

Cloud Storage at Computer Technology Review



Don MacVittie, 2011-15-03

Since I've mentioned it a couple of times, I thought I'd offer you all a link to my [article](#) in [Computer Technology Review](#) about The Cloud Tier. The point was to delve into how/when/where/why of cloud storage usage. While there is a lot to say on that topic and the article was of limited word count, I think the idea that it can fit into your existing architecture with minimal changes and then be utilized to service the needs of the business in a better/faster/more agile manner is the key point.



Normally I keep my blogs relatively vendor-independent. This one talks mostly about our ARX solution because the article referenced was vendor independent, and I do think we've got some cool enabling stuff, so sometimes you just gotta talk about the cool toys. No worries, if you're not an ARX customer and won't be, there's still info in here for you.

For our part, [F5 ARX](#) is our link into that enabling story. Utilizing ARX (or another NAS virtualization engine), you can automatically direct qualifying files to the cloud, while pulling files that are accessed or modified frequently back into your tier one or tier two NAS storage. This optimizes storage usage while keeping the files available to your users. We call that point between CIFS/NFS clients and the storage they use one of our Strategic Points of Control. A spot where you can add value if you have the right tools in place. This one is a big one because files that move to the cloud can appear to users to not have moved at all – ARX has a [file virtualization](#) engine that shows them the file in a virtual directory structure. Where it is physically stored behind that virtual directory structure is completely unrelated to where the user thinks it is, and IT can move the file as needed – or write rules to have the ARX move files as needed. The only difference the user might notice is that the files they use every day are faster to open, and the ones they never or almost never access are slower.

That's called responsive, and it makes IT look good. It also means that you can say "We're using the cloud, check out our storage", and instead of planning for a huge outlay to put another array in, you can plan small monthly payments to cover the cost of storage in use. That's one of the reasons I think cloud storage for non-critical data will take off relatively quickly compared to most cloud technologies. The benefits are obvious and have numbers associated with them, they reduce the amount of "excess capacity" you have to keep in the datacenter to account for growth, while shifting the cost of that excess capacity to much smaller monthly payments.

What I don't know is how the long-term cost of cloud storage will equate to the long term cost of purchasing the same storage. And I don't think anyone can know what that graph looks like at this time. Cloud storage is new enough that it is safe to say the costs are anything but stabilized. Indeed, the race to the bottom, price-per-terabyte wise early in cloud storage's growth nearly guarantees that the costs will go up over the long term. but how far up we just don't have the information to figure out yet. Operations costs for cloud storage (from the vendor perspective) are new, and while the cost of storage over time in the datacenter is a starting point, the needs of administering cloud storage are not the same as enterprise storage, so it will be interesting to see how much long-term operation of a cloud storage vendor impacts prices over the next two to five years.

Don't let me scare you off. I still think it's a worthy addition to your tiering strategy, I would only recommend that you have a way out. Don't just assume your cloud vendor will be there forever, because there is that fulcrum where in order to survive they may have to raise prices beyond your organization's willingness (or even ability) to pay. You need to plan for that scenario, because at worst having a plan is a waste of some man-hours and maybe a minimal fee to keep a second cloud provider on-line in case, while the worst case if things go wrong is losing your files. No comparison of the risks there.

Of course, if you have ARX in place, moving between cloud storage providers is easier, but frankly, running that much data through your WAN link is going to be a crushing exercise. So plan for two options – one where you have the time to trickle the movement of data through your WAN connection (remember if you're moving cloud storage providers, all that data needs to come into the building and back out that same WAN link unless you have alternatives available), and one where you have to get your data off quickly for whatever reason.

Like everything else, there's benefit and cost savings to be had, but keep a plan in place. It's still your storage, so an addition to the Disaster Recovery plan is in order too – under “WAN link down” you need to add “what to do about non-critical files”. Likely you won't care about them in an extreme emergency, but there are scenarios where you're going to want an alternate way to get at your files.

Meanwhile, you have a theoretically unlimited tier out there that could be holding a ton of those not-so-critical files that are eating up terabytes in your datacenter. I would look into taking advantage of it.

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