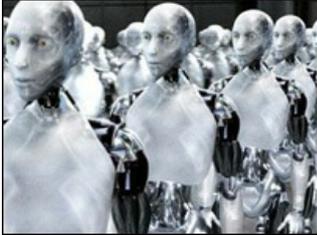# I, Cloud

**Lori MacVittie, 2010-11-10**

*Do we need Three Laws of Cloud? Not yet. Neither should we be overly concerned regarding reports of cloud leading to the elimination of IT.*



Every time a technological innovation has spurred automation – since the time of Henry Ford right up to a minute ago – someone has claimed that machines will displace human beings. But the rainbow and unicorn dream attributed to business stakeholders everywhere, i.e. the elimination of IT, is just that – a dream. It isn't realistic and in fact it's downright silly to think that systems that only a few years ago were unable to automatically scale up and scale down will suddenly be able to perform the complex analysis required of IT to keep the business running.

The rare reports of the elimination of IT staff due to cloud computing and automation are highlighted in the news because they evoke visceral reactions in technologists everywhere and, to be honest, they get the click counts rising. But the jury remains out on this one and in fact many postulate that it is not a reduction in staff that will occur, but a transformation of staff, which may eliminate some old timey positions (think sysadmins) and create new ones requiring new skills (think devops).

I am, as you may have guessed, in the latter camp. IT needs to change, yes, but that change is unlikely to be the elimination of IT. Andi Mann 🔧 (VP with CA Technologies) put it well when he says yes, IT staff reductions are always a possible outcome of better IT, "*Yet with a sunk cost in training and skills, and the seemingly endless list of projects on most CIOs' desks, is cutting staff numbers really a good outcome? How does reassignment and redeployment fit into this value too – better or worse?*" It is that "seemingly endless list of projects" that makes a mass reduction in IT unlikely along with the fact that systems are simply not ready to "take over" from human beings. Not yet.

## DATA is not INFORMATION

Any business stakeholder who dreams of a data center without IT should ask themselves this question: could a machine do *my* job? Could it analyze data and, from all those disparate numbers and product names and locations come up with a marketing plan? Could it see trends across disparate industries and, taking into consideration the season and economic conditions, determine what will be the next big seller?

Probably not, because such decisions require an analytical thought process that simply doesn't exist in the world of technology today. The "intelligence" that exists in any system today is little more than a codified set of rules that were specified by  - wait for it, wait for it – yes, a human being. It was a person who sat down and codified a set of basic rules for automatically responding to deviances in performance and capacity and specified what action should be taken. Without those basic rules the systems could not decide whether to turn themselves on or off, let alone make a decision as complex as where to direct any given application request.

When you study communication theory you discover some very basic facts about the nature of learning and intelligence that comes down to this: words, which are just data, have no intrinsic meaning. Words, numbers, data. These things by themselves carry no especial value in and of themselves. It is only when they are seen and understood by a human being that they become valuable and become "information". The same is true of the data that flies around a data center and upon which decisions are made: it's just data, even to the systems, until it's interpreted by a human being.

Not convinced? Shall we play a game? **\***

## 3.14

What is this number? Most of you probably answered, "pi" - the mathematical constant used in formulas involving circles. But it could just as easily represent the number of milliseconds it took for a packet to traverse a segment of the network, or the number of seconds it took for the first byte of a response to an HTTP request to arrive on a user's desktop, or the average number of items purchased by storefront X in the month of August. It could be the run-rate in hundreds of thousands of dollars of a technology start up or perhaps it's my youngest daughter's GPA.

So which is it?

You need context in order to interpret what that number means. I'll call that a point proven but in case you're not convinced, let's dig a little deeper. Even after it's interpreted in its proper context this number requires <u>further analysis</u> to become valuable. After all, if that's my daughter's GPA you don't know whether that's good or bad without knowing a lot more about her. Maybe she's underachieving, maybe she's overachieving. Maybe she's seven years old and in high school and that's amazing.

There's just a lot more intelligence required to make sense out of a piece of data than we realize. The reaction taken to this data once it becomes information requires analysis; *human* analysis. Even if we could codify every rule and correlate all the data necessary to make sense out of a simple number, there's still the fact that there are exceptions to every rule and there's always something we didn't consider that changes the equations.

## PEOPLE SKILLS REQUIRED

I'm not talking about the customer-service-likes-to-interact-with-others kind of people skills, I'm talking about analytics and the ability to think through a problem or, even better, simply recognize one when it happens. The best example of the continuing need for such skills is the recent outage experienced by Facebook.

> " The 150 minute-long outage, during which time the site was turned off completely, was the result of **a single incorrect setting** that produced a cascade of erroneous traffic, Facebook software engineering director Robert Johnson said in a posting to the site.
>
> "Today we made a change to the persistent copy of a **configuration value that was interpreted as invalid**. This meant that every single client saw the invalid value and attempted to fix it. Because the fix involves making a query to a cluster of databases, that cluster was quickly overwhelmed by hundreds of thousands of queries a second," Johnson said.
>
> "To make matters worse, every time a client got an error attempting to query one of the databases, it interpreted it as an invalid value and deleted the corresponding cache key," he added. "This meant that even after the original problem had been fixed, the stream of queries continued. As long as the databases failed to service some of the requests, they were causing even more requests to themselves. We had entered a feedback loop that didn't allow the databases to recover." [emphasis added]
>
> "Facebook outage due to internal errors, says company" ZDNet UK (September 26, 2010)

The first thing that comes to mind on reading this explanation is that if the configuration change had been made by a human being manually it might have been followed by an error message and all subsequent changes halted until it was determined why the system thought the setting was invalid. But the system, elegantly automated, propagated the

**THREE LAWS of CLOUD**

1. A cloud may not negatively impact the availability of an application or, through inaction, allow an application to become unavailable.

