



The Work Flow of System Administration

Prepared For:	General Release
Author:	Geoff Halprin
Reference:	SYSADMIN.Workflow-0100
Version:	V1.00
Date Created:	16 May 1997 12:03
Date Modified:	25 June 1998 12:26

A SysAdmin Group Technical Report

Abstract

The average system administrator's day consists of so many complimentary and contradictory tasks that they often find it difficult to describe to other people what it is that they do. This is not surprising, because a system administrator performs such a wide variety of tasks each and every day, that it is often difficult to remember what they did before lunch!

In this paper we examine the basic types of work flow performed by a system administrator, the characteristics of each of these, and then draw conclusions as to how to better manage these diverse work flows to reduce stress and improve both customer and personal satisfaction.

© Copyright 1998, The SysAdmin Group Pty Ltd. All Rights Reserved.

All information in this document is copyright of The SysAdmin Group ("SysAdmin"). This document may not be duplicated in any form (paper, electronic, etc.) except as permitted by a valid license agreement with SysAdmin.



Table of Contents

1.0	Preface.....	3
1.1	Scope	3
1.2	Intended Audience	3
1.3	References	3
1.4	Change Control	3
2.0	Executive Summary	4
2.1	System Administration is Different!	4
2.2	Organisation	4
3.0	The World of Service Level Agreements	5
3.1	Writing Service Level Agreements	5
3.2	The Standard Service Level	5
3.3	The Optional Service Level.....	6
3.4	The Discretionary Service Level	6
4.0	Quality	7
4.1	Quality Assurance, TQM and Quality Improvement Projects	7
4.2	So Who Pays For Quality?	7
5.0	The Workflow of System Administration	9
5.1	Customer Support	10
5.2	Customer Consulting	10
5.3	Project Consulting	11
5.4	Infrastructure Projects.....	11
5.5	The Dreaded Timesheet	11
6.0	The Product Perspective.....	12
7.0	Organisational Productivity	13
7.1	Flow Time	13
7.2	Organising Support Teams.....	13
7.3	Learning Organisations - Reducing Reactive Time.....	14
8.0	Personal Productivity.....	15
9.0	Conclusion	16



1.0 Preface

1.1 Scope

This document examines the workflows of IT system support services, and their impact on personal productivity, support team organisation and productivity and customer contract negotiations and support systems.

1.2 Intended Audience

This white paper is intended to be read by systems administrators and other support staff, their direct management and IT management including those negotiating Service Level Agreements with their customers.

1.3 References

[1] "Peopleware," Tom Demarco and Timothy Lister.

1.4 Change Control

Version	Release Date	Author	Comments
V01.00	25-Jun-1998	Geoff Halprin	Release for Comments



2.0 Executive Summary

Increasingly, companies are trying to squeeze the last drop of “productivity” out of every resource, be it a computer, an employee or a sub-contractor. The IT department is, of course, caught at the bleeding edge of this push, too. Departments have their budgets reduced, and companies tender under closed bids to provide outsourced IT services for the lowest price. All of this inevitably impacts on system administrators who are being asked to do more with less.

The only counter to this push is to proactively seek to quantify what will be provided, and hence form some basis for estimating the costs and resources required to fulfil a contract. In order to derive these costs, we must look well beyond the raw product or service being offered to the underlying infrastructure and organisational costs. Too many companies set themselves up for failure by only considering the raw materials and resource costs, ignoring the working environment which is essential for personal and organisational productivity.

In this paper we examine we examine a wide variety of topics seeking to gain a better insight into the role of systems support, and how to deliver those services to customers in a more controlled, quality assured, productive and profitable manner.

2.1 System Administration is Different!

The average system administrator’s day consists of so many complimentary and contradictory tasks that they often find it difficult to describe to other people what it is that they do. This is not surprising, because a system administrator performs such a wide variety of tasks each and every day, that it is often difficult to remember what they did before lunch!

Against this backdrop, it is clearly essential that the system administrator - the one person who understands (at least to some degree) the nature of their role - take a leading hand in creating an environment which is conducive to productivity.

In this paper we examine the basic types of work flow performed by a system administrator, the characteristics of each of these, and then draw conclusions as to how to better manage such diverse work flows to improve productivity, reduce stress and improve both customer and personal satisfaction.

This paper seeks to bring to the fore the aspects of system administration that tend to go unnoticed, such as the psychological factors and workflow inherent in the role. We can then consider ways to improve productivity in light of this knowledge. The result is a wholistic approach to system support, where the total value far exceeds the sum of the parts.

