

WebJob Breakeven Analysis

Installing and Configuring a Solaris Package

Klayton Monroe and Andy Bair

KoreLogic Inc.*

Overview

This paper describes the labor cost associated with deploying and configuring a Solaris package on several hundred systems. Three solutions are compared: manual, semi-automated, and fully-automated. The manual solution involves a system administrator remotely logging into each system, untarring a tar ball, and following a set of installation and configuration procedures. The semi-automated solution involves a system administrator remotely logging into each system, untarring a tar ball, and running a setup script, which automates the installation and configuration procedures. The fully-automated solution involves a system administrator scheduling a job to run on all (or specific groups of) WebJob [1] clients and reviewing the results once they have been uploaded to the WebJob server.

Assumptions

The following assumptions were used to carry out the analysis on which this paper is based:

- All installs go according to plan. This paper does not consider the extra time that would be required to debug any problems should a particular install fail.
- A competent system administrator could perform a manual install in approximately 10 minutes. This number reflects the time a system administrator would need to spend on each system downloading the tar ball, executing commands, and verifying each step of the installation process.
- A competent system administrator could produce (i.e., design, develop, and test) a script that automates the installation process in approximately three hours. This would reduce the time a system administrator would

*Initial Revision September 20, 2007; Current Revision September 25, 2007

need to spend on each system to approximately three minutes. Most of that time would be spent downloading the tar ball and waiting for the setup script to complete.

- A competent system administrator could produce (i.e., design, develop, and test) a WebJob script that automates the installation process in approximately three hours. Since this job can be scheduled to run on all systems from the WebJob server, the time a system administrator would need to spend on each system is reduced to zero, and the time required to schedule the job is less than a minute. Because this job is fully automated, the system administrator would need to spend time verifying the results of each install. However, this is easy to do with a WebJob framework because all output generated during the installation and configuration process is captured and uploaded to the WebJob server. If that output includes pass/fail status information, a system administrator could verify the installation process for all systems in less than a minute with a simple `grep(1)`.
- The manual and semi-automated solutions are limited to 4^1 and 12.56^2 installs per day, respectively. Both of these solutions suffer due to the inefficiencies inherent in traditional, ad hoc administrative practices.
- An operational WebJob framework already exists, and the costs associated with setting it up and maintaining it are not included in this analysis.
- All WebJob clients execute their job on a daily basis in a fixed two-hour maintenance window, and the average job execution time is 60 seconds.

¹This number is based on a case study where it took a large organization approximately 60 days to perform 240 installs.

²This number is based on the ratio between the manual and semi-automated solutions (3.15:1), which is derived from the data for Figure 1 at the point in each curve where the number of systems equals 1000.




Figure 1: Total Labor Time

Total Labor Time

Figure 1 shows the total labor time in hours for the three solutions. This is the amount of time a system administrator would be required to do real work such as manually executing commands or writing a script to automate the installation process. Based on the above graph, the following items are worth noting:

- The breakeven point for the fully-automated solution (i.e., the point where the blue line intersects the red line) is three hours and 18 systems. This means that the cost of writing the job script is fully recovered if the package is installed on 18 or more systems.
- The total time required to implement the fully-automated solution is fixed. This means the labor costs associated with developing a single WebJob script are invariant with respect to the number of systems in a given deployment. In other words, you save time and money because you only have to write the job script once.
- The manual and semi-automated solutions require approximately 4.2 and 1.3 man weeks for 1000 systems, respectively. The fully-automated solution, on the other hand, can be effected in a single day.

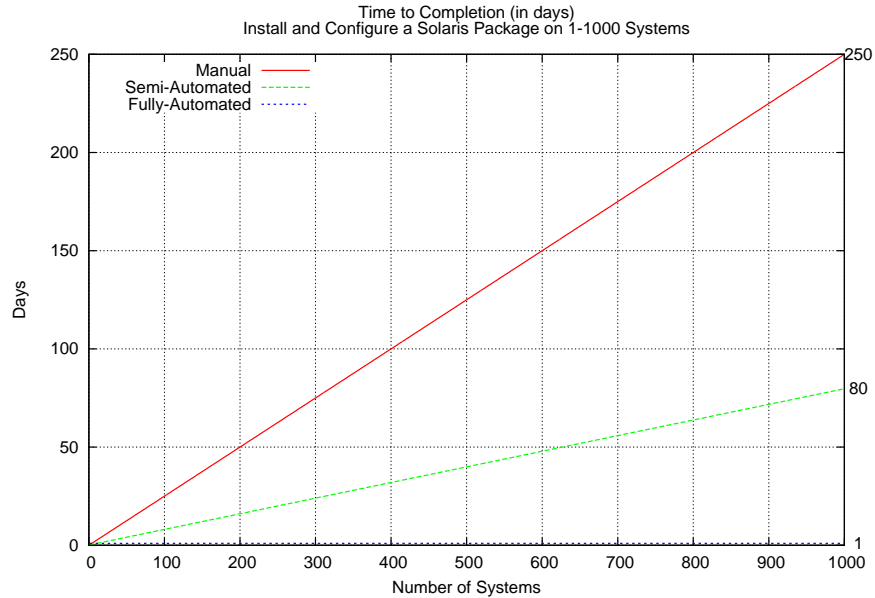


Figure 2: Time to Completion

Time to Completion

Figure 2 shows the time to completion in days for the three solutions. This is the amount time it took to complete all installs regardless of the solution or techniques used. Based on the above graph, the following items are worth noting:

- The breakeven point for the fully-automated vs. manual solutions (i.e., the point where the blue line intersects the red line) is one day and 4 systems.
- The breakeven point for the fully-automated vs. semi-automated solutions (i.e., the point where the blue line intersects the green line) is one day and 11 systems.
- The semi-automated solution is limited to the ratio between the manual and semi-automated solutions in Figure 1, which is 3.15:1. In reality things won't be that good. Generally, it doesn't matter how efficient the install process is when human intervention is required. What matters is how many installs a system administrator will actually perform in a given work day. For this reason, the semi-automated solution will likely be much closer to the curve depicted for the manual solution.

- The main distinction, therefore, between the manual and semi-automated solutions is that the semi-automated solution is much less error prone because someone took the time to create a script that automates all the commands that a system administrator would have had to enter manually.
- The manual solution requires nearly one man year (260 days) to complete. Obviously, this number could be reduced by delegating the work to multiple system administrators. However, the fully-automated solution requires no more than one system administrator.

Conclusions

Based on the analysis presented in this paper, we believe that one can draw the following conclusions:

- It pays to automate common administrative tasks such as installing and configuring a Solaris package. Given a rate of \$50/hour and 1000 systems, the WebJob solution would save \$8,200 and \$2,500 over the manual and semi-automated solutions depicted in Figure 1, respectively.
- The cost savings associated with this particular task translates to almost any other common/repetitive administrative task (e.g., config file changes, user management, certificate updates, etc.). If an organization was required to perform 100 such tasks per year, the savings would amount to \$820,000 and \$250,000 for the manual and semi-automated solutions, respectively.
- The fully-automated solution is far more efficient than the manual or semi-automated solutions because almost all the people and day-to-day complications (e.g., interruptions, email, phone calls, meetings, long lunches, office politics, etc.) have been taken out of the loop.
- The time to completion for the fully-automated solution is 250 and 80 times faster than the manual and semi-automated solutions, respectively. In a highly competitive environment where time is money, those who automate clearly have an advantage over those who do not.

References

- [1] <http://webjob.sourceforge.net/WebJob/>