

August 2001. FILE NOTE

ARE OSS SUPPLIERS FAILING THE NETWORK OPERATORS (and ultimately the users)?

At a recent SLA summit run by IIR at the Meridien Hotel in London, Intercal Mondiale Limited was in the chair for one of the days. The session was quite an eye opener.

It appears that the meaning of SLAs has been hijacked by marketing hype and suppliers of software packages with relatively limited capabilities.

An SLA is a weapon used by clients to protect themselves when they outsource. Their aim is to keep the supplier honest. We heard a number of SLA claims that have the mark of the marketeer on them, for example:

- 100% Availability (except the first hour),
- 100% Availability excluding the tail circuit etc etc

One global network operator even proudly boasted that its 100% offer assumes that they will pay out service credits when it fails to meet its SLA. This means that the client does not get what he thought he was paying for and unless we re-introduce the basic engineering design techniques that seem to have been forgotten by many in the industry, there will be a breakdown of trust between the suppliers and the users, if that has not already happened.

A number of bemused attendees, especially from the Middle East, where they are about to take the first steps into full competition, were left wondering what the array of tools on offer could do to help them avoid bankruptcy when their networks failed to achieve SLAs that they had, in all innocence, accepted on the say so of a software package salesman.

There are software packages that guarantee to discover your network (as long as it only has IP devices in it), provide a history of faults, manage your SLAs, provide trouble ticketing systems, handle instantaneous service provisioning (unless you want a tail circuit) and a range of similar claims, many of which rest on very shaky ground. Users were critical of suppliers whose systems would not scale up to the levels needed by any self respecting network operator.

These are not insubstantial packages, and they cost many hundreds of thousands of pounds to procure, more effort to configure and even more to support. Claims are made that they will reduce operating costs and payback their cost in very short time periods, but little evidence of the truth of these claims was provided at the conference. Indeed the opposite seems to be true. Multiple packages are necessary to obtain even a rudimentary functionality and users can be left with significant holes in the support of their legacy equipment.

What caused some consternation was the inability of the SLA management packages to answer some of the more basic questions posed by the attendees. There were packages that could carry out Root Cause analysis, and Failure Modes and Effects analysis. There were packages that would warn you when an SLA threshold was about to be breached. Nowhere was there any capability that enabled an operator to predict in advance what the performance would be so that SLAs could be set with some degree of confidence that they would be met. One supplier quite happily mixed up the effects of faults and overloads in a single parameter that was called Availability, and claimed to be able to predict when it was about to exceed its threshold! Another claimed to be effective against operator error, which was identified as the largest single cause of downtime, the reduction of which would yield significant savings in operational costs, more than repaying the cost of the software.

One final point in this diatribe is a plea for simplification. The desire for convergence is driving ever more complexity into the network. The introduction of MPLS and its equivalent Multi-Media protocols to allow different QoS standards across a common infrastructure is not achieving its goal. What is worse is that the pressure that was driving us towards convergence (the high cost of core bandwidth) has evaporated. Whilst this has had a drastic effect on the industry as a whole, the implications for simpler and more reliable network architectures have yet to filter through. An architecture that is converged at the edges but which splits up traffic streams at the earliest possible point and then carries them across infrastructures that are properly designed in line with good basic engineering standards optimised for specific Qualities of

Service is attractive. This will certainly save some money and we may wake up to discover that instead of a massive raft of software systems to manage our infrastructure, we can actually get away with something quite simple. A welcome return to the original philosophy of the Internet.

A bewildering array of SLAs were presented to the conference. It was claimed that these would provide the Unique Selling Points (USPs) necessary for operators to generate margin and profitability. This is complete nonsense. Users are demanding SLAs because they do not trust the operators to keep to their claims - with some justification it seems. Research carried out by IML indicates that users do not want SLAs, they actually want a system that delivers what it promised to deliver. Fancy SLAs just confuse the issue.

IML sees four types of SLA and these are not USPs other than in a negative sense, they are the price of entry to the market. Without them you will not survive in this market.

The first category covers Implementation: response to queries, agreement of a delivery date, actual delivery and generally doing what was promised.

The second covers capacity: for a given agreed traffic level, does the system deliver the capacity/ latency/ delay desired by the client. This should not be based on some abstract using a low priority ping, but should be measured in a manner meaningful to the client.

The third covers operational efficiency: how often does the equipment fail (this is not the same as overload, nor is it the same as operator error - something that is anathema to operators of the old school) and how fast is it repaired. In other words true Availability.

Finally there are those parameters that do not fit into any of the above. These include security, functionality, environmental risks and so on.

All of these parameters are susceptible to engineering design and to ignore them is both unprofessional and an abdication of a designer's responsibility to his clients.

It appears that the hiatus in the industry that is causing such disruption to the established order offers all of the entities in the supply chain; software supplier, operator and end user, the chance to step back, take stock and consider what it is that the industry really needs.

Contact steve.hodson@intercai.co.uk