2.2 Organisation

We begin our journey with a look at support organisations and how they define their relationship with their customers; the SLA. We introduce the three externally initiated workflows which result from such an agreement. We then take a quick tour of quality systems, how they impact on the relationship, and the fourth (internally initiated) workflow that they introduce.

Next, we examine each system administration workflow in more detail, how to manage support activities, and organisational productivity. Finally we turn our attention to individual productivity.

3.0 The World of Service Level Agreements

The basic mechanism for the communication of requirements and expectations between an IT support organisation and their customers is the Service Level Agreement (SLA). This document defines what the products are that the support organisation has agreed to provide to the customer (including service “products” such as desktop support), and the expectations of the customer in relation to those products and services.

The essence of the SLA is three-fold; to enumerate the services and products which are to be provided, to define the costs associated with the provision and support of these, and to quantify Service Level Objectives (SLOs) which the customer expects to be met in relation to these. The SLA can be crafted in varying degrees of legalese, depending upon the relationship between the provider and the customer.

3.1 Writing Service Level Agreements

A single SLA is often not desirable. It is difficult to craft one SLA which applies to the entire client organisation. A good rule of thumb is that an SLA should cover at least 80% of the target user base. If you find that the SLA you are drafting does not meet this measure, then you are probably looking at two or more distinct user communities, each of which should have their own SLA.

A good SLA should provide for three basic levels of service; *standard*, *optional*, and *discretionary*.

- A **standard** service is generally provided at a fixed charge (per seat), and so it is important that we (the support organisation) understand exactly what has been agreed to be supplied at this fixed charge, because we must absorb any unforeseen difficulties in providing this level of service.
- An **optional** service is one which is provided at a customer’s request. There will always be specific products which have been bundled up, but are not required by all customers, and can be cost-effectively provided on an individual or group basis. An optional product will usually consist of an initial cost component and an ongoing support cost component.
- A **discretionary** service is one where the customer requires special one-off consulting services (or a one-off product installation), and these are tracked and provided on a time-and-materials basis.

3.2 The Standard Service Level

Often, we can define a base level of functionality, a Standard Operating Environment (SOE), which we agree to provide to all desktops, or all desktops of type X. We may be required to define several classes of desktop; e.g. general, power-user, developer. These definitions may entail a particular hardware profile and software profile, with associated lease and monthly support costs. Thus, the standard support level may actually provide for several sub-levels of equipment.



This standard service level will cover all aspects of the SOE and associated services, such as desktop support. The related SLOs may include such things as:

- The defined environment will be provided with 98% end-to-end availability.
- Help desk requests will be responded to within 30 minutes.
- 95% of all help desk requests will be resolved within 5 working days.

These are all measurable objectives. In order to craft an appropriate SLA, you will need several months of statistics which allow you to arrive at appropriate figures. It is important, although not necessarily easy, to derive appropriate metrics for measuring service. (What do you mean you don't have this sort of information?)

3.3 The Optional Service Level

An optional product denotes an approved variation from the standard SOE. The process of approval will involve an investigation of the costs of providing and supporting this variation so that it is an understood cost to the support organisation. For example, a number of users may require project management software. This, in turn, will require additional support resources, which may involve hiring a suitable person, training existing staff members, or contracting this support service out to a specialist company. The number of users requiring this service indicates that we should offer it as an optional product, and the analysis of our support costs will allow us to price this option accordingly.

3.4 The Discretionary Service Level

A discretionary service is really the catch-all consulting basket. Too often, system support people are called upon to provide services well beyond the scope of their normal duties. This time is often just given away, at great cost to the support organisation.

Every time a PC has been altered from the SOE in some unauthorised way, it greatly increases the cost of desktop support. Someone investigates a problem on the PC, only to spend wasted time because their assumptions about the presence of the SOE are not correct. Their investigation and resolution times are extended, and fewer help desk calls are resolved each day.

Often, there is no timesheeting system in place, and no way to track this effort, so there is no way to bill the customer for this time. Thus, the cost must be absorbed into the standard level of service. This means that higher staffing levels are required to meet SLOs, or fewer trouble tickets are processed each month. All of this leads to decreased customer satisfaction.

Much of what a support organisation provides to its customer base is consulting in various forms. It is essential that this is recognised by both parties, and that a healthy consulting relationship exist (i.e one in which consulting is formally recognised, benefited from and paid for), so that these services are not provided at the expense of pre-agreed service levels.

