Total Cost of Ownership: Issues around Reducing Cost of Support in a Manufacturing Organization Case

Stefania Castellani, Antonietta Grasso Jacki O' Neill, Peter Tolmie
Work Practice Technology Area
Xerox Research Centre Europe
6, chemin de Maupertuis, 38240 Meylan, France
+33 4 76615050
{*name.surname*}@xrce.xerox.com

Abstract.

The notion of Total Cost of Ownership has become a particular focus of interest across a wide range of commercial and academic communities in recent years, because of the way it enlarges the understanding of what the real cost of technology amounts to, beside purchase cost. The dominant conception of TCO derives from a model created by Gartner Group that considers 'hard' and 'soft' costs that are distributed across three prime categories: acquisition costs, control costs; and operational costs. In order to measure those, the so far proposed TCO metrics tend to be based upon conventional, abstract categories of cost as defined by accounting and management literatures. These categories rarely attend to what might be understood as a cost in the thick of the everyday work of the actual members of an organisation that such categories are applied to. Our own research presented in this paper augments the field primarily by focusing upon break-fix costs, but also by relating these to the wider issues in TCO. The goal within this has been to not just take costs for granted but rather to try to understand what something like 'break-fix' amounts to in organisational life. This has been pursued by looking at the details of how costs arise and are reasoned about in actual working practice. From this we have then sought to understand what current barriers exist to the implementation of a successful break-fix strategy and what kinds of technologies might be designed to facilitate that success.

1 Total Cost of Ownership: an Overview

Whilst the notion of the Total Cost of Ownership (TCO) of devices is today almost a commonplace within various communities concerned with the purchase and maintenance of technology, one can actually find the idea being put to direct use with regard to computing technology as early as the 1960s (see, for instance [7]). At that time the interest was more focused upon whether computing technology should be leased or bought and upon arriving at ways to assess the subsequent impact of those decisions. It would seem, however, that the interest in cost, what amounts to a cost, and what one should be seeking to incorporate in cost calculations has evolved over time. Indeed, it is hard to ignore the fact that technology has itself made it easier to locate certain costs and perform more complex cost calculations as it has grown in sophistication, and the focus on TCO accompanies a trend amongst managers to use ever more elaborate computer-generated metrics in their everyday activities.

The idea of TCO we are perhaps most familiar with today is the one derived from the Gartner Group's notion of:

"a holistic view of costs related to IT acquisition and usage at an enterprise level" *[5]*

which is as much concerned with how to manage ongoing costs as it is with how to rationalise initial outlay, though outlay is an inevitable feature of that. Perhaps the principal component of the Gartner model is the distinction drawn between 'hard costs' that are directly expressed and relatively easy to trace in budget statements, and 'soft costs' that are often not clearly formulated in any one place or perhaps not even regularly 'costed' at all because they arise outside of acquisition. As an example, a study to "understand and interpret the various factors that affect the TCO of printing" was undertaken by Gartner Consulting for HP in 2001 [4]. In addition to costs of:

* hardware
* software
* consumables

this study also identified a number of other cost categories which it believed should be included in any assessment of TCO, including:

* costs of network management and administration
* infrastructure costs (floor space for devices and supplies)
* maintenance and support costs
* "productivity costs" (loss of productivity due to inadequate, inappropriate or poorly functioning printing equipment).

Over the years the original Gartner model has been refined in various ways. David *et al.* [3] for instance, have proposed that TCO should be considered to consist of the 'Acquisition Costs' of hardware and software, 'Control Costs' such as 'Centralization' and 'Standardization', and then a whole range of 'Operations Costs' such as 'Support', 'Training', and 'Downtime'. They offer on the basis of this a simple 'per-seat' model of TCO but then go on to note that this fails to take into account the various possible interrelations between costs that a sophisticated and effective model should ideally have built into it.

