

USENIX Association

Proceedings of
LISA 2002:
16th Systems Administration
Conference

Philadelphia, Pennsylvania, USA
November 3–8, 2002

**USENIX
SAGE**

© 2002 by The USENIX Association

All Rights Reserved

For more information about the USENIX Association:

Phone: 1 510 528 8649

FAX: 1 510 548 5738

Email: office@usenix.org

WWW: <http://www.usenix.org>

Rights to individual papers remain with the author or the author's employer.

Permission is granted for noncommercial reproduction of the work for educational or research purposes.

This copyright notice must be included in the reproduced paper. USENIX acknowledges all trademarks herein.

A Simple Way to Estimate the Cost of Downtime

David A. Patterson – University of California at Berkeley

ABSTRACT

Systems that are more dependable and less expensive to maintain may be more expensive to purchase. If ordinary customers cannot calculate the costs of downtime, such systems may not succeed because it will be difficult to justify a higher price. Hence, we propose an easy-to-calculate estimate of downtime.

As one reviewer commented, the cost estimate we propose “is simply a symbolic translation of the most obvious, common sense approach to the problem.” We take this remark as a complement, noting that prior work has ignored pieces of this obvious formula.

We introduce this formula, argue why it will be important to have a formula that can easily be calculated, suggest why it will be hard to get a more accurate estimate, and give some examples.

Widespread use of this obvious formula can lay a foundation for systems that reduce downtime.

Introduction

It is time for the systems community of researchers and developers to broaden the agenda beyond performance. The 10,000X increase in performance over the last 20 years means that other aspects of computing have risen in relative importance. The systems we have created are fast and cheap, but undependable. Since a portion of system administration is dealing with failures [Anderson 1999], downtime surely adds to the high cost of ownership.

To understand why they are undependable, we conducted two surveys on the causes of downtime. In our first survey, Figure 1 shows data we collected failure data on the U. S. Public Switched Telephone Network (PSTN) [Enriquez 2002]. It shows the percentage of failures due to operators, hardware failures, software failures, and overload for over 200 outages in 2000. Although not directly relevant to computing systems, this data set is very thorough in the description of the problem and the impact of the outages. In our second study, Figure 2 shows data we collected failure from three Internet sites [Oppenheimer 2002]. The surveys are notably consistent in their suggestion that operators are the leading cause of failure.

Collections of failure data often ignore operator error, as it often requires asking operators if they think they made an error. Studies that are careful about how they collect data do find results that are consistent with these graphs [Gray 1985, Gray 1990, Kuhn 1997, Murphy 1990, Murphy 1995].

Improving dependability and lowering cost of ownership are likely to require more resources. For example, an undo system for operator actions would need more disk space than conventional systems. The marketplace may not accept such innovations if

products that use them are more expensive and the subsequent benefits cannot be quantified by lower cost of ownership. Indeed, a common lament of computer companies that customers may complain about dependability, but are unwilling to pay the higher price of more dependable systems. The difficulty of measuring cost of downtime may be the reason for apparently irrational behavior.

Hence, this paper, which seeks to define a simple and useful estimate of the cost of unavailability.

Estimating Revenue and Productivity

Prior work on estimating the cost of downtime is usually measuring the loss of revenue for online companies or other services that cannot possibly function if their computers are down [Kembel 2000]. Table 1 is a typical example.

Brokerage operations	\$6,450,000
Credit card authorization	\$2,600,000
Ebay	\$225,000
Amazon.com	\$180,000
Package shipping services	\$150,000
Home shopping channel	\$113,000
Catalog sales center	\$90,000
Airline reservation center	\$89,000
Cellular service activation	\$41,000
On-line network fees	\$25,000
ATM service fees	\$14,000

Table 1: Cost of one hour of downtime. From *InternetWeek 4/3/2000* and based on a survey done by Contingency Planning Research. [Kembel 2000].

Such companies are not the only ones that lose revenue if there is an outage. More importantly, such a table ignores the loss to a company of wasting the time of employees who cannot get their work done during an outage, even if it does not affect revenue. Thus, we need a formula that is easy to calculate so that administrators and CIOs in any institution can determine the costs of outage. It should capture both the cost of lost productivity of employees and the cost of lost revenue from missed sales.

