolution of outages.

the issue. For example, carriers will often compromise on the "outage" issue by acknowledging that periods of service degradation are outages, but only if the customer is willing to release the service for testing and repair.

When carriers fight revising outage and downtime definitions, they often claim that change is unnecessary because periods of degradation are covered by alternative measures, such as latency, throughput, and jitter. They're right in one respect, but wrong in a more practical sense. Although these other measures provide useful ways to track service quality, associated remedies are weaker than those provided under the availability SLA. Moreover, these other measures are seldom tracked on a real-time or customer-specific basis, except when a problem arises or the customer has an expensive network management agreement. Availability remains the touchstone for defining and measuring performance, and customers should insist that outages fall under that category.

Fixing the definition of outage and downtime is just the beginning. The revised terms must be used correctly throughout the SLA. For example, in the definition of availability, generic references to "disruptions" or "troubles" must be replaced with the defined term, "outage." Furthermore, the calculation of availability must include and accurately apply the definition of downtime. Getting the right terms in place and using them consistently minimizes ambiguity and avoids needless fights over SLA application once the services are in place.

ACCESS? WHAT ACCESS?

The Interexchange Carriers (IXCs)—including AT&T, MCI, Sprint, Level 3 Communications, and now SBC and Verizon—hate dealing with access facilities. Usually, an IXC will assume the same level of responsibility for the access it provides over its own facilities (for example, access provided to "lit" buildings) as for the other services it sells.

This isn't the case with resold access—including access connections from local exchange carriers. Resold access often generates the most provisioning and operational problems, yet it's the service or circuit element over which the IXC has the least control. Even SBC and Verizon are affected, since most states require that the local exchange part of these behemoths deal with their long distance unit the same way they deal with AT&T and MCI.

Nevertheless, the IXC should take responsibility for all the access it sells. To the customer, the IXC is the access provider and earns the revenue. The customer has no contractual relationship with the local exchange carrier to actually furnish the access circuit. If the IXC won't take responsibility, the customer has no redress if the access facility goes down. Moreover, the IXC will usually impose additional charges if the customer arranges for access directly. If the carrier is selling the service and imposing financial (and sometimes operational) hurdles that penalize customers for obtaining their own access, the carrier should be held accountable for the access it provides.

If an SLA excludes access facilities (for example, a frame relay availability SLA that only measures uptime port to port), it's important to either modify the SLA or to negotiate a separate SLA with the IXC for access. IXCs have a hard time managing the local exchange carriers, but they're the only ones in a position to do it. If the IXC isn't accountable for that access under an SLA, it has no incentive to manage the local exchange carrier, and the customer loses an important tool for managing the IXC.

THE BLAME GAME

What carriers give in service levels, they often take away with exceptions and exclusions. A typical carrier SLA will have 10 to 20 exclusions stating when an outage isn't a capital-O outage and thus won't be factored into any availability calculations. Some of these exceptions are valid. For example, excluding outages caused by agreed "Force Majeure" conditions—events outside the reasonable control and expectation of a party—reflects common contract practice. Yet even this exclusion requires scrutiny. If Force Majeure includes acts or failures of third parties, those third parties shouldn't include the provider's contractors, suppliers, or agents.

Other exclusions are trickier. It may be appropriate to exclude outages caused by the customer from the availability measure, but should such outages be excluded from the Mean Time to Repair (MTTR) measure? Arguably, the carrier should make repairs promptly regardless of the cause.

Some exclusions are simply loopholes. A classic example is an exception for service interruptions owing to "emergency maintenance." Why should an outage resulting from emergency maintenance be treated any differently when the effect on the user is the same? No service. The very need to perform emergency maintenance suggests a problem with the service.

SLA Negotiation Do's

There are many forces at work attempting to derail your SLA negotiations. Some are the result of carrier reticence, but others can be blamed on lazy or ill-prepared customers. Here are a dozen pointers to keep you on track:

- Do read the SLAs regardless of how much coffee it takes, and do view "goals," "targets," and "objectives" skeptically.
- Do involve your technical folks early, both to define SLA needs and priorities and to review the carrier's SLA offering.
- Do focus on negotiating good metrics and remedies for the critical measures of critical services at critical sites. Don't waste time and effort demanding five-nines availability across the board.
- Do know the performance record of the incumbent, and do be prepared to show why its failings (or its success, if you're talking to the other guy) compel real SLAs and real remedies.
- Do start your SLA negotiations early.
- Do understand the carrier's technical constraints—network buildout, geography, supplier constraints, service maturity, and service nature all affect what a carrier can and can't offer.
- Do focus on setting reasonable performance expectations within the enterprise, and do negotiate SLAs and associated metrics that help manage those expectations. Don't promise five-nines availability and put yourself in the position of negotiating a meaningless metric.
- Do know what the competition offers.
- Do talk with procurement and technical peers about what they see from carriers.
- Do remember that credits are a means of motivating performance, not an end in themselves.
- Do be sure that SLA credits aren't deducted from the calculation of your attainment of minimum annual commitments, and that you don't have to "repay" them in the event of termination.
- Do negotiate provisions requiring the carrier to provide the reports and data you need to track SLA compliance, and do make the carrier a part of the SLA validation process.

