WHAT TELECOM SERVICE PROVIDERS WANT FROM THEIR SUPPLIERS

Jacques Gros

j.gros@vsnlinternational.com

VSNL International, 90 Matawan Road, Matawan, NJ 07747 USA

Abstract: What do Service Providers want from our suppliers? We want RELIABILITY, PREDICTABILITY, RESPONSIVENESS, OBSERVABILITY, and QUALITY in all its dimensions. Nothing new here. What is different is how much more vigorously and statistically we manage the supplier relationship, uncovering hidden costs, and the suite of tools that back up the analysis and discussions. We keep many statistics on supplier and equipment performance in key areas that impact network performance, and share these findings with suppliers in regularly scheduled meetings. Techniques used are described which resulted in substantial savings in operations cost and investment, and created environments that led to substantial improvements in our supply chains. (Techniques which we know other service providers use, such as project management of new installations, are not covered.)

1. INTRODUCTION

There can be many hidden costs associated with the supplier relationship and the quality of the products purchased, and we try to quantify and recover unjustified costs. We have been told by our suppliers (confirmed by other service providers) that we monitor, manage and articulate the supplier relationship more vigorously than their other customers, especially on failure rate and repair history of our equipment cards. As a global company, with facilities in many countries, we have supply chain needs that have to be satisfied globally so we can achieve everywhere the high availability that our customers expect. We also have significant investment in spares (traditionally, we have spared on both a site basis and using regional spare pools where that makes sense) and have many cards being repaired at any time. This paper outlines the techniques we have used with suppliers to maintain the extremely high availability standards our customers demand, while at the same time saving millions of dollars through better management of our assets and in some cases obtaining credits of millions of dollars of amelioration from our suppliers. While we monitor many statistics, the paper highlights via underlines the ones that have been used most successfully by us. Whereas the paper focuses on what service providers like us can do, it also suggests strategies suppliers can use advantageously that we have not observed used. Our undersea network was originally built and owned by Tyco Telecom, a supplier, and many of our operations staff formerly worked in factories[1] or in system commissioning, so we believe we may have a greater appreciation of supply chain issues than service providers with other heritages.[2]

There is an extensive literature on Supplier Quality Management, mostly focusing on manufacturing. While some issues are different--scrap, rework, manufacturing line shutdown, warranty expenses and product recall are important topics in manufacturing--several findings are applicable to our industry also: "Less than 50% of companies pursue cost recovery with their suppliers. And a majority of these companies only recover material costs from their suppliers. According to a recent report by AMR … about 65% of the costs attributed to the poor supplier quality are non-material related."[3] The evidence suggests the percentage of service providers seeking recovery is smaller in our industry, and, if anything, we only recover repair (material cost) even though there are other costs in addition to the direct costs we pay our suppliers. While the material cost percentage can be well over half of the direct expense of problems (direct costs include expensive repair costs, with additional direct costs in providing recovery to customer (SLAs) and expenses getting someone to the site), in reality, poor supplier quality cost is dominated by the non-direct costs due to

[1] Many of the concepts of this paper come from our manufacturing experience given both our network and a manufacturing line require the equipment to be working to provide outputs. Good manufacturing concepts include making problems visible, looking for bottlenecks, teamwork, dropping everything when there are problems, and having no built-in protection. We have noticed other companies have other philosophies in the topics discussed in this paper, such as having maintenance agreements rather than holding spares, and do not follow-up on supply problems as we have.

[2] We have been sensitive not to use or retain proprietary Tyco Telecom (and Lucent, Nortel and Telcordia) data (from previous experience) in our discussions with suppliers. We have observed other suppliers have benefited from their insights from owning service providers and moving people between companies, such as Ericsson in the architectural design of their AXE family of telecommunication switches. See A switch in time: AXE - creating a foundation for the information age by John Meurling.

[3] http://www.metricstream.com/insights/bestPractices_supqltymgmt.htm

holding additional spares consistent with supplier concerns.