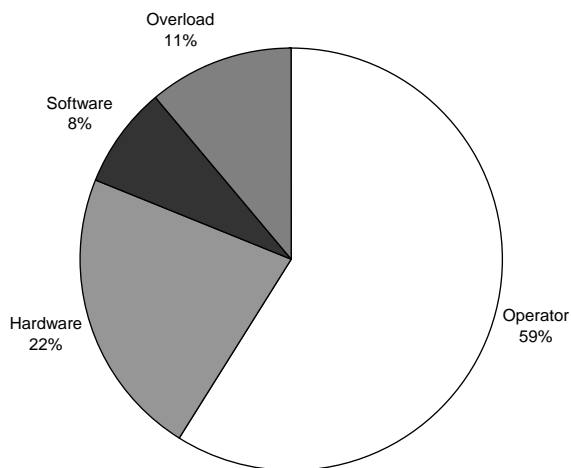


Figure 1: Percentage of failures by operator, hardware, software, and overload for PSTN. The PSTN data measured blocked calls during an outage in the year 2000. (This figure does not show vandalism, which is responsible for 0.5% of blocked calls.) We collected this data from the FCC; it represents over 200 telephone outages in the U. S. that affected at least 30,000 customers or lasted at least 30 minutes. Rather than only reporting outages, telephone switches record the number of attempted calls blocked during an outage, which is an attractive metric. The figure does not include environmental causes, which are responsible for 1% of the outages.

We start with the formula, and then explain how we derived it:

$$\text{Estimated average cost of 1 hour of downtime} = \text{Empl. costs/hour} * \% \text{ Empl's affected by outage} + \text{Avg. Rev./hour} * \% \text{ Rev. affected by outage}$$

Employee costs per hour is simply the total salaries and benefits of all employees per week divided by the average number of working hours per month. *Average revenue per hour* is just the total revenue of an institution per month divided by average number of hours per week an institution is open for business. Note that this term includes two factors: revenue associated with a web site and revenue supported by the internal information technology infrastructure. We believe these employee costs and revenue are not too difficult to calculate, especially since they are input to an estimate, and hence do not have to be precise to the last penny.

For example, publicly traded companies must report their revenue and expenses every quarter, so quarterly statements have some of the data to calculate these terms. Although revenue is easy to find, employee costs are typically not reported separately. Fortunately, they report the number of employees, and you can estimate the cost per employee. The finance department of smaller companies must know both these terms to pay the bills and issue paychecks. Even departments in public universities and government agencies without conventional revenue sources have their salaries in the public record.

The other two terms of the formula – fraction employees and fraction revenue affected by outage – just need to be educated guesses or ranges that make sense for your institution.

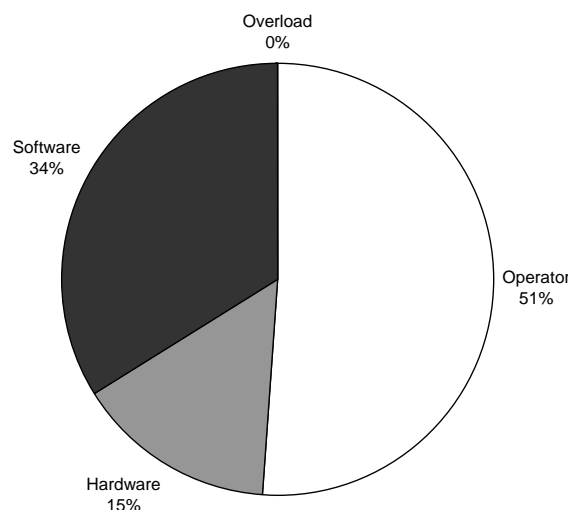


Figure 2: Percentage of failures by operator, hardware, software, and overload for three Internet sites. Note that the mature software of the PSTN is much less of a problem than Internet site software, yet the Internet sites have such frequent fluctuations in demand that they have overprovisioned sufficiently so that overload failures are rare. The Internet site data measured outages in 2001. We collected this data from companies in return for anonymity; it represents six weeks to six months of services with 500 to 5000 computers. Also, 25% of outages had no identifiable cause, and are not included in the data. One reviewer suggested the explanation was “nobody confessed;” that is a plausible interpretation.

Is Precision Possible?