4.0 Quality

4.1 Quality Assurance, TQM and Quality Improvement Projects

As has been recognised over the past 20 years, the old attitude of “quality costs” is, indeed, the opposite of the truth. Being in a position to assure and predict the quality of the products and services it provides is now an essential part of the competitiveness of any company. This has been recognised world over in the creation of the ISO-9000 quality assurance standards collection and accreditation bodies.

The basis of ISO quality is that any significant process that the company uses is documented so that the process is repeatable. This gives a predictability to the level of quality. The other major aspect to quality assurance in the ISO world is the maintenance of important records relating to work that has been performed (the execution of aforesaid procedures).

Another related quality philosophy is Total Quality Management (TQM). The basic notion behind TQM is that when a defect is discovered, we do not only repair the defect, but examine and repair the processes involved so that the defect cannot re-occur. This provides for the continuing improvement in quality of all products and services.

These quality systems do come at a price. There is a resource cost incurred in maintaining the quality systems described above. A direct result of ISO-9000 is the overheads of generating procedures and of record keeping. A direct result of TQM is that when a defect is found, we may need to generate a project to address the improvement of any number of systems or components to prevent this defect from re-occurring.¹

Thus, there should always be an allocation of resources to these quality improvement programmes, and these should be tracked and controlled as projects in their own right; *infrastructure projects*. I consider that most of these quality improvement programmes will be infrastructure projects, because they are about addressing problems systemically, and therefore about introducing change into a stable environment; changing the operational structure and systems (both manual and programmed) of the computer production operations environment to address the entropy and changing requirements of that environment.

4.2 So Who Pays For Quality?

This question is an interesting one. Obviously the customer ultimately pays for quality, or perhaps does not pay (their bills) due to the lack of it. The question really is, “how do we charge the customer for quality improvement projects?”

Infrastructure projects are about continually assessing and improving the quality of life for the staff, and the quality of service for the customers. Most of these projects directly benefit the customer, and almost all the remaining projects will indirectly benefit the customers. It is very appropriate, therefore, that the customer recognises this, and pays for these projects.

1. I am not intending to turn this paper into a discussion of the relative price of repeatedly fixing defects vs. maintaining the quality systems that prevent them from occurring. But, for the record, I believe fixing things properly is both far cheaper financially, and far more valuable in terms of staff development and stress reduction.



At one extreme, this could be done on a per project basis. A business case would be drawn up, and customer approval sought for each such project. I believe this approach would be a big mistake. It gives the customer too much power over decisions that they cannot understand the implications of. It also sends out exactly the wrong signals about quality.

At the other end of the spectrum, we do not ask the customer at all. We build a suitable percentage mark-up into all services to cater for the R&D or infrastructure component. This recognises that R&D is a vital component of remaining competitive, and gives it a budget from which we can prioritise and allocate projects. I believe this gives us an appropriate level of balance, forecasting, visibility and control. To take this approach, however, introduces a fourth workflow; infrastructure (quality improvement) projects.

4.2.1 Why Should There Be a Continuing Allocation?

Systems Management is the ongoing fight against the god of entropy. A system left alone will degrade exponentially over time. Thus, it is vital to maintaining an existing level of quality (let alone quality improvement) to regularly analyse and review existing production management techniques and decisions in the light of this continually changing environment, emerging technology and products, and changing customer conditions and requirements.

As more companies look towards outsourcing as some sort of panacea, and these same companies negotiate Service Level Agreements (where previously none had existed), the onus of maintaining a given level of service in a continually changing environment must be addressed.

It is universally accepted that every business should have an R&D budget. Why should IT companies be any different?

5.0 The Workflow of System Administration

At its most simple, there are two types of workflow; *reactive* and *proactive*. With respect to the reactive part of our workflow, it should be our intention as an organisation to (a) minimise the reactive element, and (b) to control what reactive element must remain. We will always be called upon as things break, so we must recognise that this reactive workflow is a critical part of our business, and treat it with due respect, but we should still be able to predict and control this element.

The history of systems administration as a profession means that many sysadmin groups tend to consider their sole role as being reactive. Certainly, a support role in most organisations consists of reacting to help desk calls, with an occasional upgrade. Understanding the need to limit and control the reactive work elements is a vital step towards altering the corporate culture and gaining control over the systems which you manage.