2. RELIABILITY

We take several views of product reliability: card view, product view and circuit view. Card view: We monitor and keep statistics on failure rates of each active card in our network[4], including monitoring out-of-the-box failures and short-life time failures. (Out-of-the-box and short-life-time failures have very deleterious impacts on the network, forcing us to hold many more spares. Most sparing tools assume that all spares will work when used, and given our experience with some codes having many out-of-the-box failures, we had to create a new sparing tool to measure the impact of failed spares (see figures 1a, 1b and 2 for the large impact out-of-the-box failures have, where even a small percentage of spares being bad requires adding another spare to the spare pool)). High rates also strongly suggest problems in the repair process.[5] We also monitor if the same card has failed several times.) For each card type, we compare the actual failure rate with what was expected (by multiplying the number of cards in service by the failure rate of the contract (often by converting Failures in Time, FIT, to failure rate per year)). We also monitor whether failure rates are getting worse over time, along with failure rates after they have been repaired.[6] For cards we calculate the number of spares needed to meet our sparing criteria, along with the outage rates (meaning in this case not providing service due to not having a spare when needed, see figure 3 how outage rate varies with number of spares available) associated with the number of spares suppliers provide in contracts. We especially focus on inconsistencies with the failure rates actually experienced, and the failure rates that the suppliers said would be the case in their manuals. When there is a discrepancy, with the experienced failure rate much higher than predicted, we have discussions with our suppliers on causes, often aiming for remediation in the form of free repairs (so we only pay for the predicted number of failures), free upgrades if there are generic problems, or free spares.[7] Of course, part of the discussion is on the quantified impact of the higher than expected failure rate in terms of additional repairs we pay for, additional disruption to our network (including calling out technicians to do repairs in the middle of the night) and customer circuits, and in the additional spares needed in the spare pools (see figure 4 how number of spares needed varies with the failure rate expressed in FITS, failures per billion hours).

Product View: Using card failure rate data, we create reliability models of equipment, especially from the view of a circuit crossing the equipment. The calculations take into account cards that operate redundantly, cards on circuit path and also cards that do not touch the circuits. Of course, for our equipment, the cards a circuit crosses depends on the size of the circuit—different cards touch different size circuits and have different failure rates, and impact different numbers of customers.

Circuit View: We model the availability of circuits, given that is what our customers are interested in, and we provide Service Level Agreements guaranteeing service levels are achieved. These circuit availability models use the reliability models of the equipment on the circuit's path, and also take account of the transmission medium (e.g., cables), along with operator error. Similar to what was done to model reliability of circuits crossing equipment, the reliability of circuits across our network depends on the size of the circuits and the path the circuit follows

3. PREDICTABILITY

We monitor how well equipment performance matches documentation that was used to sell and operate equipment, in terms of features, functionality and meeting standards. For upgrades, we confirm backward compatibility with previous versions, and when there is not backward compatibility, as we found in a recent case, we need to better manage the cards used, making sure spares are available for all versions in service (we quantify the additional spare cards needed). We keep track of how consistent suppliers are compared to contracts and promises. For repairs, do cards come back within the contractual repair internals? When repairs are delayed, or we have many more failures at an individual site than anticipated, we have to make workarounds so that there remains sparing coverage. When repairs take longer than anticipated, we need to

[4] We are very conscious that many of our cards have very low failure rates, so one or two failures, while implying an actual higher failure rate than anticipated, are not significant from a statistical basis. In addition, some cards share the same design, with just differences in frequencies (e.g., lasers), so we do analysis of all these cards together, obtaining a more statistically significant population for the analysis.

[5] Changes in repair processes have resulted from our finding high out-of-the-box failure rates and early-life failures. We have also changed our inventory database to capture out-of-the-box card failures.

[6] For cards with suspected high failure rates after repairs, we keep track of when these cards are returned to service and when they fail, and find a failure rate based on these in-service intervals.

[7] What is striking is that suppliers provide FIT (Failure in Time) rates for individual cards, rather than for equipment as a whole. On a statistical basis, one would expect that some cards would fail at greater than the published failure rate, while others at less, with the overall product still meeting an overall rate.

hold more spares in our spare pools, and quantify these additional needs. (Histograms of when cards are actually returned have been very effective in the discussions (see figure 5 of an example), along with quantified additional needs for spares due to longer than anticipated Return & Repair intervals (see figure 6).[8]) We also provide feedback to the repair organizations, and have worked with them to streamline our interactions, especially for a supplier who had quite a few repairs with very long repair intervals.