As two of the four terms are guesses, the estimate clearly is not precise and open to debate. Although we all normally strive for precision, it may not be possible here. To establish the argument for systems that may be more expensive to buy but less expensive to own, administrators only need to give an example of what the costs might be using the formula

above. CIOs could decide on their own fractions in making their decisions.

The second point is that much more effort may not ultimately lead to a precise answer, for there are some questions that will be very hard to answer. For example, depending on the company, one hour of downtime may not lead to lost revenue, as customers may just wait and order later. In addition, employees may simply do other work for an hour that does not involve a computer. Depending on centralization of services, an outage may only affect a portion of the employees. There may also be different systems for employees and for sales, so an outage might affect just one part or the whole company. Finally, there is surely variation in cost depending when an outage occurs; for most companies, outages Sunday morning at 2 AM probably has little impact on revenue or employee productivity.

Before giving examples, we need to qualify this estimate. It ignores the cost of repair, such as the cost of overtime by operators or bringing in consultants. We assume these costs are small relative to the other terms. Second, the estimate ignores daily and seasonal variations in revenue, as some hours are more expensive than others. For example, a company running a lottery is likely to have rapid increase in revenue as the deadline approaches. Perhaps a best case and worst-case cost of downtime might be useful as well as average, and its clear how to calculate them from the formula.

Example 1: University of California at Berkeley EECS Department

The Electrical Engineering and Computer Science (EECS) department at U. C. Berkeley does not sell products, and so there is no revenue to lose in an outage. As we have many employees, the loss in productivity could be expensive.

The employee costs have two components: those paid for by state funds and those paid for by external research funds. The state pays 68 full-time staff collective salaries and benefits of \$403,130 per month. This is an annual salary and benefits of \$71,320. These figures do not include the 80 faculty, who are paid approximately \$754,700 per month year round, including benefits. During the school year external research pays 670 full-time and part-time employees \$1,982,500, including benefits. During the summer, both faculty and students can earn extra income and some people are not around, so the numbers change to 635 people earning \$2,915,950. Thus, the total monthly salaries are \$3,140,330 during the year and \$4,073,780 during the summer.

If we assume people worked 10 hours a working day, the Employee costs per hour is \$14,780 during the school year and \$19,170 during the summer. If we assumed people worked 24 hours a day, seven days a week, the costs would change to \$4300 and \$5580.

If EECS file servers and mail servers have 99% availability, that would mean seven hours of downtime per month. If half of the outages affected half the employees, the annual cost in lost productivity would be \$250,000 to \$300,000.

Since I was a member of EECS, it only took two emails to collect the data.

Example 2: Amazon

Amazon has revenue as well as employee costs, and virtually all of their revenue comes over the net. Last year their revenue was \$3.1B and it has 7744 employees for the last year. This data is available at many places; I got it from Charles Schwab. Alas, annual reports do not break down the cost of employees, just the revenue per employee. That was \$400,310. Let's assume that Amazon employee salaries and benefits are about 20% higher than University employees at, say, \$85,000 per year. Then Employee costs per hour working 10 hours per day, five days per week would be \$258,100 and \$75,100 for 24x7. Revenue per hour is \$353,900 for 24x7, which seems the right measure for an Internet company. Thus, an outage during the workweek that affected 90% employees and 90% of revenue streams could cost Amazon about \$550,000 per hour.

We note that employee costs are a significant fraction of revenue, even for an Internet company like Amazon. One reason is the Internet allows revenue to arrive 24x7 while employees work closer to more traditional workweeks.

Example 3: Sun Microsystems

Sun Microsystems probably gets little of its income directly over the Internet, since it has an extensive sales force. That sales force, however, relies extensively on email to record interactions with customers. The collapse of the World Trade Center destroyed several mail servers in the New York area and many sales records were lost. Although its not likely that much revenue would be lost if the Sun site went down, it could certainly affect productivity if email were unavailable, as the company's nervous systems appears to be email.

In the last year, Sun's revenue was \$12.5B and it had 43,314 employees. Let's assume that the average salary and benefits are \$100,000 per employee, as Sun has a much large fraction of its workforce in engineering than does Amazon. Then employee costs per hour working 10 hours per day, five days per week would be \$1,698,600 and \$494,500 for 24x7. Since Sun is a global company, perhaps 24x5 is the right model. That would make the costs \$694,100 per hour. Revenue per hour is \$1,426,900 for 24x7 and \$2,003,200 for 24x5, with the latter the likely best choice.

