

The Yearly Review, or How to Evaluate Your Sys Admin

Carrie Gates and Jason Rouse – Dalhousie University

ABSTRACT

While some work has discussed hiring system administrators, and other work has focused on the technical and mechanical requirements for terminating a system administrator, there has been very little published regarding how to review or evaluate a system administrator. This paper presents one approach to doing such a review, followed by scenarios that explore the approach. The system developed in this paper has the aim of creating measurable goals that a competent system administrator should be able to achieve. We also discuss when the use of this model is appropriate, its strengths and weaknesses, and the responsibilities placed on management if this model is used.

Introduction

There are several publications that talk about how to hire a system administrator (e.g., [5, 6]). In these cases, the emphasis is on how to determine the applicant's problem solving ability, general knowledge, and fit within a company. There are fewer publications dealing with how to fire a system administrator (e.g., [5, 7]). Those publications that do discuss how to fire a system administrator concentrate on the technical aspects: how do you ensure that the system administrator no longer has system access, and that there are no backdoors, for example.

However, there is no previous work that discusses how to evaluate the effectiveness of a system administrator. The learning of system administration seems to be based largely on the apprentice system, with the implicit assumption that junior system administrators will learn from senior system administrators. It is further assumed that senior system administrators, by virtue of years of experience, are competent. But what if the junior administrators are taught poor practices? Unfortunately, it is not always the case that the senior administrator is competent, regardless of years of experience. Nor is it always the case that the system manager is himself an administrator, and so knowledgeable of the area and able to teach or evaluate his staff. Thus, guidelines need to be developed to assist in evaluating the effectiveness of system administrators.

This paper presents one approach to the evaluation of system administration personnel. We believe that there are three main criteria against which system administrators can be evaluated: achievement of goals (e.g., installing and deploying a back-up system), achievement of a specified service level (e.g., that the time between a user request and fulfillment of that request is less than some specified value), and general competence.

The first criterion focuses on the development of work-related goals that are mutually acceptable to the administrator and manager, where performance is later measured against the achievement of these stated

goals. This emphasizes coordination and cooperation between administrator and manager, and is geared towards allowing non-technical managers to understand and evaluate administrator performance in an objective manner.

The second criterion concentrates on meeting "standard" service levels. Service levels include minimizing unscheduled downtime and having a guaranteed response time for users. In this section we define four main components – availability, usability, security and customer services – and outline suggested service levels and measurements for each one.

The third criterion deals with the problem solving and general competence of the system administrator. We believe that the best form of system administration is to solve problems correctly the first time, rather than continuously "hacking" a system, as this latter approach leads to later problems with various services. It also compounds the complexity of fire-fighting and can make third party trouble shooting nearly impossible. Thus, the third portion deals with measuring how often an administrator fixes the same problem.

It should be noted that the ultimate goal of this paper is to start discussion within the community on how to evaluate system administrators. To stimulate this discussion, this paper presents some guidelines in the formation of an evaluation system, as well as a proposed system that as much as possible tries to quantify the art of system administration.

Criteria

Before addressing the specifics of system administration, it is appropriate to consider performance appraisal in general. While it is often the case that a manager, who may or may not be familiar with the details of an employee's job, must evaluate that employee, this process *should* be one that involves both manager and employee. It is important that evaluations be performed in an objective manner that evaluates the work performed, and the quality and outcomes

of that work, and not the personality of the employee or the personal biases of the manager. In addition to avoiding potential legal problems from a subjective review, this also provides the best approach to a review that is *fair*.

Approach

The ultimate goal in the design of this evaluation procedure is fairness, to the employee, the manager, and the organization. We feel that the best way to attain this goal is to evaluate strictly on the work performed, and not on any of the more subjective criteria (e.g., how well the employee gets along with others, whether the employee is cooperative, and the like).

How can a system administrator's work be evaluated? System administration can often seem like more of a black art than a definable job. And no two days in an administrator's life might ever be the same, but instead consist of multiple tasks that are often unpredictable, making concrete evaluation trickier. Perhaps the most important skill of any system administrator is problem solving, yet this skill also comes from experience: experience with similar problems, with similar software, and with the organization of the system in question. There are, sadly, no standardized or widely accepted tests available that will grade an employee's problem solving skills, let alone these skills in relation to system administration. Further complicating the matter is the distinct possibility that the manager evaluating the system administrator may not be familiar with the systems for which the administrator is responsible (e.g., the manager might have a Windows NT background while the administrator works with Solaris and AIX), nor is there any guarantee that the manager has ever been a system administrator.