One way to address exclusions is to eliminate egregious loopholes and limit the remaining ones. For example, say an availability or MTTR measure excludes all instances where the carrier can't gain access to a customer's facilities. Replace this all-or-nothing exclusion with a substitute provision that excludes from applicable measures (such as availability and MTTR) only the period when the carrier requires but can't get access to the premises to fix the problem.

For exclusions involving scheduled maintenance, negotiate a reasonable notice period and cap the time frame during which service can be affected. Scheduled maintenance shouldn't mean hours of outage time when the maintenance is botched.

Finally, if the provider insists on an exception for emergency maintenance, limit it to proactive maintenance, where the carrier is attempting to prevent a more serious problem through some kind of intervention, such as installing a patch.

Private vs. Internet

We're sometimes asked if there's a difference between negotiating SLAs for private network services (for example, frame relay, ATM, and their IP-enabled counterparts) and Internet-based services. The short answer is, it depends. For dedicated Internet access, the customer is basically buying a gateway to the Internet and its associated access. Up until the data gets to the Internet, the customer should expect reliable performance and be able to negotiate meaningful availability and MTTR SLAs, just as it would for any data service.

That said, it's hard to negotiate meaningful throughput and latency measures regarding the Internet itself, and these are probably largely irrelevant anyway, unless the dedicated Internet connection is for a do-it-yourself, point-to-point VPN connection or is serving as a hub for remote users. For end-to-end services over the public Internet—think managed point-to-point VPN connections—the customer should look for SLA measures and associated metrics similar to those available for their functional equivalent in private network services. Thus, if a point-to-point VPN link is to replace a frame relay connection, seek service levels and remedies comparable to those for a frame relay connection. If the carrier won't provide comparable performance assurances for the VPN link (and until recently few carriers would), consider that fact in assessing whether the cost savings associated with the Internet service are adequate compensation for the difference in service guarantees.

SEEING THE FOREST AND THE TREES

Service element-specific SLAs are adequate for stand-alone service connections such as private lines and dedicated access for voice services. However, SLAs for WAN services need performance standards at both the individual connection level and the network level.

Connection-specific and aggregate measures each have strengths and weaknesses. Neither approach alone gives a complete sense of the carrier's performance or its ability to address systemic problems. Aggregate WAN performance metrics (for example, all of the Multiprotocol Label Switching [MPLS] connections in the WAN) may "hide" problems with individual components, but they provide a useful gauge of overall performance. Connection-specific measures allow for identification of trouble spots, but seldom provide meaningful remedies.

Increasingly, carriers are offering availability metrics for individual service connections. Although this is a big step forward in helping customers evaluate the operation of individual connections, it isn't an entirely altruistic move. The change from overall availability measures has been accompanied by diminished credits, decreased proactive reporting, and reduced emphasis on overall performance. Individual connection specifications are also often substantially lower in terms of performance than aggregate measures. For core data services, customers should still seek performance metrics at both the individual and aggregate levels.

Certain carriers now offer MTTR metrics on individual connections as well. Whether connection-specific or aggregate measures are more useful varies by customer. Connection-specific measures provide some insight into geographic holes in the carrier's service support. Aggregate measures may provide a general sense of performance, but hide specific problems. Again, it's better to measure both than to choose between imperfect options.

PROCESS, PROCESS, PROCESS

The best SLA means little without the ability to track, report, and resolve outages. With carriers shedding customer support staff and trimming account resources, it's becoming harder to obtain assistance in these tasks. Not only is it more difficult to report and collect credits than it was a few years ago, but it's also harder to track and resolve persistent service problems. Most IXCs offer Web portals where customers can track the performance of some core services, but without monthly or quarterly service meetings with a carrier representative, it's difficult to translate that performance data into concrete action to resolve service problems.