2 The Dimensions of TCO in the Case of Multi-Functional Printing Devices

There are a number of main components that are presumed to amount in sum to the TCO of some technology. Some of these have been outlined briefly above to provide a sense of what a conception of TCO might be considered to consist of. However they can remain quite abstract if not related to a specific case and even a type of technology. In the rest of the position paper we will therefore focus how these dimensions apply to a printer manufacturing company in general and then we will further detail one sub-category, that of break-fix support.

The cost of purchase: There is evidently a cost associated with actually going out and buying output equipment fleets such as printers, copiers, fax machines, and scanners [5]. Frequently, the type of support you are opting for as a customer is now an integral part of the sales package offered when purchasing a device. Thus, to take a Remote Call Assist type contract with a device implies certain kinds of services one might turn to in break-fix situations. And to buy such a contract is a cheaper option than seeking to purchase full service engineer support regardless of the circumstance. There is scope in that case for tackling break-fix costs at the point of sale, though, contractual obligations are not necessarily enough to persuade all customers to troubleshoot.

The cost of consumables and associated supplies: The regular purchase of things like toner and ink for printers or copiers is one of the prime ongoing costs associated with owning such a technology [5]. One way of managing such costs might be seen to be the production and purchase of consumables at the lowest possible price or the optimisation of the supply chain. One feature of break-fix that could be considered to fall within this category is increasing the number of potential customer replaceable units (CRUs), thereby minimizing the need for service engineers to go out and replace parts. Basic CRUs like fuser and xerographic units are a standard consumable for many machines and for some machines there is a push to extend this to certain types of rollers as well. Thus one aspect of a strategy for reducing break-fix is to increase the scope of consumables.

The cost of support, maintenance and management: This can cover a vast range of things from the provision of a helpdesk, to specific maintenance agreements, to costs of downtime and repair, not to mention activities associated with providing and monitoring things like printing facilities. It can also cover things like upgrades, for instance providing a high capacity input tray for a printer, or additional software that might support activities like printing [5]. This area, in particular, is the subject of a considerable current thrust in enterprises in the document production industry. Machine self-diagnosis, the increase in customer replaceable units, the quest to optimize troubleshooting activities and technical callouts, all represent key endeavours to bring down cost in this domain. This, of course, is the primary TCO component associated with break-fix and where break-fix costs would traditionally be located.

The cost of access: The physical amount of time it takes to get access to a device may be opened up to measurement, but so too may the administrative cost of managing how rights of access are allocated in the first place [4]. In break-fix terms this can appear a highly complex consideration. For instance there are associated costs with just how long it might take for support personnel, engineers, or troubleshooters to gain access to some device to effect a repair, with the longer it takes the more costly the repair in both labour terms and downtime.

The degree and range of utilisation: A favourite metric of many modern TCO calculations is the extent to which a technology, such as a printer is currently used, how much time it is idle, and how much of its provided functionality is drawn upon. In that these may reveal under- or over-utilisation of machines relative to others is seen to be a critical component of calculating the optimal number of devices and their position in an organisation. This, together with other features such as the cost of replication (replication of functionalities across the total number of devices in a given organisation), the cost of mix (having mixed devices from different manufacturers requiring different kinds of support), and the cost of access, is a component in the assessment of optimal device configuration from a number of companies. Whilst the relationship between this aspect of TCO and break-fix may be less direct there are a number of ways in which the two may be seen to be related. For instance, the extent of utilisation of a machine will have some kind of impact upon the frequency of downtime to be expected as no parts will simply last forever. This also affects, of course, the consumption of CRUs and the longer-term reliability of a machine.

Time costs: A huge number of things may be assimilated as costs in time: pretty well anything that involves labour may be expressed as a pro rata division of the salaries of those involved; downtime may be expressed in terms of loss of productivity and lost profit, and so on. This is another core component associated with break-fix concerns because many break-fix calculations include the cost of downtime quite explicitly.