To address these issues, this paper adopts a three-part approach that is centered around the ideal of administrators and managers working together. The first part adopts the concept of goals that are developed by both the administrator and manager. In this case, the administrator is measured by his progress toward goals he helped to set. The second part presents four key components of system administration work and suggests relevant measurements within these components. The third part attempts to measure how effective a system administrator is by measuring how much time he spends revisiting problems.

It must be stressed at this point that the evaluation procedure requires communication between the manager and system administrator. A manager cannot simply tell an administrator that they are about to be evaluated without having previously explained the criteria they are expected to meet. Similarly, a manager should not be kept in the dark regarding current or potential problems on which the system administrator is working.

Goals

The first part of the three-part approach presented here draws heavily on the suggestions made by King [2].

King defines "performance plans" where the manager and employee work together to plan for the coming year and to define what needs to be achieved. King notes that performance plans require five characteristics in order to be considered achievable: specific, measurable, time limited, realistic, and challenging.

That is, the plan must be *specific* so that both the manager and employee are clear on what is expected. For example, "keep servers up" is too vague, whereas "ensure that web site X has no more than two hours of downtime in the coming year" is more specific, and therefore *measurable*.

Time limited ensures that the employee and manager are aware of any deadlines. For example, a goal of "install a new back-up system" might be given lower priority by an administrator if the current system is still in place and working, whereas the manager might expect that the new system is online within a month and have promised as much to clients. Thus, *time limits* are required on all aspects of a performance plan.

Finally, the plan must be both *realistic* and *challenging*. A realistic plan should be obvious; it is unfair to expect an employee to complete tasks that are not possible or are unnecessarily stressful (e.g., install and configure the new back-up system, and deploy to 150 clients, by the next day). Challenging may be a little less obvious. King argues that employees should be challenged in their jobs so that they remain interested in their work and are given the chance to excel. Thus, any performance plan drafted by the manager and employee should provide this opportunity.

In this article, we use the term *goals*, which we define as having an equivalent meaning to King's definition of a performance plan. More recently, goals have received attention in the human resources literature, where they have received the moniker of "SMART" goals: Specific, Measurable, Action-Oriented or Attainable or Aggressive, Realistic, Time-Constrained or Tangible [8, 1, 9]. In this case "challenging" has been replaced with "action-oriented," where it is expected that the goal requires that the individual must perform some action in order to attain the goal, rather than to define a goal and then not have any means for pursuing it.

In our evaluation procedures, we expect the manager and administrator to work together to define SMART goals. Due to the constantly changing nature of most system administration environments, it is suggested that this meeting take place multiple times per year, perhaps every three or four months, depending on the natural cycles of the environment in question. This allows for a complete view of the environment and the evolution of facilities or capacity planning, as well as enabling administrators and managers to note when goals begin to move "off the rail."

By having the manager and administrator work together to set goals, we allow the manager to ensure

that the overriding client¹ requirements are understood by the administrator (e.g., two-hour response time to client problems). It provides the manager with a chance to ensure that the goals the organization considers important are being met. Conversely, it allows the administrator to ensure that the goals are not unrealistic (e.g., 100% uptime for the next two years). The administrator also has the opportunity to make the manager aware of any issues that might otherwise be missed (e.g., a major server is starting to have hardware problems and will need to be replaced soon).

In addition to providing input to the manager that will enable him to make more effective decisions, and providing the administrator with a sense of the larger picture in terms of corporate or organization goals, having the manager and administrator work together to determine objective goals will make the evaluation process consistent and fair. For example, a manager can not simply evaluate an administrator using the criteria outlined here, or any other criteria, without informing the employee of the criteria against which they will be measured. By working with the manager to create goals, not only is the employee aware of the criteria against which he will be measured, but he is also ensured a priori that both of them agree that these criteria are reasonable.

The authors recognize that there are three types of goals: personal, professional and organizational. Personal goals involve those goals a person might have that are in no way related to his profession (e.g., to become a better painter). Professional goals encompass the development of skills which are not necessarily directly related to the organization (e.g., to learn Perl by the end of the year) but may be used to the benefit of the organization. Organizational goals are those that are directly related to what the organization requires from the individual.

While all of the goal examples previously described were organizational in nature, it is also appropriate to include professional goals in the goal-setting section of the performance review. This gives the administrator the chance to learn new skills that, while not necessarily directly related to the immediate organizational goals, will indirectly benefit the company by allowing the administrator to remain current, and by keeping the administrator happy with the company.

