Never attribute to technology that which is explained by the failure of people

Lori MacVittie, 2012-09-05

#cloud Whether it’s Hanlon or Occam or MacVittie, the razor often cuts both ways.

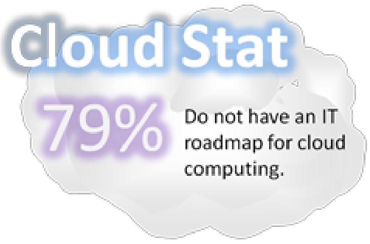

http://www.cloudtweaks.com/2011/07/cloud-computing-advantages-disadvantages-infographic/

I am certainly not one to ignore the issue of complexity in architecture nor do I dismiss lightly the risk introduced by cloud computing through increased complexity. But I am one who will point out absurdity when I see it, and especially when that risk is unfairly attributed to technology.

Certainly the complexity introduced by attempts to integrate disparate environments, computing models, and networks will give rise to new challenges and introduce new risk. But we need to carefully consider whether the risk we discover is attributable to the technology or to simple failure by those implementing it.

Almost all of the concepts and architectures being “discovered” in conjunction with cloud computing are far from original. They are adaptations, evolutions, and maturation of existing technology and architectures. Thus, it is almost always the case that when a “risk” of cloud computing is discovered it is not peculiar to cloud computing at all, and thus likely has it roots in implementation not the technology. This is not to say there aren’t new challenges or risks associated with cloud computing, there are and will be cloud-specific risks that must be addressed (IP Identity Theft was heretofore unknown before the advent of cloud computing). But let’s not make mountains out of molehills by failing to recognize those “new” risks that actually aren’t “new” at all, but rather are simply being recognized by a wider audience due to the abundance of interest in cloud computing models.

For example, I found this article particularly apocalyptic with respect to cloud and complexity on the surface. Digging into the “simple scenario”, however, revealed that the meltdown referenced was nothing new, and certainly wasn’t a technological problem – it was another instance of lack of control, of governance, of oversight, and of communication. The risk is being attributed to technology, but is more than adequately explained by the failure of people.

> The Hidden Risk of a Meltdown in the Cloud

> Ford identifies a number of different possibilities. One example involves an application provider who bases its services in the cloud, such as a cloud -based advertising service.
>
> He imagines a simple scenario in which the cloud operator distributes the service between two virtual servers, using a power balancing program to switch the load from one server to the other as conditions demand.
>
> However, the application provider may also have a load balancing program that distributes the customer load.
>
> Now Ford **imagines the scenario in which both load balancing programs operate with the same refresh period, say once a minute**. When these periods coincide, the control loops start sending the load back and forth between the virtual servers in a positive feedback loop.

Could this happen? Yes. But consider for a moment how it could happen. I see three obvious possibilities:

1. IT has completely abdicated its responsibility to governing foundational infrastructure services like load balancing and allowed the business or developers to run amokwithout regard for existing services.
2. IT has failed to communicate its overarching strategy and architecture with respect to high-availability and scale in inter-cloud scenarios to the rest of the IT organization, i.e. IT has failed to maintain control (governance) over infrastructure services.
3. The left hand of IT and the right hand of IT have been severed from the body of IT and geographically separated

3. The left hand of IT and the right hand of IT have been severed from the body of IT and geographically separated with no means to communicate. Furthermore, each hand of IT wholeheartedly believes that the other is incompetent and will fail to properly architect for high-availability and scalability, thus requiring each hand to implement such services as required to achieve high-availability.

While the third possibility might make a better “made for SyFy tech-horror” flick, the reality is likely somewhere between 1 and 2. This particular scenario, and likely others, is not peculiar to cloud. The same lack of oversight in a traditional architecture could lead to the same catastrophic cascade described by Ford in the aforementioned article.

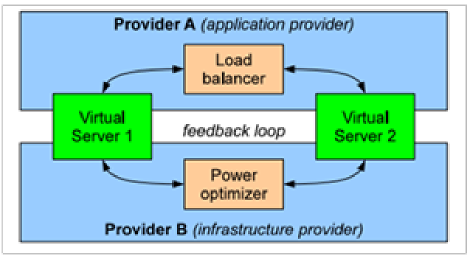

Diagram from “The Hidden Risk of a Meltdown in the Cloud”

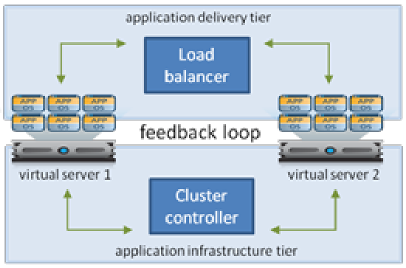

Same scenario within a traditional architecture

Given a load balancing service in the application delivery tier, and a cluster controller in the application infrastructure tier, the same cascading feedback loop could occur, causing a meltdown and inevitably downtime for the application in question.

Astute observers will conclude that an IT organization in which both a load balancing service and a cluster controller are used to scale the same application has bigger problems than duplicated services and a failed application.

This is not a failure of technology, nor is it caused by excessive complexity or lack of transparency within cloud computing environments.

It’s a failure to communicate, to control, to oversee the technical implementation of business requirements through architecture.

That’s a likely conclusion before we even start considering an inter-cloud model with two completely separate cloud providers sharing access to virtual servers deployed in one or the other – maybe both? Still, the same analysis applies – such an architecture would require willful configuration and knowledge of how to integrate the environments. Which ultimately means a failure on the part of people to communicate.

THE REAL PROBLEM

The real issue here is failure to oversee – control – the integration and use of cloud computing resources by the business and IT. There needs to be a roadmap that clearly articulates what services should be used and in what environments. There needs to be an understanding of who is responsible for what services, where they connect, with whom they share information, and by whom they will (and can be) accessed.

Maybe I’m just growing jaded – but we’ve seen this lack of roadmap and oversight before. Remember SOA? It ultimately failed to achieve the benefits promised not because the technology failed, but because the implementations were generally poorly architected and governed. A lack of oversight and planning meant duplicated services that undermined the success promised by pundits.

The same path lies ahead with cloud. Failure to plan and architect and clearly articulate proper usage and deployment of services will undoubtedly end with the same disillusioned dismissal of cloud as yet another over-hyped technology.

Like SOA, the reality of cloud is that you should never attribute to technology that which is explained by the failure of people.

- BFF: Complexity and Operational Risk
- The Pythagorean Theorem of Operational Risk
- At the Intersection of Cloud and Control…
- What is a Strategic Point of Control Anyway?
- The Battle of Economy of Scale versus Control and Flexibility
- Hybrid Architectures Do Not Require Private Cloud
- Control, choice, and cost: The Conflict in the Cloud
- Do you control your application network stack? You should.
- The Wisdom of Clouds: In Cloud Computing, a Good Network Gives You Control...

F5 Networks, Inc. | 401 Elliot Avenue West, Seattle, WA 98119 | 888-882-4447 | f5.com

F5 Networks, Inc.
Corporate Headquarters
info@f5.com

F5 Networks
Asia-Pacific
apacinfo@f5.com

F5 Networks Ltd.
Europe/Middle-East/Africa
emeainfo@f5.com

F5 Networks
Japan K.K.
f5j-info@f5.com