The unitary cost (as a function of others in order to work out optimal total cost): Finally, it is worth noting that a great many TCO assessments find their ultimate expression in some kind of unitary cost, for instance the cost per page for a printer. Compound costs are thereby reconfigured so that they can be expressed in ways directly associable with throughput of some kind. Optimal total cost may then be expressed in similar terms. Furthermore, in that break-fix components were found to be present implicitly or explicitly in most of the above this unitary cost will inevitably contain these components in its ultimate expression.
Other costs relate: to specialisation (having different devices or multiples of the same); legacy devices (in that they might break down more, be difficult to service and may not interact well with other newer devices); expertise and training (to use a new technology); costs of processing and pre- and post-processing (these are so-called ‘soft

costs' in the Gartner model, used by more sophisticated analyses of TCO – focusing on the impact on work processes of the implementation/use of technology); and unauthorised costs (where technology is used for purposes that offer no productive consequences for that organisation, such as using printers for personal documents (described by Gartner as 'futz')).

It is of course the case that many of the above costs either overlap with one another or are wholly contained within one another. The distinctions between them amount to distinctions made at different times by different groups researching or seeking to manage TCO in some way. It is clear that no model is absolutely comprehensive and that some categories are rather vague. It is of no small significance that what gets wrapped into TCO is often: a) what those involved in assessing it believe they can measure and quantify in some way; and b) what is taken to be in some respect variable and therefore open to being 'fixed'. And this, of course, is no different with regard to the calculation of break-fix costs which are quite clearly implicated throughout the range of possible TCO components that are conventionally worked upon. This underscores the fact that measuring TCO is typically motivated by some kind of desire to either 'reduce' it or facilitate improved productivity (which, as a 'doing more with the same' amounts to an overall reduction in unitary cost), not just to know it as a state that might be informative for completely different organisational ends. However, existing research into technological solutions for the reduction of cost in organisations has largely been constructed upon taken for granted understandings of what cost amounts to as an organisational concern. Whilst this offers certain obvious advantages with regard to the assessment of measurable and a generically coherent body of reasoning that amounts to an exhortation to minimize cost and maximize productivity, there is more to organisational life than such abstractions are able to capture.

A central tenet of how we have conducted the research underlying this paper has therefore been one of arriving at an understanding of what break-fix costs might amount to in their detail, not just in the above kinds of abstractions. To illustrate, one of the metrics used to calculate TCO is to calculate the distance between devices and users. That is, how far must the various users walk to get to the device (this can easily be calculated using floor plans)? The suggestion for optimizing TCO that then comes from such a calculation might be to rationalize the position of the printer in relation to its users so that in a cumulative sense the distance people are walking to a printer is minimized. This might then appear to reduce TCO, by reducing time costs. However, this analysis may sometimes prove too simplistic in that it only takes into consideration device usage - it does not consider how printing might be situated in the broader flow of work people are undertaking. An example of this might be a user who regularly passes the mail room on his way to the printer. As this is a regular occurrence and he is responsible for distributing mail around his section he takes the opportunity when he goes to the printer to drop into the mail room to check for new mail. The new position of the printer may change his route, thus requiring additional work in going to the mail room as well. Nor is it inconceivable that this change in printer position might also have impact upon other users' work practices. Thus by not considering the work that is done in 'going to the printer' but instead basing printer location upon relative distance as an abstract but quantifiable category of TCO, any judgments of reduction in TCO may be flawed. This is what motivated our study and it will be found that one of the central differentiators between our approach and other approaches is the concern to ground our explorations of TCO, and more specifically of 'break-fix', in a close attention to how costs are made manifest and oriented to in people's actual working practices, rather than restricting ourselves purely to some abstract set of pre-givens.

3 Reducing the Cost of Support

In order to apply an observational approach to the problem we are going to focus on one specific aspect of TCO, the one of support, and for doing so we are going to look at how cost stands as an issue for support organizations. We will be using a typical large printer company's technical support as our articulating case, but many of the issues we touch upon here are the kinds of issues that any support organisation might encounter.

In the following pages we will be examining the costs borne by a printer supplier in supporting its document input and output devices and examining the main themes in customer support. Reducing support costs has a direct impact on profit margin on support contracts and/or improving the competitive positioning of the company's products and services. However, it is important to remember that this must be weighed against certain other things such as costs in production and/or a possible loss of competitive advantage if it is seen to lead to a concomitant drop in the level of service.