Service Levels

The first stage in designing evaluation criteria for a system administrator is to identify the critical components of system administration. Here, we define these general components to be availability (of hardware, software, services and data, including backups), usability (whether users are able to perform the tasks

¹We do not want to constrain this to a business environment; academic and government facilities have clients too! We use the term client in a general sense, encompassing users and lower-level management, as well as the traditional business meaning.

they need to complete), security, and customer service. It should be noted that, ultimately, system administration is a service, and as much as we complain about "lusers," we are ultimately responsible to our users.

Within these four components, outcomes need to be defined. These outcomes need to be easily quantifiable and measurable, without necessarily requiring the services of an expert system administrator. The outcomes should be easy to gather, and should follow the KISS (Keep It Simple, Stupid) principle. Further, the outcomes should also build in a "benefit of the doubt," so that it can be recognized that the administrator made a legitimate mistake or that there were extenuating circumstances. In particular, junior administrators should be given more latitude than senior administrators.

The following questions in each of the four components have been identified as meeting the design criteria:

1. Availability:
 - a. How often has the system, including hardware and key services such as web or ssh, had unscheduled downtime within the past year?
 - b. Can a file that was deleted yesterday be reliably recovered from backup? Deleted last week? Last month? From both servers and desktops?
 - c. How many mistakes has the system administrator made that directly led to system or service downtime?
2. Usability
 - a. What is the average time between a user request being made and the fulfillment of that request?
 - b. What is the average time to install a new server and have it operational?
 - c. What is the average time to install a new service and have it operational?
 - d. Identify the top 10 software and services used on a system. How far out of date is this software? Is there a legitimate reason for using old software (e.g., gcc 2.95 still required over 3.0)?
3. Security
 - a. What is the average time between a relevant security patch being released and being installed?
 - b. What is the average time between a problem occurring with the system and the administrator noticing? (Underlying question: Does the administrator monitor the system?)
 - c. Are appropriate security measures installed and monitored? For example, is there an intrusion detection system, and are alerts monitored?
 - d. Is confidential material treated appropriately?
4. Customer Service
 - a. What is the average response time when a user emails a system administrator?

- b. Does the administrator notify users of events such as expected downtime or policy changes?
- c. Does the administrator make an effort to stay current, either through reading appropriate mailing lists, taking training courses, reading books and magazines, etc.?
- d. Does the administrator follow a practice of both learning from and teaching co-workers? (Or, does he have job security through obscurity?)

In some of these cases, reasonable values need to be determined, and may be dependent on the organization. For example, the ideal time between a relevant security patch being released and being installed should be relatively short (e.g., perhaps three days). The average response time to user email might be organization dependent, where some organizations expect a one business day turnaround, and others might only require a one week turnaround.

Competency

This section was the cause of much discussion among the authors and other system administrators. How does one evaluate the competency of system administrators? At the very least, how can one differentiate between a competent system administrator and one who may need to be appropriately trained and/or disciplined, or even replaced. This is not easy. For example, we have seen systems where there were five backup copies of the passwd file going back two years in /etc, along with a directory /etc/passwd.backup containing more backups of the password file. This was

on a server that was also being backed up by two different systems: Amanda and TSM. How does one quantify system administration practices in such a manner that recognizes these practices as undesirable and unnecessary?

The one consistency that the authors could find is that poor system administration practices lead to the same problem being revisited multiple times. That is, if the system or service was installed correctly the first time, there should be minimal problems reported with the service. This is not to say that there will be no user requests. Rather, the user requests should take the form of “please install . . .” rather than “please fix . . .”!

Therefore, in order to provide some measure of general competency of the system administrator, the manager will need to follow user requests and track when those requests are due to a piece of software that was incorrectly installed or configured.

Measurement

As there are three parts to the evaluation, there are also three parts to the measurement section. The first two sections are worth 10 points each, while the third section is worth 5 points, for a total of 25 overall. The lower the overall score, the better the performance of the administrator. In the first two sections, 0 points implies that the standard (either goals or service level) was met. As it is possible to exceed the expectations for goals, it is possible to score lower than 0 points in this section. The third section is only worth 5 points. This is NOT an indication of its importance relative to the other two measures. Rather, it is a recognition of

