

# The Challenges of Cloud: Infrastructure Diaspora



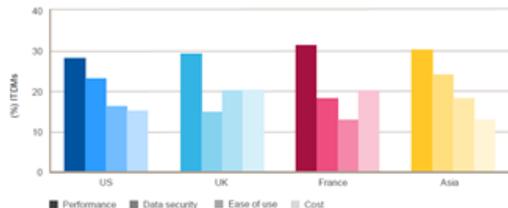
Lori MacVittie, 2012-27-08

#webperf #cloud With performance rising as a concern for [cloud computing](#) adoption, the disparity between services in the data center and the cloud needs to be addressed.

## STRONG PERFORMANCE IMPROVEMENTS ARE AT THE TOP OF ITDMs' DEMANDS ACROSS MARKETS

While data security is a concern, it is trumped by a desire to see greater, more reliable performance, and in Europe it falls behind demand to see lower costs

QUEST: WHAT SINGLE ASPECT OF YOUR CURRENT CLOUD SOLUTION NEEDS THE MOST IMPROVEMENT?



One of the negative's of cloud computing is it's one-size-fits-all approach to infrastructure. A single [load balancing](#) system (and subsequently configuration) is considered acceptable for all applications. After all, it's just about distributing requests, isn't it?

Except it isn't, and neither are myriad other infrastructure services that provide not only customized services for applications but additional benefits not currently offered by what are commoditized versions of functionality.

Even assuming an organization is using a fairly non-customized [Load balancer](#), there is a disparity between the algorithms supported by

the industry and those supported today by cloud computing providers. If you don't think [something as simple as the choice of a load balancing algorithm has an impact on availability and performance, think again](#). The reason there's a list of more than six "industry standard" algorithms is the maturation of distribution algorithms over time. Different methods are better suited to specific types of applications and usage patterns, while those same algorithms are wholly unsuited for others. Determining the best algorithm is part of the process of deploying said solutions, and one that's completely ignored by providers of cloud computing load balancing services.

Similarly, organizations that have deployed [web application firewall](#) or web filtering (web secure gateway in today's vernacular) solutions, recognize that the policies created and enforced by such solutions are not just application but URI specific, making shared, generic configurations almost completely useless. Such solutions must be deployed and configured on a per-application basis at a minimum, and the time and effort involved in doing so is generally non-trivial (though collaborative efforts around [Persistent Threat Management](#) offer a potential solution to drastically reducing the time required to configure WAF solutions for the most common threats).

## NOT JUST COSTS, CAPABILITIES

Thus when organizations look outward to the cloud, it's not just a matter of costs but also **capabilities** that becomes important. Replication of infrastructure services is beginning to be recognized as an imperative. Given the rising importance of performance as a concern for cloud computing deployments, the impact of infrastructure diaspora on application performance should be treated with the seriousness it deserves.

"I don't feel that sticking your servers out there and saying, 'OK, you've got cloud now,' is the way to go," said [Tom Hollingsworth](#), a senior network engineer with United Systems, an Oklahoma City-based value-added reseller (VAR). "I want to replicate [in the cloud with] as much functionality [customers] have for load balancers, firewalls and things like that."

Hollingsworth described a hypothetical situation where an enterprise has a mail server that has been tuned to a specific in-house load balancer and then wants to move that mail server to an IaaS provider that offers fundamentally different load balancing capabilities. Attempting to recreate those Layer 4-7 services from a data center to the cloud is complex, time-consuming and difficult to manage once you've got it up and running.

Many [IaaS providers](#) sell Layer 4-7 [cloud networking services](#) (firewalls, load balancers, application accelerators) to customers, but these services tend to be monolithic, feature-limited and in some cases proprietary.

-- [Layer 4-7 cloud networking still scarce in IaaS market](#)

There are myriad options in the TCP RFC that enable organizations to tune networking stacks to improve performance for a given application and its unique usage patterns. TCP window sizes, turning on or off Nagle, and controlling time-out values have a significant impact on not only performance but capacity of web applications. Eliminating the ability to tweak and tune these settings in a cloud computing environment removes a very important set of tools upon which the enterprise relies to address performance issues in the data center.

This infrastructure diaspora has other consequences, as well, including the introduction of a separate set of operational processes that must be managed along with existing procedures. This burdens operations with more management and monitoring duties, and introduces additional risk in the form of mis-configuration or missteps in deployment processes.

While some application delivery vendors have addressed this disparity with [cloud-enabled ADN offerings](#), these are still not universally available or supported across all cloud computing offerings. Similarly, some customers will have no complementary offerings in their own data center (if they have a data center) but will still experience the same performance-degrading scenarios which could be addressed by more robust Layer 4-7 services in cloud computing environments.

The challenge for providers is balancing costs of their services versus costs to organizations who lose revenue due to applications exhibiting poor performance when deployed in their environment. The cost-benefit analysis for enterprises will certainly include this value, and thus providers who move to address the use of more robust application delivery services as a means to redress potential performance problems will be better positioned to vie for enterprise customers for whom performance is as important – or more so – than other inhibiting concerns.

---

Referenced blogs & articles:

- [It's 2am: Do You Know What Algorithm Your Load Balancer is Using?](#)
  - [Persistent Threat Management](#)
  - [Layer 4-7 cloud networking still scarce in IaaS market](#)
  - [F5 Friday: Avoiding the Operational Debt of Cloud](#)
  
  - [The Conspecific Hybrid Cloud](#)
  - [Complexity Drives Consolidation](#)
  - [Cloud Computing and the Truth About SLAs](#)
  - [Curing the Cloud Performance Arrhythmia](#)
- 

F5 Networks, Inc. | 401 Elliot Avenue West, Seattle, WA 98119 | 888-882-4447 | [f5.com](http://f5.com)

F5 Networks, Inc.  
Corporate Headquarters  
[info@f5.com](mailto:info@f5.com)

F5 Networks  
Asia-Pacific  
[apacinfo@f5.com](mailto:apacinfo@f5.com)

F5 Networks Ltd.  
Europe/Middle-East/Africa  
[emeainfo@f5.com](mailto:emeainfo@f5.com)

F5 Networks  
Japan K.K.  
[f5j-info@f5.com](mailto:f5j-info@f5.com)