There are five broad categories of activities which contribute to controlling support costs while providing efficient support to the customer, all of which can be seen to overlap and interrelate in certain ways:

* Making devices more reliable
* Facilitating device diagnostics (both at the device and remotely)
* Making it easier for customers to resolve problems themselves (both through troubleshooting and through making certain parts easier to replace)
* Improving the efficiency of technicians in the field (primarily by reducing contact time with customers and a reduction in the numbers required to service the same territories)
* Maximising the efficiency of telephone support to customers

3.1 Making Devices More Reliable

Improving the overall mechanical reliability of devices is one of the constant preoccupations of the engineers designing new models, together with how they might improve the number of pages per minute that are printed, and in recent years the lifetime of the Xerographic module has increased from about 20,000 prints to over 220,000 prints. Although further improvements can be expected, there will always be a trade-off between enhancing reliability and keeping down the cost of a device, especially with current pressures to reduce the time to market of new products.
Other, more recent research efforts have been focused on maximizing printing device up-time by adding redundancy. One such effort is called TIPP (Tightly Integrated Parallel Printing). This new breed of printers includes two, and potentially more, print engines that work in parallel to provide higher-speed printing machines with lower speed (and cost) engines. One of the obvious benefits of this approach is the ability to provide some redundancy: if one of the print engines fails or needs to be maintained, the second engine is still able to continue printing, although at a lower speed. Although this does not directly reduce cost, since the maintenance has to be performed anyway, this can reduce the penalties incurred in certain maintenance contracts based on down-time, or alternatively afford a more efficient usage of the maintenance force by leaving more time to react.
However, against all of these things it is important to remember that whilst enhancing reliability may reduce the number of calls received over time it may also increase the initial price of the product, as may adding components or additional machines to improve up-time. In some concerns where continuous printing is critical up-time might be given a disproportionate value, but in other cases it might be seen that this is contrary to the optimization goal of TCO by creating redundancy.

3.2 Facilitating Device Diagnostics

An effective diagnostic capability can reduce support costs, especially if the diagnostic information can be made available to service engineers at a remote site, thus potentially avoiding a site visit.
Over the years, a growing number of sensors have been embedded in printing devices, and this has enabled an increasingly precise diagnostic capability.
Initially available for the maintenance technician to diagnose a problem on-the-spot, a subset of the data collected by these sensors is also available through the network on the machine's built-in server to provide status information. This data can also be used to provide simple, customized troubleshooting instructions to the end user on the device's front panel.
This growing number of sensors can enable a number of initiatives to turn such devices into intelligent devices. The basic idea is that by endowing devices with greater intelligence they are able to monitor their health and diagnose faults as they occur. Monitoring the device health may enable preventative maintenance, while fault diagnostics should provide more precise information on faults to end users, or to service engineers, thus shortening the break-fix cycle. Device intelligence can also be applied

to monitoring of media and consumables, with the possibility of automatic integration into the order process.
One of the difficulties that have been experienced with intelligent devices is that they can err on the side of producing too many false positive results – declaring non-existent faults. This leads to the conclusion that it is dangerous to rely completely on such techniques, and it is preferable to use them as supporting tools, for instance by providing rich information to support centre troubleshooters or engineers. Furthermore, it is also clear that, as with projects designed to improve reliability, focusing upon improving diagnostic capability may also come at a price in design and production which has to be weighed against other possible savings. On top of all this, it should be remembered that the efficacy of such a technology turns upon the problems being open to being sensed in that fashion. It is questionable whether some issues, such as a loss of image quality, can be sensed readily as a fault by the machine.

3.3 Making it Easier for Customers to Resolve Problems

One clear way of minimizing the cost of support is to enable customers to repair the printing devices themselves, thus saving the costs of technical support. This can be achieved in a number of ways, including designing machines to be easier to repair by the customer, providing efficient online or off-line resources for the user, or providing the user with customized instructions for troubleshooting a problem, principally via the telephone (with the caveat that the customer is actually willing to undertake such self-repair – something that cannot be taken for granted).

