

Inter ‐ Cloud: Let’s talk data.



Don MacVittie, 2010-11-01

Funny thing happens when you start talking about things like inter-cloud standards, those who are looking at it from the IT guy's perspective start to see issues that are as-yet unresolved.

We have an excellent [screencast](#) on moving VMs between clouds, and [Lori](#) has written a ton about inter-cloud standards, but neither goes far enough. Yet.

George Crump of [Storage Switzerland](#) talks about moving data over on the [Network Computing Blogs](#) too, but he also is missing some important bits in the Inter-Cloud story.

Simply put, you're gonna have downtime if you're trying to do it today. Or tomorrow. Yeah, probably next year too. After that the vision gets a bit cloudy.

Your website "WeSellStuffForBigProfits.com" is not going to be around for a while, and that means "WeDontSellStuffAtAll.com" is a better name.

We at F5 have a ton of the puzzle in place – our [GTM product module](#) can dynamically redirect users to the new instance no matter how far across the globe it has moved, our [LTM product's](#) iSessions can create back-end tunnels to transfer data (assuming both cloud vendors have BIG-IP LTMs, which at this point in time is a relatively safe assumption), and we have the whitepaper on how to move your VMs and bleed off users to the new cloud provider.

But there is the problem. If there is no cutover, what to do about changing data? Whether it is in files or in a database, users in two places on an interactive system means you have two sets of data that don't match and are both changing. Bad Juju.

We can move your users, we can move your app, we can move your files and databases. What we can't do is guarantee that the new file system or database is the only place that changes are being made – because you either migrate users (and thus they're potentially updating on two systems simultaneously) or you cut them over (and they lose their connection).

But all is not lost. Several years ago, quite a few companies started approaching data replication from a new perspective – Continuous Data Protection (CDP). While most of the time CDP is overkill (every DB transaction replicated as-it-happens, in essence the call being replicated rather than the data, each write to the file system the same), moving between clouds might just be the golden problem for CDP to solve. Turn CDP on for the old DB/Filesystem and make the new DB/Filesystem the target. Then whenever someone runs a transaction or uploads a file to the old site, it is automatically copied to the new site also. I do have some questions about changes coming at the new site from both users and the old site – there is a potential there for conflict – but that's the type of stuff I would ask the CDP vendor how they resolved.

I've not tried this of course. Intercloud is not yet in such a state that you can come up with a good idea and pop-off to test it, but the theory is sound as long as both providers offer you APIs for getting at your files and data. Indeed, the Cloud Interoperability crowd should be taking steps to make certain this happens.

Why? Because for the people Inter-Cloud is supposed to serve – IT shops who don't want to be locked in – moving IPs and VMs is okay, but not a complete solution. The need is for the ability to seamlessly move applications, VMs, and users. And that won't happen if "bleeding off" users causes your data – both structured and unstructured – to become out-of-synch between source cloud and target cloud. And the CIO doesn't want to hear that there's going to be downtime for their site from the moment the move starts until it is finished. That's just not viable for most online applications.

There are still quite a few CDP vendors out there. I have in-depth knowledge of the CDP solutions for one vendor, but I'll skip mentioning them here (I have a compensated relationship with them and they're not F5, so it saves me having to put a disclaimer in my blog ;-)). You can do a bit of research into CDP and find several companies with offerings.

Replication will only take you so far... It's not real-time enough to handle things like primary/foreign key mismatches, though you can work around this, it is work, and I've seen even those workarounds (like separate ranges for primary keys that are auto-increment) fail. So we need something more, and CDP or massively distributed databases and filesystems are the only real answers.

Until next time,

Don.

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

©2016 F5 Networks, Inc. All rights reserved. F5, F5 Networks, and the F5 logo are trademarks of F5 Networks, Inc. in the U.S. and in certain other countries. Other F5 trademarks are identified at f5.com. Any other products, services, or company names referenced herein may be trademarks of their respective owners with no endorsement or affiliation, express or implied, claimed by F5. CS04-00015 0113