| Service | Response | Value Range |
|---------------------|---|-------------|
| 1. Availability | | |
| | (a) Count the number of non-hardware related down-times (Hardware and power failures should not be counted here) | 0 - 8 |
| | (b) Count the number of no responses (out of 6) | 0 - 6 |
| | (c) Count the number of known occasions (Note that the manager might not know these!) | 0 - 6 |
| 2. Usability | | |
| | (a) Count the number of days | 0 - 5 |
| | (b) Count the number of days | 0 - 5 |
| | (c) Count the number of days | 0 - 5 |
| | (d) Add 0.5 for every unjustified full release difference Add 0.1 for every point release difference | 0 - 5 |
| 3. Security | | |
| | (a) Count the number of days | 0 - 5 |
| | (b) Count the number of hours | 0 - 8 |
| | (c) If no, add two points, else add zero | 0 - 2 |
| | (d) If no, add five points, else add zero | 0 - 5 |
| 4. Customer Service | | |
| | (a) Count the number of days | 0 - 8 |
| | (b) If no, add five points, else add zero | 0 - 5 |
| | (c) If no, add five points, else add zero | 0 - 5 |
| | (d) If no, add two points, else add zero | 0 - 2 |

Table 1: Suggested measurements for the service levels.

the difficulty of measuring a system administrator's competence, and of the volatility involved in the process and subsequent discussions.

Measuring Goals

The first part, measuring the achievement of the goals set throughout the year, is the easiest. Given that there are performance evaluations annually, the measurement consists of "met some goals" (10 points), "met most goals" (5 points), "met all goals" (0 points), "exceeded some goals" (-2 points) and "exceeded most goals" (-5 points). As the process of setting goals requires the administrator and manager to meet periodically through-out the year, it allows time to adjust deadlines if necessary. For example, a goal might have been to install a new server by a particular date. However, if the server arrived one month later than expected, this should be factored into the goal deadlines.

Measuring Service Levels

Table 1 provides a suggested measurement scheme for part two of the evaluation. It should be noted that this is a suggestion only, and will need to be adjusted appropriately for each organization. For example, condition 4a, the average response time when a user emails an administrator, might more appropriately be measured in hours rather than days depending upon the environment. Or condition 2b, the average time to install a new server and have it operational, might better be measured in weeks rather than days, again depending on the complexity of the environment. The table presented provides suggested measurements based on a medium-sized faculty in a research-based university. Most organizations will likely want more strict values.

The measurements for this section result in a measure out of 80, so will need to be divided by 8 before being added to the total. Each of the four components is weighted equally in this scheme. However, some organizations might place a higher premium on some components over others (e.g., security over customer service). In these cases they should adjust the weights accordingly.

Some of these questions require further explanation at this point. First, it is assumed that a tracking system is in place so that the manager is aware of items such as the time between a request being made and the fulfillment of that request. If no such system is in place, then the manager requires a more hands-on involvement in the system administration of the organization in order to be able to answer some of the questions. In any case where the manager is not aware of the answer, the benefit of the doubt should *always* be given, and so a value of 0 should be assigned. The manager should never assign a higher value without the documentation to support his evaluation. Otherwise, the manager is no longer evaluating the work performed, but rather his personal beliefs of the worker's performance.

In the same vein, items such as 3d, is confidential material treated appropriately, should not be based

on "gut feeling." Rather, a non-zero value should only be assigned if the administrator has actually performed some action that violates a user's confidentiality (e.g., printing credit card information and then leaving it in a public recycling bin rather than shredding it). Similarly, items 4b, email of notification of downtime or policy change, 4c, personal improvement, and 4d, collaboration with peers, should receive a value of 0 unless the manager has documentation otherwise. Documentation in this case might include a complaint from a user of not being informed of a system change (4b), or written confirmation from the administrator that he does not follow appropriate mailing lists or magazines, etc., (4c), or complaints from coworkers that the administrator does not share information or help with systems changes (4d).

Measuring Competency

For the third section of the evaluation, the manager needs to review the system administration tasks for the period being reviewed (this should ideally be performed every time the manager meets with the administrator regarding the goals for the period, rather than only once per year), and compare these against the user requests made, as well as the jobs put in the job tracking system. The number of times that a user has repeated the same request before final resolution should be counted. For example, a user might request a Perl module be installed, which, once installed, might be followed by a comment that the module still does not work. This would count as one, as the user needed to repeat the request once before the service was usable to the user. Similarly, if a backup system was installed, and the task was set as completed, yet job tickets followed stating that machines X and Y were not backing up properly, this would count as two: one for machine X and one for machine Y.

It is recognized that administrators will make mistakes. It is further recognized that junior administrators will make more mistakes than senior administrators. However, mistakes should be minimal as administrators should test their installations, configurations and changes before checking the job as completed in the job tracking system. The manager should determine what he feels is an acceptable level of repeat incidents.