A further analysis will show that there are two distinct types of reactive workflow; troubleshooting and consulting. Troubleshooting consists of standard support for user and system problems, and consulting support is for where it has been deemed that the problem being investigated is outside the SLA, but still requires urgent attention. Note that both of these workflows are reactive in their nature. We can, hence, only statistically analyse and predict resource allocations for these workflows.

Clearly, given that the vast component of our reactive support is standard system support, our efficiency in dealing with this workflow, and whatever systems we put in place to reduce the amount of time spent in this mode will directly affect our ability to be pro-active, and ultimately allow us to be more effective with less resources.

Turning to proactive workflow, we again discover two workflows which can be so characterised. The first of these is consulting workflow of a project nature. This includes work commenced at the request of the customer and also work we initiate such as system upgrades, additional products, etc. This work can be proactively scheduled based upon customer priorities. This is day-to-day proactive work.

The second type of proactive workflow are the infrastructure projects. These are projects which are found to be necessary by internal quality assurance and improvement processes and programmes. They are proactive, and can be scheduled as such. The creation of tools to automate sysadmin functions falls into this category.

So, in summary, we have four workflows within an IT support organisation:

1. Customer Support - reactive. Standard service level.
2. Customer Consulting - reactive. Out of scope of SLA. Discretionary service level.
3. Customer Project Consulting - proactive. Project work at request of customer. Discretionary service level. This includes product evaluation and product installation (on the way to production support).
4. Infrastructure Projects - proactive. Always initiated by an internal requirement. Many of these will be to directly improve the quality of service for the customers.



5.1 Customer Support

A large proportion of time and resources are, of course, spent in customer support. That is, a customer has a problem, and contacts the help desk. Support personnel attend to this, hopefully bringing about a successful resolution.

I avoided the blanket statement that “the majority of resources” are dedicated to this support service, although I suspect that is how most people read it. I think you may be surprised how much time is spent in the customer consulting role, once you have a mechanism for tracking this.

The customer support workflow is entirely reactionary. Our goal is to reduce the time spent in this workflow as much as possible, through standards, processes, technology, and training. Some techniques include;

- Review recent trouble tickets looking for trends. These might suggest (a) training, or (b) a flawed process, or a requirement for a new process.
- Review a particular process which led to a number of trouble tickets.
- Automate a procedure which has been producing inconsistent results, or taking too long to execute.
- Run a regular introductory class for new staff members (customers), covering many of the topics which you have found you are required to provide assistance on.
- Publish a newsletter and/or create an on-line newsgroup to distribute helpful tips, overviews of upcoming products and improvement projects, and other information of value to your user community.

The common thread through many of these is that you have implemented a trouble-ticket system which allows you to maintain history and statistical information on support calls.

5.2 Customer Consulting

Whatever the environment (SOE) that is implemented as a base functionality, there will be many needs for variations from that environment. Many of these will be approved variations (optional products), but a significant number will be one-off customisations. It is important that you do not provide support for these variations free of charge - this defeats the whole purpose of an SLA, and will result in a more diverse environment which grows in an uncontrolled way; the ultimate result being the exact chaos you were attempting to address.

There are two distinct ways in which customer consulting is initiated;

- In response to a help desk call, when a problem has been identified as being outside the scope of the SLA (during problem investigation); and
- At the request of a customer, in relation to a problem or proposed technology implementation. This is discussed shortly.

When responding to a help desk call, the moment you have assessed that the system in question has been altered from the baseline without approval, it is time to start the clock. Likewise, when you’ve been called to a customer who actually wants to discuss matters outside the scope of the SLA (or the reason you were called), it is time to start the clock. (As always, a healthy dose of discretion is vital to maintaining a healthy customer relationship.)

Before commencing a consulting project, be sure to get the customer to sign an acceptance form that they will pay for the consulting - an “authorisation to proceed”. This will prevent a lot of problems when it comes to getting them to pay their bills.

5.3 Project Consulting

As your relationship with your customer develops from merely being a help desk to being a technology partner, your opportunities to provide more meaningful consulting will further increase your consulting operations.

This workflow is differentiated from the previous one, because it is proactive in nature - you have warning of this need and can schedule this work appropriately. Also, it may be of a significant nature, requiring a project plan and the scheduling and coordination of other resources.

5.4 Infrastructure Projects

Infrastructure projects are about continually assessing and improving the quality of life for the staff in the support organisation, and the quality of service to its customers. Most of these projects directly benefit the customer, and almost all the remaining projects will indirectly benefit the customers.