Credit reporting requirements and processes are becoming more onerous as well. The major IXCs have operational requirements in their SLAs that, at best, discourage claiming credits and, at worst, may bar remedies entirely. Two of the Big Three IXCs require customers to file a written claim for a credit within five

days after an outage occurs. This is an unrealistically short time frame. However, there's an even bigger problem: Most metrics are monthly averages, and credits are based on the failure to meet those numbers. How, then, can customers know if they're entitled to a credit based on a single outage? Form SLAs also give the carrier exclusive discretion in determining whether a credit claim is timely and meritorious, eliminating any need to be fair or objective.

Although these process problems aren't insurmountable, they can be frustrating to fix. Large customers or customers that have had some pre-existing account support with the incumbent carrier can usually negotiate an alternative, such as monthly or quarterly performance reviews and a dedicated carrier representative to push through credit requests. Smaller customers will have a harder time and must be creative in addressing the process problem. Usually, smaller customers must track and enforce SLA compliance with minimal help from the carrier, a difficult task given personnel and resource constraints.

Big or small, all customers should fix the glaring process barriers. The most obvious of these is negotiating reasonable time frames for seeking credits. Customers should have both sufficient time and the necessary data to verify compliance and credit eligibility after a performance measurement period ends. Also, customers should have the right to challenge credit determinations and to escalate serious service problems within the carrier and the customer. The carrier shouldn't be judge and jury on credit eligibility or the treatment of a service problem.

Finally, if obtaining dedicated service support is impossible, propose periodic meetings with designated representatives in both the carrier's sales and technical organizations.

MANAGING YOUR EXPECTATIONS

Service credits can't compensate for the costs that an outage imposes, and customers shouldn't expect them to. Instead, SLA remedies should motivate the best possible performance.

Most standard carrier SLAs lack meaningful performance incentives. It's both amusing and sad to read through a carrier credit calculation example. The example and the myriad calculations of credit might yield a paltry \$50, a sum not worth applying for.

The question becomes one of, am I better off expending time and energy negotiating greater credits, or seeking other remedies? When making this decision, remember that carriers hate credits and will fight to limit their liability for service failures. Remember, too, that the goal of SLAs isn't outage credits for the customer, but good performance and prompt resolution of outages.

Consider remedies that address the underlying service problem. For example, if a service misses an availability measure by a certain amount or in consecutive months, a useful remedy could be mandatory reprovisioning of the

resources :::

• *Network Computing's* article, "How SLAs Are Used," offers a practitioner's view of SLAs. Go to www.nwc.com/showArticle.jhtml?articleID=15000495.

• *InformationWeek* offers some helpful tips as well. For more information on avoiding SLA pitfalls, read "SLA Pitfalls And How To Avoid Them" at www.informationweek.com/showArticle.jhtml?articleID=6507989. If you're interested in SLAs for managing outsourcing relationships, read "Smart Advice: Consider Using SLAs To Manage Outsourcing Vendors' Performance" at www.informationweek.com/showArticle.jhtml?articleID=23901083.

service at the carrier's expense, or mandatory escalation of service problems to decision makers within the carrier.

Credit accruals can also serve as a proxy for additional relief. For example, if credits are paid out on a certain number of connections, that may serve as a trigger for additional credits, or even service termination rights.

Having this last option is important. Never give up the right to terminate for a service-related material breach. Beware of provisions in the SLA or the master agreement stating that the SLA sets forth the customer's sole remedies. If the carrier's performance is so bad that there's a claim for material breach, the customer should have the right to assert it.

Finally, tailor your SLAs to the service. A common SLA provision allows termination for chronic and catastrophic failure of a connection after three or more outages of 30 minutes, totaling 12 or more hours in a three-month period. Although fine for a private line, this may be useless for a frame relay connection. Unless the customer has multiple networks or a network split between two providers, it can't eliminate a troubled connection. Solutions for switched packet services should provide an incentive to address problems with both individual connections and overall network quality.

PATIENCE, PERSISTENCE, AND PLANNING

This article isn't an exhaustive discussion of negotiating SLAs, but it does demonstrate the need to devote the time, effort, and resources necessary to get the SLAs right. A final thought is that there's no substitute for competition. However good your relationship with your carrier, nothing motivates a carrier more than the knowledge that another carrier is trying to take away a customer. Whether the issue is an SLA, pricing, or other operational or contractual matters, the incumbent will be far more responsive if it knows that the customer is seriously considering moving its business to another carrier.

• **Mark Johnston and Justin Castillo** are partners at **Levine, Blaszak, Block & Boothby**, a law firm that focuses on negotiating telecom and technology agreements for enterprise customers. Send comments to editor@networkmagazine.com.