A suggested guideline might be that incidents such as those described above should happen no more than once per month on average for senior system administrators. Junior administrators should be given more latitude, allowed to make three mistakes per month on average. Therefore the administrator receives a value of zero for this section if mistakes were made no more often than listed for his level. For each additional month of mistakes (that is, assuming a junior administrator, for every three additional mistakes per year above and beyond the 36 allowed) one additional point is assigned, up to a maximum of five points.

Usage

The score achieved on each of these three sections can be added together to provide an overall view of the effectiveness of the system administrator, where lower numbers indicate better performance. While it is tempting to provide absolute values that categorize an administrator as either good or bad, the authors feel that the manager has the responsibility of determining what are acceptable values for each of the three portions of the evaluation, as well as the overall value. The manager is also responsible for communicating this to the administrator *before* the evaluation.

The results of the overall evaluation can be used to assist the manager in determining how best to use the administrator. For example, the administrator might perform strongly everywhere except for customer service related goals and service levels. In this case, the manager might move the administrator to a less visible administrative role and have another administrator be the primary contact for users.

Similarly, the manager might also be able to use the results of the evaluation to determine what training courses might be appropriate. For example, if the system administrator consistently misses some goals, then the manager should look for some common element in the goals missed. Were they all related to AIX systems? Or were they all related to Solstice Backup Server? By searching for the common element, the manager can recognize where particular training might be beneficial. A careful review of where a system administrator makes mistakes (as noted in the third part of the evaluation) can also lead the manager to determine where further training might be appropriate.

The results from the evaluation might also show that the administrator consistently exceeds the majority of the goals set, but might not perform so well under the service level section. In this case it might simply be a matter of structuring goals for the administrator so that the service levels are also achieved. Or perhaps the administrator was not aware of the service levels as separate goals that needed to be achieved, and so did not prioritize appropriately.

Finally, this evaluation tool can be used to determine the administrator's strengths and weaknesses. This allows both the manager and administrator to work on the weakness, and to take advantage of the strengths.

Scenarios

This section describes five different scenarios, with some typical measurements, and the responses to the evaluations. It is provided to give the reader a sense of how this process can be deployed, the types of responses it might generate, and how the administrator and manager can work together to solve issues.

The Happy Helper

The first case involves an administrator in a small university, working with a team of administrators. The

university provides a number of servers and services to its students, faculty and staff.

The established goals centered around technical requirements, such as deploying a new printing system. At the end of the evaluation period, the administrator had met all goals, receiving 0 points.

For the availability and usability sections of the service levels evaluation, the administrator scored well, receiving 0 points for each. However, in the security section the administrator received 1 point for the time delay between a security patch being released and being installed, another point for the time between a problem occurring with the system and the administrator noticing, and 2 points for not having additional security measures installed and monitored. Under the customer service section, the administrator received 4 points for the response time to user email and another 2 points for not always notifying users of changes. The total number of points awarded was 10, resulting in 1.2 points (out of 10) for this service levels section.

For the third section, the administrator, who was expected to perform at the level of a senior administrator, had acquired 3 points (having, throughout the year, made 15 errors that resulted in services needing to be reconfigured). This resulted in an overall evaluation score of 4.2 (out of 25).

During the performance review, the administrator complained that he did not have time to monitor the system or respond to email because he was so busy helping the people who dropped by his office. Similarly, services were often deployed without complete testing due to these interruptions. He felt that his job included helping people, and so other areas suffered, and yet there was no recognition in the process of this service. In this case, the manager recognized that the administrator was correct (having often walked by his office and seen people in it!).

The manager agreed that helping people was part of the administrator's responsibilities. However, he felt that many of the questions could be handled by other personnel, freeing the administrator to perform more administration duties. A compromise was agreed to where the administrator would hold office hours, during which time he was available to others for consultation, while the manager would email everyone to inform them of this change.

The goals for the administrator were centered around eliminating the distractions during non-office hours by asking people to enter their problems into a ticket system or to come back during office hours, and to close his door and put down the blinds during non-office hours. Other goals related to improving user response time and service. While security was also identified as important, the manager did not want to overload the administrator with goals, thereby decreasing his chances of meeting them, so it was agreed that the manager would ask another administrator to take over the

security duties. The last goal set was that the administrator was to reduce the number of errors per year to an acceptable level. It was felt that this goal was complimentary to the first, since the number of distractions was a significant contributor to the number of mistakes made.

