

I say cloud, you say grid



Lori MacVittie, 2008-17-07

With more and more focus on cloud computing one theme seems to be running consistently: the "cloud" is public, and anyone who claims to be building a "private" cloud, a.k.a. mini-cloud or enterprise cloud, is just doing it wrong.

John Foley @ InformationWeek has it mostly right when he says that what's important is the technology.

[The Rise of Enterprise-Class Cloud Computing](#)

That's an oxymoron since cloud computing, by definition, happens outside of the corporate data center, but **it's the technology that's important here, not the semantics.** [emphasis added]

Focusing on what you call your compute environment based on whether it's public or private seems a bit silly. There are some who try to elucidate the difference between "grid" and "cloud" computing, and do manage to make a technical distinction between the two.

[RightScale's blog quoting Rich Wolski](#)

Grid computing has been used in environments where users make few but large allocation requests. For example, a lab may have a 1000 node cluster and users make allocations for all 1000, or 500, or 200, etc. So only a few of these allocations can be serviced at a time and others need to be scheduled for when resources are released. This results in sophisticated batch job scheduling algorithms of parallel computations.

Cloud computing really is about lots of small allocation requests. The Amazon EC2 accounts are limited to 20 servers each by default and lots and lots of users allocate up to 20 servers out of the pool of many thousands of servers at Amazon. The allocations are real-time and in fact there is no provision for queueing allocations until someone else releases resources. This is a completely different resource allocation paradigm, a completely different usage pattern, and all this results in completely different method of using compute resources.

Given that definition there **are** enterprises engaged in building their own mini-clouds: private, real-time on-demand data centers that service only one entity (the organization) with potentially many customers (business units).

"Private" clouds employ metering (chargebacks), rely heavily on multiple forms of virtualization, and provision resources in real-time using an on-demand model. That's a cloud as much as it is a grid.

Like SOA, it's not as if there's some certification board that's going to tell you that you're doing it "wrong", or "right" for that matter. SOA, grid, cloud - it's all about meeting the needs of the business in an operationally and financially efficient way. If that means a private cloud, than that's what you build. If that means using a public cloud, that's what you use.

Cloud computing is no more required to be public than any other computing model. It's just that - a model - and where it is implemented is of no consequence.

Besides, I thought part of cloud computing was that we weren't supposed to care about location anyway.

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