3.3.1 Customer Replaceable Units

Customer Replaceable Units (CRUs) are parts of the device that are easy enough for the customer to replace. For most current device models the Xerographic unit and the fuser module are CRUs, in addition of course to toner cartridges. Some machines also offer some rollers as CRUs. Not all devices currently make a clear distinction between CRUs and consumables and clearly the principle is much the same: when the part is exhausted the customer replaces it.
The way in which CRUs are managed is somewhat different in the US and Europe. In the US customers are encouraged to stock CRUs on site and change them themselves when necessary. In Europe, CRUs are delivered and installed on demand by a courier service. The US system has the disadvantage that the immobilized stock of CRUs at customer sites represents a significant investment, and for customers with a service contract it is a charge to the supplier.
While the concept of CRUs is clearly a sound approach to reducing support costs, a successful implementation depends upon a high level of reliability of the corresponding replaceable modules, and it could not always be the case. For instance, it was reported to us by some service engineers that there are machines in one of the flagship product lines for the concept of CRUs, that are subject to an excessive number of callouts. This pushes up the cost of these machines both in terms of physical callout time and in terms of potential customer dissatisfaction.

3.3.2 Off-line Resources

The quality of user documentation is also a significant factor affecting the ease of repair, and here again significant progress has been made in recent years in improving the quality of user documentation resulting in several product lines where a CD-ROM containing documentation is provided with the device.

In a similar way, many printers provide simple troubleshooting instructions from the control panel of the machine, to solve a number of problems.

Customer surveys have indicated that, although customers who use documentation are generally satisfied with its quality, many users do not turn to such documentation to solve printer problems, perhaps because they are unaware of its existence or because it is not readily available. This view was somewhat supported by troubleshooting staff at a Welcome and Support Centre we studied. Whilst applauding the quality of the aforementioned CD-ROM they expressed scepticism about its use. In their view customers typically receive a number of CDs these days with any device like a printer. These CDs are understood to be, and usually oriented to as 'technical information' and so understood to be of little relevance to the ordinary user. Therefore one of the primary problems here might be getting people to take such materials out of their case in the first place.

3.3.3 Online Resources – the Support Knowledge Base

Online databases provide the clear advantage that they can be updated on a regular basis, as opposed to paper or CD-ROM documentation. Typically accessed through the Web, they may be easier to query than paper or CD-ROM documents, but they are often less convenient at the point of need (i.e. accessible whilst right in front of the machine). Print devices and PCs with network access are rarely co-located, especially now that the emphasis is upon shared MFDs rather than individual, stand-alone devices.

We have studied a Support Knowledge Base of a printer supplier. In it the user can choose from a number of device families and eventually select the precise model of interest. Once the device model has been selected, online help is provided either from menus or by a natural language query search.

On the national (non-English) Web sites the search tool accepts queries in the national language for some classes of device. There is considerable variation here and some languages, notably Nordic ones, are not currently supported. The online support resource is at its most comprehensive in English. A search by device fault code is also available, and there are FAQs (Frequently Asked Questions) which provide solutions to commonly occurring issues.

We have studied the use of the support knowledge base for troubleshooting device problems, and have identified a number of issues which impair its efficacy. Not least of which is the fact that the user needs to have a networked computer to access the knowledge base, and this may not always be available close to the device in question. Indeed, it is not clear that users are by any means always aware of the availability of this resource. Other issues relate to the fact that a lay-user's language may not always map well to the terms used in the knowledge base.

3.4 Improving the Efficiency of Technicians in the Field

When a fault cannot be handled by the troubleshooters and requires engineer intervention the details of the fault are keyed in and the system is automatically able to dispatch this information to the correct Regional Centre on the basis of the machine's serial number. There is then a database of reported faults on a server at each regional centre where it is someone's responsibility to re-dispatch information to more local groups. Technicians at the local centre may attempt to resolve the problem over the phone, but otherwise a site visit by a support technician will be organised. Technicians have a policy of always phoning prior to making a visit and, in many cases, they will attempt to troubleshoot the problem with the customer, even if some troubleshooting has already happened.