Mr. Job Security

The second case involves a small, for profit, Internet-based software company with one administrator. The company relies on a number of NT servers to provide a service to their customers, and so requires a high level of uptime (99.999%).

Before implementing the performance evaluations described, the manager felt that the administrator was doing an adequate job. However, he did not like the administrator's attitude. The administrator felt that he had job security, and so would not deal with user requests in a timely fashion.

For the first section of the performance review, the goals that were set for the administrator centered around ensuring service for the clients, such as bringing new servers online, and developing specifications for new servers that would meet demands, such as hot failover. The administrator met all goals set, and so received 0 for this part of the evaluation. Similarly, for the third section of the performance review, the administrator did not often revisit the same problem, and so received 0 points.

The second part of the evaluation centered around the service levels. In the section on availability, the administrator performed well, although it was noted that the backups were not always reliable, scoring 2 points here. For the usability section, the administrator also performed well, with the average time for installing servers and services being acceptable. It was noted, however, that the usual time between a user request and fulfillment of that request was three days, and that some of the software was out-of-date, resulting in a score of 4 points. For the security section, it was noted that while patches were applied on time, and there had been no issues with confidential material, there was also no initiative on the administrator's part to improve security, and so there were no intrusion detection systems, etc., resulting in another score of 2 points. Finally, on the customer service side, the average response time for email was two days and the administrator did not make an effort to stay current, resulting in 4 points. The final point regarding the administrator both learning and teaching co-workers was deemed irrelevant and so was ignored.

The final score for the service levels evaluation was 12 out of 80 (which reduced to 1.5 out of 10, resulting in an overall evaluation score of 1.5 out of 25), which was felt by the manager to be unacceptable. The manager sat with the administrator and discussed how most of the points dealt with user-related issues and general response time. The result was goals

for the following evaluation period being set that centered around user response time. It was made clear to the administrator that he was expected to achieve these goals, or else go through a discipline process.

In this case, the manager was able to articulate in an objective manner the requirements for the administrator, without resorting to statements such as "I don't like your attitude." The administrator became aware of the importance placed on timely responses to users, and was given the chance to correct his behavior.

The following year, the performance review had much the same results. There was no noticeable improvement in the response to users, and so no change to the service levels review, receiving a score of 1.5 (out of 10) again. As in the previous year, no noticeable mistakes were made, and so a value of 0 was assigned. However, the goals section had been changed from the previous year to include a number of goals ensuring improved response to user requests. As these goals were not met, a score of 10 (out of 10) was received for this section. The overall evaluation was therefore 11.5 (out of 25), as compared to 1.5 the year before. This provided the management with documented grounds for dismissal, and so the administrator was fired.

Just Plain Overworked

The next case involves a system administrator, working as part of a team, in a government office. The administrators were responsible for a large server farm, as well as user workstations and laptops, for a very large and demanding user group.

During the performance evaluation, the administrator performed well in the third section, having made few mistakes of note. For the first section, the administrator had completed all goals. However, he had not met all of them on time, having missed some of the less critical goals by a few days.

For the service levels section, the administrator scored well on availability and security, with zero points for each. However, for usability, the administrator received two points for the time required to respond to a user request, five points for the time required to install a new server, and three points for the time required to install a new service. In the customer service category, the administrator received a further two points for the delay in responding to user email.

With a total of 12 points, the manager felt that the service level score was too high, and also found it disturbing that some of the goals were missed. In general, it was felt that this administrator, along with the others on the team, was very good.

Similarly, other members of the team received scores with deficits in both the customer service and usability categories, and some of them had also failed to meet some goals on time. This alerted the manager to the fact that there was a more general problem with the system administration team. Although the team members had individually reported problems with

overallocation of work, the manager had believed that the workload was appropriate. With these new measures, however, the team was able to establish a reasonable benchmark for their output. The manager was able to identify the need for a larger team. Along with individual reports and these measures, the manager presented a balanced case to his superiors for the addition of another administrator.

After careful evaluation, another team member was hired, and, during the next review period, customer service and usability scores improved greatly, while the achievement of goals by individual team members reached 100%.

The No-Win

The fourth case involves a system administrator for a large telecommunications company. The administrator was responsible for the testing and deployment of new applications for the Internet service provision division of the company.

At the beginning of each year, goals for the entire year were established. Rather than individualizing each set of goals, and meeting periodically to review them, overarching goals of the company were used, and the administrator needed to choose one or more of these pre-defined goals. For example, one goal might be to save the company money. The onus was then on the administrator at the end of the year to list the projects worked on, and how he had saved the company money on those projects.

