The Thing Private Clouds Can Do that Public Clouds Can’t

Lori MacVittie, 2009-09-10

When an admin brags they can do some task with their eyes closed there may be hidden process inefficiencies that orchestration can uncover. But the orchestration in a public cloud is effectively done for you, with little opportunity to design based on your organization's operational processes. Orchestration in a private cloud, however, is all up to you.

I was doing the laundry a few weeks ago, folding the clothes before I took them upstairs and hung them up when I realized just what I was doing. What I had been doing for, well, a very long time now. The "process" of doing laundry had become so automatic that it was ritualized, a habit. I didn't even notice that folding the clothes I was going to hang up anyway was terribly inefficient; a waste of the very limited time afforded to someone with a toddler.

Embarking on a private cloud initiative will, hopefully, illuminate something very similar with your operational processes: embedded inefficiencies. Until you have a reason to look at what has nearly become a ritualized task, you probably don't even notice there's anything "wrong" with the way you're doing things. But when you have to examine it because you're going to be codifying those processes through automation of tasks and hopefully, eventually, orchestration you might uncover some operational inefficiencies you could definitely do without. Inefficiencies that can free up time that you can spend doing other more innovative things, or at least things that might show some business value.

Only private cloud can really do this. Public cloud, with its mantra of relieving the burden of operational management and "you don't need to worry about the details" does what it promises, but also effectively removes the visibility into the very processes in which costly operational inefficiencies are deeply buried. If I outsourced the laundry and *they* folded the clothes first, I'd never notice how inefficient it is. Only because I have to "do it myself" are such wastes of my limited time uncovered and subsequently eliminated.

Does it matter? It can, depending on the process being automated. For some tasks maybe it is irrelevant and doesn't impact the total cost to deploy in the cloud at all. If it's related to auto-scaling an application, redundancies and inefficiencies in the way application instances are launched or infrastructure adjusts to the new instance may introduce enough latency to be problematic. Despite the "magical" nature of auto-scaling it isn't an instantaneous process. It takes time to boot up, start up, and hook up all the pieces of the application and its supporting infrastructure before its ready to be accessed. That time is highly dependent on the order of operations and the steps involved in completing each individual task.

VISIBILITY: MOSTLY CLOUDY

In the public cloud you don't see the gears moving and thus you'll never notice that one task is slower than the other and causes other processes to yield, waiting for some *other* part of the process to complete. And if you don't notice that one is slower you can't figure out why, and determine whether or not there's something you can do about it. You can't discover any inherent inefficiencies because you can't see the big picture. You're folding the clothes, not considering that one task is part of a "bigger" picture: doing the laundry.

Consider BitBucket's recent experience with Amazon in struggling against a targeted DDoS attack that left its site unavailable for nearly a full day. It took nearly 8 hours for merely the acknowledgement that "something was wrong" and 16-17 hours *after* the attack began before it was discovered. In the meantime, the IT folks responsible at BitBucket were sitting around waiting while someone else tried to figure out what was causing their downtime. Now to be fair to Amazon it could possibly have taken anyone that long to ferret out a UDP-based DDoS attack.

Alright, no, no it probably wouldn't. Network engineers would have seen the increase in traffic, alarms would have sounded, and after the logs showed a tremendous increase in UDP traffic someone would have done *something,* even if it was to cut off the traffic at the firewall. Processes and procedures would have been invoked that are designed to discover the source of the problem as quickly as possible, all leveraging the level of visibility you have into all aspects of your application, application network, and network operational state. It may not be perfect visibility, but it's certainly clearer than that offered by most public clouds.

What happened inside Amazon? No one knows because no one has visibility into the processes and tools they use to manage, monitor,and troubleshoot these kinds of problems. No visibility? No insight into the process, no chance to improve it in any way. Private cloud gives you the visibility into the process because *you* have to implement it, you have to specify it, you have to figure it out on your own.

There's a lot of things public cloud can do that private cloud can't, but if one of your primary motivators for cloud is in automation and orchestration, in the reduction of costs through elimination of operational inefficiencies, it may be that having someone else take care of your laundry is not the right answer.

- On our extended downtime, Amazon and what's coming
- Infrastructure 2.0 Is the Beginning of the Story, Not the End
- WILS: Automation versus Orchestration
- Your Network is Not My Network
- The Infrastructure 2.0 Trifecta
- The Revolution Continues: Let Them Eat Cloud
- Do you control your application network stack? You should.
- Dynamic Infrastructure: The Cloud within the Cloud