4. RESPONSIVENESS

How responsive suppliers are when there are problems is critically important in the resolution of problems. Part of responsiveness is how open the supplier is when there are problems, for example, for repairs, keeping us informed when there are delays, compared to us finding out for the first time when cards don't meet their dates. There have been situations where we had too few spares, sometimes due to delayed repairs or other times insufficient numbers were included in the supply contract, and were appreciative when suppliers worked rapidly to fix the situation. Regularly scheduled meetings provide an opportunity to discuss issues before they become problems.

5. OBSERVABILITY

By observability, we mean the remote, accurate identification of problems via the Element Management System (EMS), Network Management System (NMS) and via other means, including SNMP. Under observability, we also mean that one problem does not mask other problems—that all are identified. The tools should also indicate how fast problems need to be fixed, identifying which are customer impacting and which can cause other problems if not fixed quickly (e.g., a fan problem could result in overheating cards, or too high signal levels could impact amplifiers down the line). Another aspect of observability is that there are not false indicators of problems—we have experienced many times the EMS indicates a problem and when the card is sent into repair the card comes back No Trouble Found—leading to the expense and inconvenience of sending someone to a remote site for a repair, the cost of sending the card in, the greater number of spares required given so many cards are in the repair cycle. When No Trouble Found percentage is a significant fraction of the cards sent in for repair, we have discussions with the suppliers what we can do differently to cut the rate; in some instances they suggested we capture more data and share it with them before returning the card; in other cases, do additional tests or check cards which drive the cards that are reported failed.

6. QUALITY

Quality improvement in supply chains is usually a collaboration among a service provider and its suppliers. We strive for better ways of incentivizing this collaboration, especially in areas where our suppliers have insights and data we do not have access to (access to product plans, developers, repair facilities, product histories). Given quality in all its dimensions covers many topics, the focus here is on a few topics where quality can be improved through enhanced cooperation. One key area for us is data latency—cards fail infrequently, and it can take a while for enough data to be available for statistically significant data to exist to confirm whether cards are failing at greater than expected rates. Of course, our suppliers observe failure experience from all their customers, which provide earlier insights when there are problems. We encourage our suppliers to share what they see overall, and have instituted improvements in our processes (e.g., field testing) in response to what suppliers reported works in other networks. What we are trying to achieve is to have appropriate systems and processes to manage and improve supplier quality in a repeatable and predictable manner. We therefore hold regularly scheduled meetings, focusing discussions on statistics on outlying products, and, to avoid discussions on the underlying data, tend to use supplier data (which we confirm). In addition to expecting learnings from other customers, we expect these meetings to also cover whether there are generic underlying problem causes, and then work together to resolve the issues (it is the suppliers who have the information on the underlying root causes which come from their incoming repair tests and how they repaired the failed cards).

7. CONCLUSIONS

We have successfully quantified the costs of poor supply chain issues and used these quantified costs in our discussions with suppliers. Utilizing our heritage, we have more aggressively than most other service providers managed our supplier relationships, including monitoring the failure rate of our cards. Many of the techniques and statistics we use were described, highlighting those used most successfully, and the quantified impact is in millions of dollars of offers for recovery of the additional expenses we bear.

[8] Suppliers have provided us shorter R&R intervals in response to the quantification of how many more spares would be needed—the impact of R&R intervals on number of spares needed is large.