The manager maintained a hands-off approach to the administration of the systems, and so was unable to comment on many parts of the service levels section of the performance review. Additionally, as the administrator was responsible for the testing and integration of new applications, much of the service levels section was not applicable. He therefore received a score of 0 for this section. Similarly, using the method described here, he received a score of 0 on the competency section.

While the manager recognized that large portions of the current review system were not applicable, no effort was made to modify the review to better reflect the administrator's responsibilities.

Additionally, the manager would often not inform administrators of upcoming projects, preventing them from preparing. As a result, much of the testing and integration was performed with little notice and tight deadlines. While not listed as part of the goals section, the manager would still note any time one of these deadlines slipped, and add it to the administrator's file.

As a result, the administrator's reviews would result in an overall score of 0, indicating good performance. However, his reviews would still list missed deadlines, with no explanation as to why these deadlines were missed included, nor with any admittance on the manager's part that better communication would solve many of these issues. The administrator

was put in a defensive position where he needed to indicate how his performance had been outstanding, yet not measured by the current system. As a result, the evaluation system being used was essentially ignored, and provided little guidance or feedback to either administrator or manager on how to improve the company's environment.

Awesome Admin

The final case involves a system administrator in another small, Internet-based company. However, in this company, there are multiple system administrators, each responsible for different portions of the system (that is, the database administrator is separate from the Unix administrator, is separate from the web administrator, etc.).

The goals for the Unix administrator again centered around technical accomplishments, such as installing a new backup system. The administrator easily met all goals set, receiving a score of 0. Similarly, he made very few mistakes, and so received a score of 0 for the third part of the evaluation. For the service levels section, the administrator again scored 0, having met all of the expectations set forth. The total score for the entire evaluation was therefore 0.

In this case, the manager sat with the administrator to determine how to better challenge the administrator, and to determine areas that could be improved. The administrator noted a desire to learn more about security, and so goals were set that included training and obtaining appropriate certification. Additionally, the administrator was given the task of performing a security audit of the entire company, intended to both allow the administrator to learn as well as provide the company with valuable feedback on how it could improve its security processes. As such, the goals shifted from technical accomplishments directly related to the job to be performed (organizational goals), to professional goals that also benefited the company.

Recommendations

Using this approach implies a certain level of maturity in both the people and the processes. The manager must be willing to work with the administrator, particularly in terms of setting appropriate goals. This approach is not intended to be sprung on an unsuspecting administrator as a means of termination, but rather to allow a manager to identify areas needing improvement and to determine how best to address these short-comings. Similarly, it allows an administrator to show how good his performance is (self-promotion of system administrators can't hurt!) to a manager who might not otherwise understand what the administrator does.

Given this, the authors have two recommendations for both managers and administrators. First, always treat each other with respect. This promotes

more open communication. For example, the manager should not talk about how the employee needs to improve, but rather about how the work needs to improve (e.g., more testing should be performed before stating that a task has been completed). This helps to take what can become a very antagonistic situation and keep it focused and non-personal.

Secondly, get everything in writing and signed. Once goals are agreed to by both parties, they need to be put in writing, and both the administrator and the manager need to sign that this is what they have agreed to. Both the manager and administrator should receive copies of this form. This serves to protect both sides: the administrator in the event the manager claims goals that had not been stated, and the manager in the event that the administrator states that he was never told that he needed to perform a particular duty.

Finally, a word of advice to managers in particular: talk to your human resources department! If you are in a large organization, it is likely that they already have recommendations, and possibly even requirements, on how to perform evaluations, as well as being able to provide advice on how to give an appraisal interview (see also [3] for advice on this) and how to handle difficult situations, such as the “poor performer.”

Disclaimer

It should be noted that this approach, while intended to be general, will not necessarily be appropriate for all situations. For example, it assumes that there is already a job tracking system in place, which is an indication of an environment that is following a more mature process, rather than, for example, a small company struggling to get started. (See Kubicki [4] for a good description of mature system administration processes.)

There was also some discussion on where it would be appropriate to deploy this process. In the situations where there is a good system administrator and a good manager, it was felt that this approach would likely only formalize the arrangements that were already likely in place informally. In the cases of a poor administrator and a good manager, this approach provides an opportunity for the manager to work with the administrator to address his shortcomings and help with prioritization. In those cases where the administrator is unable to fulfill all of his required duties, this process provides documentation of objective measures in the event the administrator is to be disciplined or terminated.