3.1.1 Eureka

Eureka is a database of "tips" which can be used by Customer Support Engineers (CSEs) to help them resolve problems they encounter in the field. Its relevance to TCO of devices is that it has been very successful in improving the efficiency of CSEs in their role of resolving printer problems at customer sites.

The Eureka system grew out of a PARC research project [1] which studied the work of technicians in the field, and has since gone on to consider support call centres. The preferred solution of a "tips" database resulted from observations of how field technicians often help each other by sharing tips as to how various problems can be tackled.

The Eureka system is also used in the Welcome and Support Centres (WSC) at the second level of support in their escalated hardware and software groups that are primarily there to underpin the work of the service engineers.

3.5 Maximizing the Efficiency of Telephone Support to Customers

In cases where customers are not able to solve the problem themselves they can, depending upon their contract, turn to the next level of support, which is usually a telephone support centre.

The WSC is the unique interface for customers seeking help to resolve device problems. In order to obtain a first hand understanding of the workings of a Welcome and Support Centre, we carried out a field study of a WSC whose structure is depicted in Figure 1.

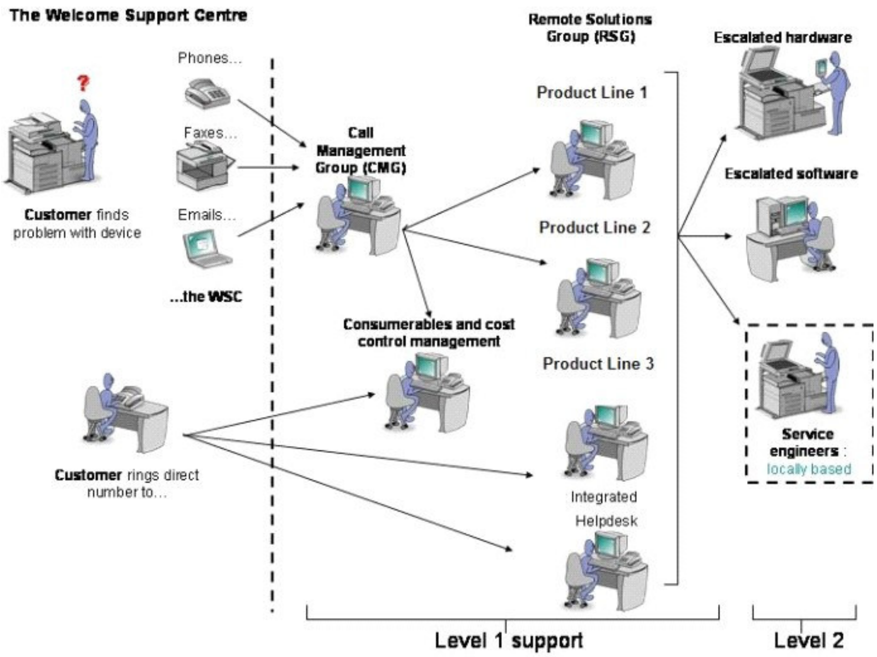

Figure 1: The Organization of a typical Welcome and Support Centre.

This study revealed a number of shortcomings in the troubleshooting process using telephone support, mainly concerning the difficulties in arriving at a shared view of the situation concerning the afflicted device between the user and the remote technical support operator. Other key difficulties relate to engaging the customer and making the troubleshooting effective in the face of possible resistance. It is also the case that the WSCs are relatively young and face ongoing problems such as how to escalate from one service level to another most effectively and managing a proliferation of resources that are not yet fully consolidated.

4 Analysing how to reduce Break-fix Cost and where areas for improvement are

The work of troubleshooting sits at the heart of every technology suppliers provision to its customers. In addition, it is tractable to technological support and development because the existing provision already incorporates a range of technologies, unlike some of the more abstract components of TCO. This, added to opportunities for improvement with regard to certain elements, has made break-fix cost reduction strategy of prime interest to us as researchers.