8. FIGURES

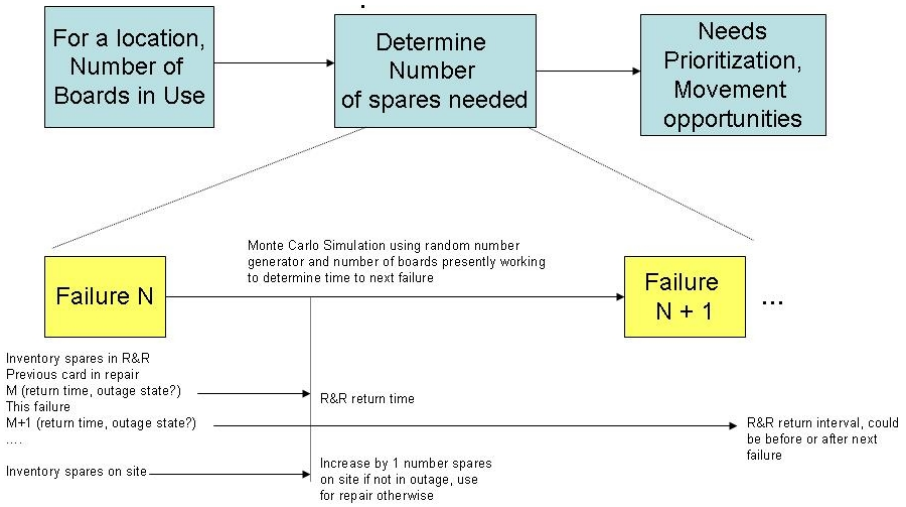

Figure 1a: Determination of Spares Needed Assuming No Spare Failures

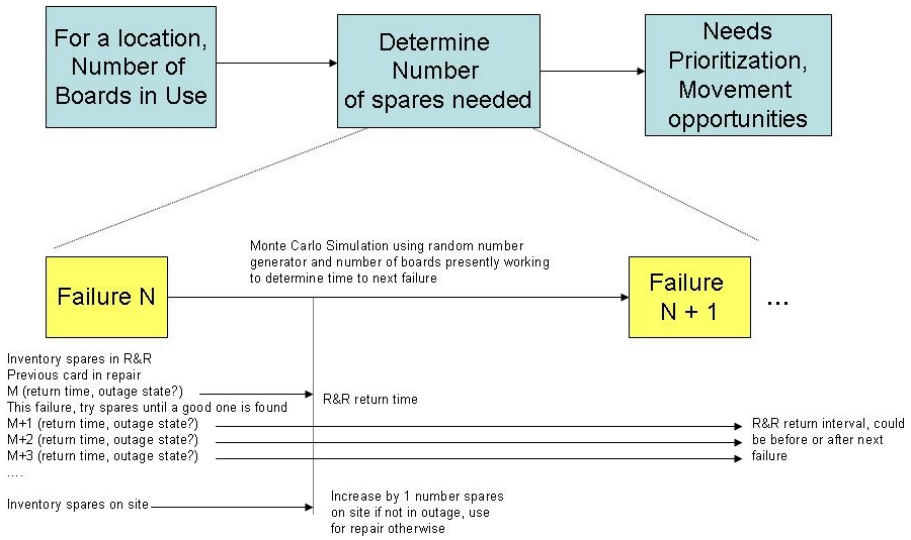

Figure 1b: Determination of Spares Needed Assuming with Spare Failure-Large Impact

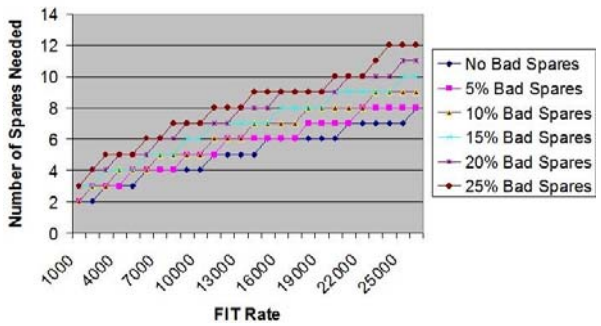

Figure 2: Number of Spares Needed versus %Bad Spares and Fit rate (20 cards operating)

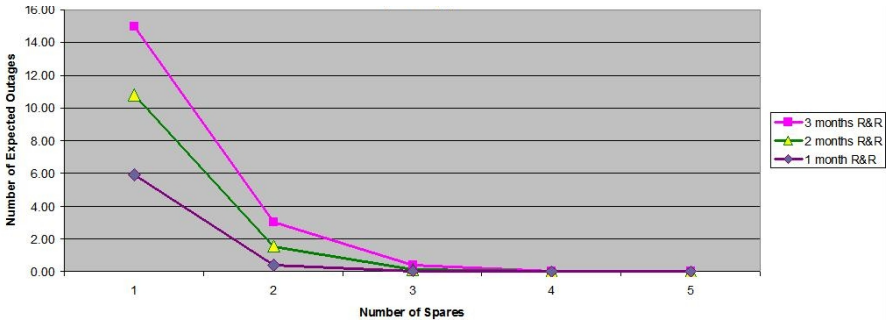

Figure 3: Expected Number of Outrages in 25 Year System Life due to Not Havung a Spare (10,000 FITs, 20 operating)