These projects include reviewing business processes, investigating and implementing software support systems such as a help desk system, automating system functions and key processes, reviewing system administration practices, benchmarking, and evaluating new technologies.

It is important that these are treated as projects, and all support staff given an opportunity to contribute to these projects where they can.

5.5 The Dreaded Timesheet

As may have become clear by this stage of the proceedings, the time sheet is an essential tool for the accurate tracking and billing of time to the customer. If the support organisation is to charge their customers for consulting effort, then they must be able to substantiate those charges. The signed timesheet is the primary method of achieving this.

Moreover, it can also assist in building up statistical work profiles for an individual and the organisation, and hence assist in better understanding and improving services and staff productivity. It can help clarify whether the SLA is being effective as a support structure, and many other statistics which are of value to the organisation.

It is essential to design a timesheet which has a minimal impact on the support staff (and hence is most likely to be filled out correctly), whilst ensuring the accounting and statistical information is properly maintained.


6.0 The Product Perspective

As we have seen, there are four basic categories of workflow which we undertake as an organisation. What has not been made clear are the interactions and movement between these categories, and what links these categories together. For this we need to take a product perspective on life.

We supply our customers with a number of products. Each product includes all three SLA components (standard/fixed, optional, discretionary). In order to supply a new product, we must kick-off a proactive investigation project to learn about the product which we are about to support, and how it interacts with other products we support.

Once the product has been placed into a production environment, we have reactive support and reactive consulting which we perform in support of that product. So, in essence, the SLA defines three classes of *service* which we provide for each *product* we support. The fourth service class (infrastructure) represents internally run projects, and so is not reflected directly in the SLA.

Each product runs on a supported *platform*, so we must also reflect all the indirect platform support costs. These, too, fall into the same four workflow categories. The diagram below illustrates the relationship between products, services, platforms, the service level agreement and workflow categories;



When a new service is introduced on behalf of a customer, the initial work in analysing and integrating the product into the existing production environment should be charged to the customer as discretionary project work.

Each product introduced will have costs associated with it in terms of fixed support costs. If the product is optional, then there are two cost components; the discretionary install component, and the fixed support component.

Each product we introduce mandates one or more platforms, and there are the indirect costs associated with management of each of these. These must also feed back into the chargeback scheme of the SLA. In this context, the platform costs include networking hardware, site costs such as power and rent, and the actual platform management costs (e.g. server lease and disk space).

7.0 Organisational Productivity

Now it is time to look at the problem from the other side. We must now look at support from inside the mind of a system administrator - the psychology of system administration. This is, by no means, meant to be a definitive analysis of how system administrators think, but is intended to provide insights which may help us in establishing a productive environment for support staff to thrive.

7.1 Flow Time

The following (slightly abridged) quotes come from Peopleware [1]:

TABLE 1. How Developers Spend Their Time

Work Mode	Percent of Time
Working alone	30%
Working with one other person	50%
Working with two or more people	20%

During single-minded work time, people are ideally in a state that psychologists call *flow*. Flow is a condition of deep, nearly meditative involvement. In this state, there is a gentle sense of euphoria, and one is largely unaware of the passage of time.

Unfortunately, you can't turn on flow like a switch. It takes a slow descent into the subject, requiring fifteen minutes or more of concentration before the state is locked in. During this immersion period, you are particularly sensitive to noise and interruption.

The effects of this understanding are many and significant. They include:

- Don't mix work flows in the office environment. Keep those who must react for a living (the help desk) away from those who must focus (project teams). Open plan offices, with front-line support staff talking on the phone all day can destroy any chance of productivity for project teams located nearby.
- Organise offices (where they are available) to be either (a) dedicated to a single person, or (b) shared by those working on the same project. The times they work alone are generally shared, as are the times they need to speak with each other.
- If you have open plan offices, set aside a couple of full height offices which can be booked, like meeting rooms, when people need to get away.
- Give people phones with "DO NOT DISTURB" buttons that work. All too often, the DND button doesn't block intercom calls which blurt out over the top of whatever you're doing.

7.2 Organising Support Teams

A typical support structure is a three-tiered help-desk model, where each tier filters out 80% of the calls, and the remaining 20% are escalated to the next level. The first two tiers are typically dedicated help desk personnel, and the third tier are the platform operations staff (e.g. UNIX operations team).