As seen earlier a central component of an approach to the reduction of break-fix costs can be to put a greater onus upon the customer for repair. Conventional support through engineers is expensive so the focus has to be upon the provision of online support for self-troubleshooting and troubleshooting over the telephone. In order to gain a proper insight into how this kind of support provision is realised and just where the current barriers to its effective implementation might be we conducted a range of studies of both online and telephone support. Ways to overcome these barriers are then able to stand as recommendations and opportunities to provide enhanced technological support.

User-testing consisting of troubleshooting a diverse set of problems on a range of different devices using the online knowledge bases, revealed a number of issues. On the positive side the printer interface helps users to disambiguate some instructions, the step-by-step guides are reasonably effective, and alternative options to troubleshooting are easy to locate. However, we also found that:

* The need to access troubleshooting advice on a PC is problematic because PCs and printers are not often co-located.
* The correct location of the online support needs to be easier to find.
* Users without deep technical knowledge come to the knowledge bases with symptoms, but the knowledge bases are so constructed as to oblige users to try to interpret what technical problems the symptoms might be indicative of. In addition the causes of symptoms typically need to be deduced from the problem resolutions that are proposed.
* Results of searches on the knowledge base are hard to refine and it is often difficult to uncover their relevance to the original query.
* The style and presentation of the knowledge bases is not consistent.

A field study of a printer supplier Welcome and Support Centre, observing the work of troubleshooting over the telephone, uncovered a similar range of issues that need to be handled in order to make the interaction between a customer and a troubleshooter successful:

* As with the direct use of the knowledge base by a customer, the problem descriptions provided by customers are initially symptomatic in character. The troubleshooter and the customer therefore have to work together to constitute the exact nature of the problem and its possible provenance before some solution can be seen to be appropriate. This takes interactional work that is wholly outside the knowledge base. Troubleshooters therefore have to mediate between the customer and the knowledge base to uncover what solution might be appropriate in this case.
* The step-by-step instructions on the knowledge base are by necessity devised to cover multiple possible circumstances of use rather than any particular one. However, the troubleshooters have to do considerable work to situate those instructions in the actual context of their interaction with the customer. Instructions have to be tailored to what customers report they are doing now and modified on-the-fly to make them relevant to what the customer can see. This has to be done progres-

sively and involves many spatial instructions and descriptions of parts beyond the content of the knowledge base.

* Problems have to be situated in relation to prior events. Relevant events have to be uncovered during the course of the interaction. The customer record is always partial because it has little detail of how the problem has become manifest now. The customer has this detail, but often lacks the longer history of problems and action undertaken with the device. The fact that the history might be relevant does not suggest itself automatically. Instead, the value of turning to the history has to be occasioned, regardless of whether the history may have some bearing upon the fault being experienced now.
* The absence of mutual access to the ailing device causes difficulties where the resolution turns upon getting adequate and relevant customer feedback and upon finding other ways to visualise the ways in which a problem has presented itself and how to proceed. Particularly, troubleshooters often 'mime' the actions the customer must engage in or else turn to actual machines to embody a solution by recreating it for themselves. If troubleshooters must mediate between the customer and the knowledge base, it is also the case that customers must mediate between the troubleshooter and the machine.
* Troubleshooters have similar problems with the search functionality on the knowledge base to ordinary users. Effective use therefore turns upon their greater technical competence, prior experience of what does and does not work, and capacity to spot blind alleys and do further searches quickly. Here too, the troubleshooters have to mediate between the user and the knowledge base to make the work effective. 'Effective use' is not, of course, acquired instantaneously but rather involves troubleshooters in learning the 'tricks of the trade'.
* The knowledge bases are far from comprehensive and a variety of other resources – menu maps, pictures, manuals, price lists, etc. – have to be woven into its use, often entailing opening additional browsers or turning to hardcopy.
* Sometimes the best source of additional information is one's colleagues. Troubleshooters therefore occasionally extend the interactional cohort. This is done in order to overcome either an absence in the knowledge base or trouble locating something. It is also done to draw on local knowledge of particular customers or of particular kinds of problems with particular machines.
* Troubleshooters also have to be able to judge when 'enough is enough' and terminate a troubleshooting session. This turns upon registering interactional evidence of frustration or difficulty. It also turns upon being able to recognise when the call has gone on for long enough for any longer to effect the service level. The termination also has to be handled in an appropriate way with a promise of additional action. The knowl-

edge base is not able to assess and make use of these kinds of judgements and shape its offerings to the particular situation of use.

