

## The Paradox of the 9s

By [Hank Marquis](#)

Hank is EVP of Knowledge Management at Universal Solutions Group, and Founder and Director of NABSM.ORG. Contact Hank by email at [hank.marquis@usact.com](mailto:hank.marquis@usact.com). View Hank's blog at [www.hankmarquis.info](http://www.hankmarquis.info).



While customers don't understand "9s" they do think that "more '9s' is better." What they don't understand is that too many "9s" actually costs more than it returns. Which leads to the question, how many "9s" do customers really need?

The IT Infrastructure Library® (ITIL®) includes Availability Management. The eyes of most students (and practitioners) glaze over when confronted with the concepts the ITIL describes.

Yet, within its complexity there lie the seeds of business IT alignment, one of the most sought after and seldom attained IT achievements. The idea is really pretty simple: understand what downtime means to the business, and then devise plans to prevent downtime from occurring.

Key to attaining this understanding is qualifying (validating with the business) what downtime actually means, and then quantifying (measuring) downtime accordingly. It is the combination of the two that can yield business IT alignment.

Quantification is usually done using "9's", like 99.999%, and most IT departments can do this at some level. However, qualification is the part most IT organizations simply have not mastered.

Based on my own experiences working in highly available systems spanning from the military to Wall Street, this article provides a matrix showing what "9s" mean to the business (qualification), explains how the "9s" work (quantification), and finishes off describing how to choose the right availability targets.

### Qualification or, Why You Should Care About Availability

While customers don't understand "9s" they do think that "more '9s' is better." What they don't understand is that too many "9s" actually costs more than it returns. Which leads to the question, how many "9s" do customers really need?

Customers don't actually care about "9s" regardless of what they may say. However, what they are essentially taping into here is that IT systems exist to improve the performance and productivity of workers. They know how well an IT system increases productivity is a function of the quality of the IT service provided, and that when the system is down they lose profits and productivity.

In this way they have come to learn about availability and have been led to believe that "more '9s' is better" for them. However, this is only true to a degree.

The problem is trying to explain IT statistics and mathematics to business managers and customers. We need to help customers understand what availability really means to them since they are not capable of doing so themselves (they are not dumb, but they are not IT engineers.)

What they do understand is their business, their marketplace, their costs, and their sales. We in IT have to discuss and justify technical things in business terms, and one way to start talking about availability with customers is in the cost of downtime (CoD).

There are many descriptions of CoD. A Meta Group report stated that in industries such as banking, telecommunications, financial institutions and manufacturing, CoD easily surpasses \$1 million per hour. Contingency Planning Research placed the number from \$28,000/hour on the low end to \$5.4 million on the high end for IT of all types.

Various reports and studies back this up:

- Find/SVP polled 450 Fortune 1000 companies and reported an average hourly downtime cost of \$82,500/hour.
- TechWise Research pegged it at \$71,000/hour after interviewing 93 poll respondents.

- ComputerWorld reported a study of industries that determined an average of \$44,000/hour.

The average of these three studies is a CoD of \$65,833/hour, or \$1,097/minute. But wait, it gets worse! Computerworld has reported that over 37% of those polled reported unplanned downtime of one hour or more per month -- that's at least \$789,996 per year.

Numbers like these make you think '5 9s' are a must! And while downtime costs money, the amount of money varies widely from industry to industry, and from service to service within an industry. Several years ago, Dataquest (Gartner) developed a model of the cost of an hour of downtime by industry, which I convert in table 1 to the cost of a minute of downtime. Find the industry closest to your own and see what an average minute of outage costs your business.

Industry	IT Service	CoD per Min and per Hr
Financial	Brokerage Operations	\$107,500/\$6,450,000
Financial	Credit card/Sales authorization	\$43,333/\$2,600,000
Financial	ATM fees	\$241/\$14,500
Media	pay-per-view	\$2,500/\$150,000
Media	teleticket sales	\$1,150/\$69,000
Retail	Home shopping	\$1,883/\$113,000
Retail	Catalog sales	\$1,500/\$90,000
Transportation	Airline reservations	\$1,483/\$89,000
Transportation	Package shipping	\$466/\$28,000
	Mean CoD/min	\$17,784

Table 1. Downtime Cost per Hour for Various Industries (source Dataquest)

Table 1 proves that not all applications require the same level of availability. In fact, high availability can cost the business more than it returns!

Still, a Standish Group study of downtime numbers found that the average application that IT labeled as “mission-critical” had the following availability experience. Note that these numbers are PER MONTH.

- 9% -- less than one hour down per month
- 24% -- from one hour to less than nine hours (60 to 540 minutes) per month
- 67% -- nine hours or greater (540 or more minutes) per month

Based on reported statistics, the average IT organization is costing its business more than \$592,000 month, or \$7,108,560/year due to unplanned downtime. Depending on your business you could be losing millions of dollars every hour, and even more per year. Money is not all the matters in CoD. In other applications, such as emergency dispatch 911 systems, CoD is in terms of life and death.