Conversely, this approach can also help good administrators who have poor managers. It allows a formalization that indicates how well the administrator is doing, and requires the manager to consider the requirements for the position and what should be done for planning (via the goals process). This protects the

administrator from personality conflicts with their manager by providing a semi-objective measure of their performance. Finally, in the case of a poor administrator with a poor manager, it was felt that in most cases this approach would not improve an already bad situation, and that neither administrator nor manager would be willing to adopt it due to the accountability and effort requirements.

Comments from the Field

When speaking with a manager we knew, his comments were centered primarily on the first stage of the review, the goal setting portion. His concern was that he would rather not set goals, then come back to them some time later and ask if they were met or not. Instead, he would prefer a mechanism that would allow him to ensure that all goals set were on track for being met.

The authors agree with him completely on this and note that, while a performance review often happens yearly, proper management happens daily. In this spirit, the goal section is not intended to be done in one year chunks, but rather should be a process that occurs approximately every three or four months, at an interval that is sensible for the organization. It is especially important in system administration that this process is visited often, as often the major tasks to be performed can not easily be predicted for an entire year. It should also be noted that, while the goal setting and goal review process happens every few months, a good manager will follow up regularly with his employees to ensure that they are on track to meet their goals and to provide any assistance they need.

When speaking with an administrator, comments focused on the pros and cons of evaluation by results. Using the time delay for a server install or customer problem resolution, he argued, was somewhat unreliable. His thoughts were that the most valid method of measurement was the number of times an administrator revisited a problem that had been previously solved.

In this case, we believe that the time measures can be valid, but must be used in context with the other aspects of the review in order to be interpreted meaningfully.

Conclusions

This paper presents an approach to the creation of performance review standards for the area of system administration. This approach is divided into three parts: achievement of goals, service-level requirements, and general competency. It is organized such that administrators can show that they are meeting a standard level of accomplishment, and that managers can know what they should expect from their administrators.

However, this approach provides a framework only. That is, the actual values for the service level requirements, for example, will be organization dependent, and so will require that appropriate values be

determined by each organization before this performance evaluation can be deployed. Unfortunately, there is no “blue book” of standard values available, such as how long it should take to install various pieces of software, etc.

This paper was written to address what the authors perceived as a lack of information for the system administration context. We invite feedback on this approach and encourage discussion and further publications on the subject.

Author Information

Carrie Gates began her system administration career as the sole administrator at a small not-for-profit organization. She left that position for a similar position at Dalhousie University, where she later became the System Manager for the Faculty of Computer Science. She held the manager position for three years before leaving to pursue a Ph.D., specializing in network security. She is currently in the final research phase of her degree, and can be reached at gates@cs.dal.ca.

Jason Rouse has been a system and network administrator for a wide range of companies, such as universities and private businesses, in both Canada and Holland. Most recently he was a Systems and Security Architect for a small private company in Halifax, Canada. He is currently pursuing a Masters degree specializing in networks and network management, and can be reached at rouse@cs.dal.ca.

References

- [1] Donohue, Gene, Creating S.M.A.R.T. Goals, <http://www.topachievement.com/smart.html>, Last visited: 20 May 2003.
- [2] King, Patricia, *Performance Planning and Appraisal: A How-To Book for Managers*, McGraw-Hill Book Company, 1984.
- [3] Kirkpatrick, Donald L., *How to Improve Performance Through Appraisal and Coaching*, AMA-COM Publishing, New York, 1982.
- [4] Kubicki, Carol, “The System Administration Maturity Model – SAMM,” *Proceedings of the Seventh System Administration Conference (LISA 1993)*, Usenix Association, Monterey, California, November, 1993.
- [5] Limoncelli, Thomas A. and Christine Hogan, *The Practice of System and Network Administration*, Addison-Wesley Publishing Company, 2001.
- [6] Phillips, Gretchen, *Hiring System Administrators*, Usenix Association, 1999.
- [7] Ringel, Matthew F. and Thomas A. Limoncelli, “Adverse Termination Procedures -or- ‘how to fire a system administrator’,” *Proceedings of the 13th Systems Administration Conference (LISA 1999)*, Seattle, Washington, USA, 1999.
- [8] Rouillard, Larrie A., *Goals and Goal Setting: Achieving Measured Objectives*, Third Edition, Crisp Publications, 2003.
- [9] Smith, Douglas K., *Make Success Measurable!*, John Wiley & Sons, New York, 1999.