* Troubleshooters have to make use of the knowledge base in interaction. Detailed reading of text is hard to accomplish at the same time as talk so strategies have had to evolve to facilitate this. These include searching whilst the customer is talking or otherwise occupied, anticipating the knowledge base with 'standard questions', and making use of top level lists of solutions as visual mnemonic devices. More direct interactional strategies are to put the customer either directly on hold, or to put them on 'interactional hold' by asking them to 'give them a moment'. Both of these methods are accountable if they go on for too long.
* Many customers are not immediately willing collaborators in a troubleshooting episode. Troubleshooters therefore have to do interactional work that will encourage customers to troubleshoot anyway, and to locate ways in which to render the customer accountable if they should fail to do so. Troubleshooters are just as accountable to the customers for the production of 'good' advice - they have to provide ways for the customer to be able to see that what they are being asked to do is reasonable. This is not something the online resources are currently able to do.

Consideration of the above observations indicates three key issues:

* There is currently a separation between the site of the problem and site of problem resolution. This is a barrier to the use of both the online resources by individual users and to the interaction with troubleshooters. It should not be under-estimated and will not disappear in the face of unrealistic beliefs that users will take portable PCs or phones to the device.
* The knowledge base has a number of usage issues and currently technical support do a lot of work to make the knowledge base work. Thus any expectation of migrating more and more users to the online knowledge base may not be realistic as user's experiences may deter future engagement with the knowledge base.
* It is unlikely that all problems will be solvable for all users purely through user-machine interaction. The work of technical support could be better facilitated thus improving the efficacy of the operator-customer interaction.

5 Conclusion

In this paper we have presented how the concept of TCO unfolds when applying it to the case of printer manufacturing companies. Of the many attributes that can be referred to TCO, we have further selected the one of support cost and presented the variety of dimensions where such cost could be affected. Finally we have further focussed our analysis by presenting areas for improvement in those support dimensions, on the base of some field studies that have highlighted the multiplicity of areas where one could intervene.

6 References

1. Bobrow, D G, and Whalen, J, 2002, "Community Knowledge Sharing in Practice: The Eureka Story", *Reflections*, Published by the Journal of the Society for Organizational Learning and MIT Press, Volume 4 Issue 2, Winter 2002
2. Cappuccio, D, Kirwin, B, *et al.*, 1996, Gartner view: PC LAN costs: best practices can lower your total cost of ownership and help stop the budgeting black hole, *CIO-Framingham*, Vol.9, No.17
3. David, J S, Schuff, D, and St. Louis, R, 2002, Managing your IT total cost of ownership, *Communications of the ACM*, 45(1), 101-106
4. HP MPS, 2002, Calculating the True Cost of Printing: HP White Paper, Hewlett-Packard Company
5. Gartner Consulting, 1997, TCO Analyst: A White Paper on Gartner Group's Next Generation Total Cost of Ownership Methodology, Gartner Group Inc.
6. Lanier, 2004, http://www.lanier.com/page.php/docutivity, Lanier Worldwide Inc
7. Lazar, C W, 1968, 'Lease/buy decisions for computer acquisition under conditions of uncertain technological change', in *Proceedings – 1968 ACM National Conference,* 685-690
8. Lundy, J, 2001, "Rightsizing Output Fleets: The Hidden Goldmine," Gartner, Inc.
9. Tolmie, P., Grasso, A., O'Neill, J. & Castellani, S. (2004) Supporting Remote Problem-Solving with Ubiquitous Computing: Research Policies and Objectives. workshop paper for 'Giving Help at a Distance', Workshop at the Sixth International Conference on Ubiquitous Computing, Nottingham, September, 2004