2. A cloud must obey any instructions given to it by operators, except where such orders would conflict with the First Law.

3. A cloud must assure its own availability as long as such assurance does not conflict with the First or Second Law.

*Adapted from Asimov's Three Laws of Robotics*

erroneous setting and as it cascaded through the system, which was automated in a way as to try to fix the problem on its own, it just made things worse and worse. At no point did the system even recognize that something was wrong, that took a human being. In the aforementioned post by Robert Johnson on the outage, he stated, "*An automated system for verifying configuration values ended up causing much more damage than it fixed.*"

In the end, the system needed to be shut down and restarted. A decision that was made by a human being, not a machine, because it was only when a human being looked at what was happening, when a human being *evaluated* the flow of data across systems and networks, that they were able to determine what the source of the problem was and, ultimately, how to resolve it. The system saw nothing wrong because it was acting exactly the way it should; it followed its programming to the letter despite the fact that it was ultimately destroying the system it was supporting.

It was just acting on data because that's all it can do; it cannot analyze and interpret that data into information that leads to the right action. Until it can, we don't really need a "Three Laws of Cloud" because the systems are not capable of performing the kind of analysis necessary to even recognize its actions might be harming the very applications it is built to deliver (an adaptation of Asimov's Second Law of Robotics).

## AUTOMATION CREATES new OPPORTUNITIES

Facebook's system, like many of those being designed and developed in organizations around the world, is the codification of processes. It is the digitized orchestration of many disparate tasks that can be performed individually and as a whole result in a specific work-flow execution. Its purpose is to minimize the amount of manual intervention required while maximizing the efficiency of such processes. The processes aren't new, but their form is. They are chunks of conditional logic that takes as parameters one or more pieces of *data* and acts upon that data. That's it. There's no intelligence, no analysis, no gut instinct to guide the system into making a choice other than the ones codified in its scripts and daemons and services.

History teaches us that assembly line technologies, which is as close a real-world analogy to automation and IT as we're likely to get, do not reduce the number of human beings required to monitor, manage, and improve the processes codified to achieve such automation. Instead, it frees human beings to do what they are best at: analyzing, innovating and finding new ways to do what we've always done that are more efficient. It allows us to do things faster and, eventually, commoditize the output such that we can focus on leveraging those processes to build better and more awesome versions of the output.

What automation of "IT" does is create new opportunities for IT, it does not erase the need for it. Until our systems are able to analyze and interpret data such that it becomes information and then act on that information in ways that may be "outside the existing ruleset" then IT – and more specifically the people that comprise IT – will not only be needed they will be necessary to the continued growth and evolution of not only IT but the business.

> " …it was Ransom E. Olds and his Olds Motor Vehicle Company (later known as Oldsmobile) who would dominate this era of automobile production. Its large scale production line was running in 1902. Within a year, Cadillac (formed from the Henry Ford Company), Winton, and Ford were producing cars in the thousands.
>
> -- Wikipedia, History of the Automobile

Notice that "large scale" one hundred years ago meant "in the thousands." In 2009 alone 61 million cars were produced (Wikipedia, Automotive Industry). And you can bet that there are more people employed today globally in the automotive manufacturing business than there were a century ago. People are still an integral part of that process and, as the technology has become more complex and sophisticated, they have become even more important than ever. The same will be true with automation and cloud computing; as the technology matures and becomes more sophisticated and complex, people will be as essential a part of the equation as when they had to manually enter the commands themselves. They will be able to recognize when processes are inefficient, when they could be improved or applied elsewhere. They will be able to take the time to build out systems that take on the burden of mundane tasks, which is what we've always relied upon machines to do.

What will be new and should be exciting is that the people involved will actually be freed to act like **people** rather than machines. And if we're lucky, that means that the business stakeholders will stop treating them as though they're machines and start leveraging their people skills instead.

**\*** YOU GET 20 GEEK POINTS IF YOU RECOGNIZED THAT QUESTION AS ONE FROM THE MOVIE "WAR GAMES", IN WHICH AN "INTELLIGENT" COMPUTER SYSTEM IS UNABLE TO DIFFERENTIATE BETWEEN A "GAME" AND "REALITY" AND NEARLY STARTS WORLD WAR III BY LAUNCHING NUCLEAR MISSILES. KEEP TRACK OF THOSE POINTS, SOME DAY THEY MIGHT BE WORTH SOMETHING, LIKE A T-SHIRT.

Related blogs & articles:

- Will DevOps Fork?
- CIOs split on future destiny of the IT department
- How Do You Measure 'Value' in IT Innovation?
- GoogleCL Puts Another Tool in the Devops Integration and ...
- Devops: Controlling Application Release Cycles to Avoid the ...
- Infrastructure Scalability Pattern: Partition by Function or Type
- Like Garth, We Fear Change
- Agent Smith Was Right: Never Send a Human to do a Machine's Job
- The New Distribution of The 3-Tiered Architecture Changes Everything
- Applying Scalability Patterns to Infrastructure Architecture
- The Multi-Generational Datacenter: From Toddlers to Teenagers
- Cloud + BPM = Business Process Scalability