The first two tiers in such a structure are entirely reactive, their only responsibility being to resolve issues as they are assigned them. This makes these positions repetitive, reactionary and career limiting. Once an initial learning curve has been mastered, there is relatively little to learn or contribute, and challenge must be sought elsewhere. This directly contributes to the turnover of staff in these roles.

By contrast, the third tier must perform reactive help-desk related support, system-induced reactive work (e.g. a system crash), and proactive work such as upgrades and installations. In their spare time, they're allowed to automate things to make their jobs easier.

The only transition between these roles is a sharp cutover from one to the other, for those who show promise. This presents them with an entirely new, steep learning curve.

Splitting people into roles based upon workflow is a mistake. Having someone solely responsible for second tier support creates a culture of being reactive and never doing anything that isn't in reaction to a help desk call. This, in turn, tends to breed a lack of follow through and a lack of vision of the broader picture.

Another problem is that the third tier, operations team appear to impose arbitrary standards and procedures from above. This means that the very people who are expected to follow these procedures have had no involvement in their development, nor intellectual sign-off on them, and so often do not understand and resent such efforts.

The resulting culture is one of organisational islands, with little or no inter-group support, and a degraded support organisation.

Cross-pollinate the support team. Rotate first-tier help desk staff through second-tier roles. Allow second-tier staff to contribute to the operations team and to projects. Round table problems with the whole team at regular meetings. Allow everyone to be proactive. This provides everybody with a path for advancement, a sense of achievement, and a different viewpoint to consider. This also breaks down the "us vs. them" mentality.

7.3 Learning Organisations - Reducing Reactive Time


It is vital that every IT support organisation implement a comprehensive trouble-ticketing and problem management system. This system must do two things:

- Maintain an accurate history of all activity on each trouble ticket (resolution diary); and
- Produce statistics on the tickets, and the resolution process.

It is preferable if the system supports automatic escalation of trouble tickets which have not been attended to within a specified period.

Part of the ongoing process of improving quality is the regular review of all trouble tickets, looking for patterns. These patterns usually indicate a deeper problem which requires, or would benefit from, a more systemic solution. Regular support team meetings should be held to perform this review, and to look for infrastructure projects which would have a significant impact on operations and/or support. These projects, once identified, must be evaluated and approved/prioritised against other projects and the R&D budget.

It is important to find a champion for each such project - someone who is interested in driving the project through to completion. This significantly improves the chances that the intended results will be seen.



8.0 Personal Productivity

So far we have looked at the Service Level Agreement as a mechanism for two organisations to define the services that one will provide to the other, and what this means to a support organisation. We have also looked at how, at an organisational level, we can improve productivity. The last area we must examine is personal productivity.

The ideas presented below largely reflect the lessons that can be applied at the organisational level, just as they could be applied to oneself.

1. Avoid being solely in reactive mode. Plan a day a week for pro-active work.

Even if you're front-line support, try to get involved in something proactive such as documenting a procedure to help others.

Organise one of your peers to cover all of your calls during this pro-active time. Do the same in return. If you can only get half a day, that's still better than none.

2. Mixing reactive and proactive workflows doesn't work.

Set aside a solid amount of time of at least half a day when you need to perform a planned project. Don't let yourself be distracted with questions, phone calls or anything else during this time. If there's a spare office so that you can get away from all the surrounding distractions, even better.

3. Rotate quiet (flow) times.

Ensure everybody in the team gets a chance for some flow time.

4. Meet fortnightly to blue sky.

Have the entire platform team (including help desk staff) get together once a week or fortnight to discuss problems and ideas for improvements.

5. Automate something, even if it's in your spare time.

The act of automation does two things; (a) it reduces your workload because that task will never take as long to do again (and the results will be more consistent), and (b) it gets you thinking and working proactively. This gives you a reason to go into work each day (to find time to work on your pet project), and leaves you with a sense of achievement when you complete the task.



9.0 Conclusion

In this paper we have looked at a wide variety of topics, from an external perspective, from a management perspective and from a system administrator's perspective. The results of all of these suggestions is a wholistic approach to system support where the customer appreciates what system support staff do, and system support staff have the best chance at long-term success and productivity.

I have tried to provide suggestions at the highest levels (contract negotiation) through team management to personal time management.

System administration is, ultimately, about knowing when to take a step back and see the bigger picture. Being at the coal face of system support it can often be difficult to apply the same sky-like attitude to yourself and your profession as you do to your customer's problems. I hope this paper has assisted people in that capacity.