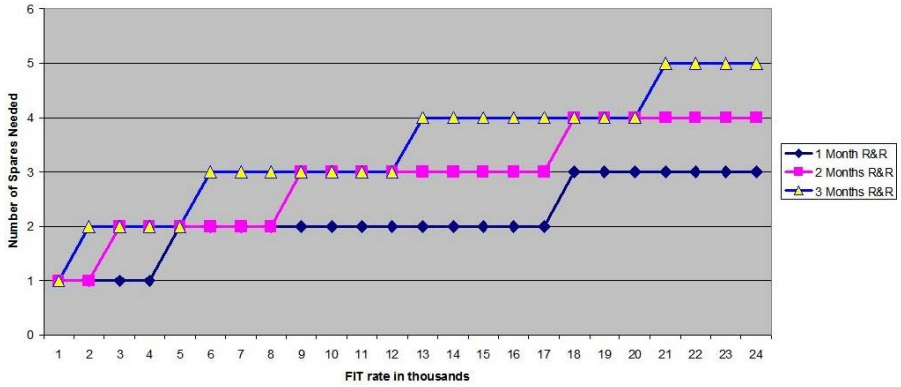

Figure 4: Number of Spares Neededd versus FIT rate (20 operating, various R&R Intervals)

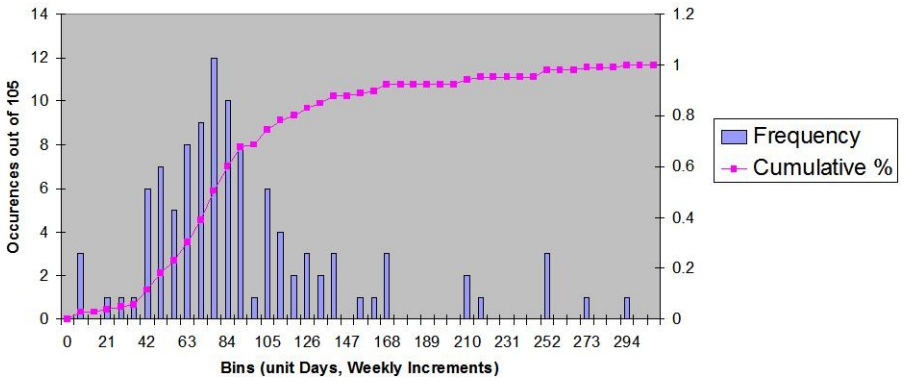

Figure 5: Supplier ‘B’ R&R Interval Histogram

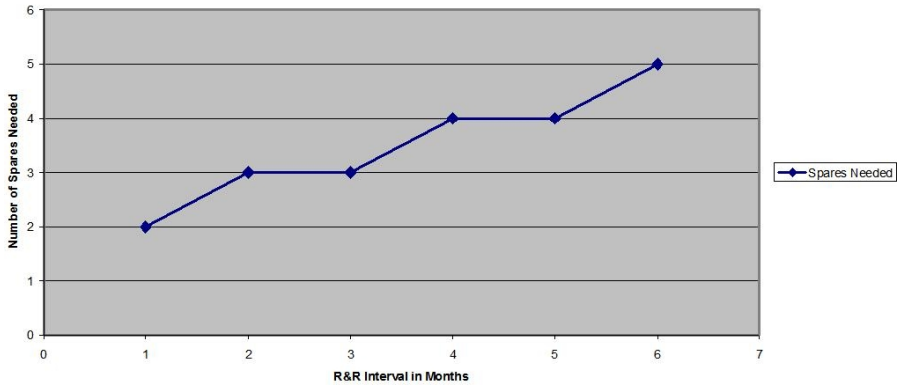

Figure 6: Number of Spares Needed versus R&R Interval (1000 FITs, 20 operating)