Let's assume that a workweek outage affected 10% of revenue that hour and 90% of the employees. The cost per hour would about \$825,000, with three-fourths of the cost being employees.

Conclusion

The goal of this paper is to provide an easy to calculate estimate of the average cost of downtime so as to justify systems that may be slightly more expensive to purchase but potentially spend much less time unavailable. It is important for administrators and CIOs to have an easy to use estimate to set the range of costs of outages to make them more likely to take it into consideration when setting policies and acquiring systems. If it were hard to calculate, few people would do it.

Although a simple estimate, we argue that a much more time-consuming calculation may not shed much more insight, as it is very hard to know how many consumers will simply reorder when the system comes back versus go to a competitor, or how many employees will do something else productive while the computer is down.

We see that employee costs, traditionally ignored in such estimates, are significant even for Internet companies like Amazon, and dominate the costs of more traditional organizations like Sun Microsystems. Outages at universities and government organizations can still be expensive, even without a loss of a significant computer-related revenue stream.

In addition to this estimate, there are may be indirect costs to outages that can be as important to the company as these more immediate costs. Outages can lead to management overhead as the IT department is blamed for every possible problem and delay throughout the company. Company morale can suffer, reducing everyone's productivity for periods that far exceed the outage time. Frequent outages can lead to a loss of confidence in the IT team and its skills. Such change in stature could eventually lead to individual departments hiring their own IT people, which lead to direct costs.

As many researchers are working on these solutions to the dependability problems [IBM 2000, Patterson 2002], our hope is that these simple estimates can help organizations justify systems that are more dependable, even if a bit more expensive.

References

- [Anderson 1999] Anderson, E. and D. Patterson, "A Retrospective on Twelve Years of LISA Proceedings," *Proc. 13th Systems Administration Conference - LISA 1999*, Seattle, Washington, p. 95-107, Nov 1999,
- [Enriquez 2002] Enriquez, P., A. Brown, and D. Patterson, "Lessons from the PSTN for Dependable Computing," *Workshop on Self-Healing, Adaptive and self-MANaged Systems (SHAMAN)*, New York, NY, June 2002.

- [Gray 1985] Gray, J., *Why Do Computers Stop and What Can Be Done About It?* TR-85.7, Tandem Computers, 1985.
- [Gray 1990] Gray, J., *A Census of Tandem System Availability Between 1985 and 1990*. TR-90.1, Tandem Computers, 1990.
- [IBM 2000] *Autonomic Computing*, <http://www.research.ibm.com/autonomic/>, 2000.
- [Kembel 2000] Kembel, R., *Fibre Channel: A Comprehensive Introduction*, 2000.
- [Kuhn 1997] Richard Kuhn, D., "Sources of Failure in the Public Switched Telephone Network," *IEEE Computer*, Vol. 30, No. 4, <http://hissa.nist.gov/kuhn/pstn.html>, April, 1997.
- [Murphy 1995] Murphy, B. and T. Gant, "Measuring System and Software Reliability Using an Automated Data Collection Process," *Quality and Reliability Engineering International*, Vol. 11, pp. 341-353, 1995.
- [Murphy 2000] Murphy, B., "Windows 2000 Dependability," *Proceedings of the IEEE International Conference on Dependable Systems and Networks*, June 2000.
- [Oppenheimer 2002] Oppenheimer, D. and D. A. Patterson. "Studying and using failure data from large-scale Internet services," *Tenth ACM SIGOPS European Workshop*, Saint-Emilion, France, September 2002.
- [Patterson 2002] Patterson, D., A. Brown, P. Broadwell, G. Candea, M. Chen, J. Cutler, P. Enriquez, A. Fox, E. Kiciman, M. Merzbacher, D. Oppenheimer, N. Sastry, W. Tetzlaff, J. Traupman, and N. Treuhaft, *Recovery-Oriented Computing (ROC): Motivation, Definition, Techniques, and Case Studies*, U. C. Berkeley Computer Science Technical Report UCB//CSD-02-1175, March 15, 2002.