This is what customers intuitively feel, and why they think “more ‘9s’ is better.” It is also why we have to help them understand how many 9s they actually need and can afford.

## Quantification, or Measuring Availability

Now that you have qualified (or understood) the potential business impact of downtime, and why the business cares about availability, you can begin to consider how to measure it. Once you can measure availability, you can measure downtime. Only then can you take steps to eliminate it if appropriate.

Calculating technical metrics is the job of Availability Management. Using “9s” lets you describe availability in a technical manner, which while pretty useless to business people, is ultimately vital to business productivity and good decision making. Availability is typically measured as a percentage of total agreed uptime available over a period. The formula is:

$$\text{Availability\%} = \frac{(\text{Agreed Service Time} - \text{Unplanned Downtime})}{\text{Agreed Service Time}} \times 100$$

Here are some examples of downtime calculations:

- 99.999% availability = 5 minutes/year downtime
- 99.99% availability = 53 minutes/year downtime
- 99.9% availability = 528 minutes (8.8 hours)/year
- 99.5% availability = 2,628 minutes (43.8 hours)/year
- 99% availability = 5,256 minutes (87.6 hours)/year

As you can see, customers are onto something as “9s” do matter! Gartner (Dataquest) has assigned classifications or labels to these ranges to make it easier to understand, as shown in table 2.

Availability Classification	Level of Availability (%)	Annual Downtime
Continuous Processing	100%	0 minutes
Fault Tolerant	99.999%	5 minutes
Fault Resilient	99.99%	53 minutes
High Availability	99.9%	8.8 hours
Normal Commercial Availability	99 - 99.5%	87.6 - 43.8 hours

Table 2. Availability Classification

Based on the preceding data, if you are like most you experience 9 hours (540 minutes) or more per month. Using the previous table, you can see how much unplanned IT outages costs per year in table 3.

Industry	IT Service	Cost per Min	Cost/IT Service/yr
Financial	Brokerage Operations	\$107,500	\$58,050,000
Financial	Credit card/Sales authorization	\$43,333	\$23,399,820
Media	pay-per-view	\$2,500	\$1,350,000
Retail	Home shopping	\$1,883	\$1,016,820
Retail	Catalog sales	\$1,500	\$810,000
Transportation	Airline reservations	\$1,483	\$800,820
Media	teleticket sales	\$1,150	\$621,000
Transport	Package shipping	\$466	\$251,640
Finance	ATM fees	\$241	\$130,140

Table 3. Average cost of unplanned downtime at 99.9% availability

Now you know why the business cares. But what do they care about? Not everything can operate at “5 9s,” it would simply cost too much.

Part of the job of availability management is to produce service and business views of availability. Those IT services that are most critical underpin what ITIL calls a Vital Business Function, VBF. Focusing on the IT services that underpin a VBF, the business can make an informed decision regarding importance, and thus the return on investment (ROI) calculation about availability. The issue of course is ‘cost to improve’ vs. ‘loss from downtime.’

Table 3 shows the average cost per year of downtime for a service with 3 ‘9s’ (99.9% availability.) If the cost to attain 4 ‘9s’ (99.99%) or 5 ‘9s’ (99.999%) is greater than the loss, then it does not make sense to do it.

## Translating Business Needs Into 9s

Now that you have an understanding of the ‘9s’ from a customer perspective the next step is to negotiate with them. You have to be clear here, since the availability targets you negotiate usually get documented in Service Level Agreements (SLAs.)

According to the Availability Index from “Blueprints for High Availability,” there are four stages of availability:

Availability	Index Name	Example(s)
1	Basic	no specific extra measures
2	Data	data backup
3	System	redundancy, fail-over, etc.
4	Organization	alternate sites, systems, services, etc.

Table 4. Availability Index

Using the availability index you can focus on which areas of service delivery to focus. Obviously as you progress higher into the availability index costs increase. However, once you understand the VBF, and the business availability needs of the VBF, ITIL offers several methods to improve availability, driven by risk analysis and continuity management in conjunction with the customer.

Work with your customers to help them understand what ‘9s’ really mean in terms of downtime per year (and per year is important here.)

Have customers define the cost of downtime so that they can qualify the value of reducing it. Remember that sometimes, the cost of increased availability is more than the value the increased availability delivers.

After quantifying the cost of downtime and then qualifying the need for availability, select the ‘9s’ the customer really needs and the organization can afford.

Once you know what you need to deliver in terms of '9s', engineer solutions based on the broad ITIL categories of recovery options ("do nothing", "manual backup", "reciprocal agreement", and "immediate", "intermediate" and "gradual") focusing on the availability index shown in table 4.

Follow these directions and you will make a positive contribution to the bottom line, and be seen as enabling business IT alignment!

Entire Contents © 2009 itSM Solutions® LLC. All Rights Reserved.  
ITIL ® and IT Infrastructure Library ® are Registered Trade Marks of the Office of Government Commerce and is used here by itSM Solutions LLC under license from and with the permission of OGC (Trade Mark License No. 0002).