

# Filling the SLA Gap in the Cloud

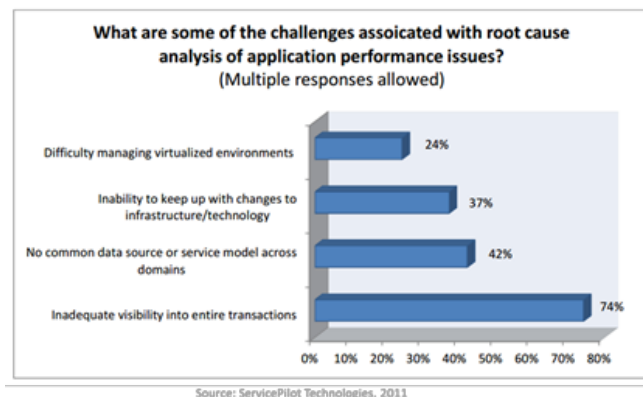


Lori MacVittie, 2012-07-11

#webperf #ado Meeting user expectations of fast and available applications becomes more difficult as you relinquish more and more control...

User expectations with respect to performance are always a concern for IT. Whether it's monitoring performance or responding to a fire drill because an application is "slow", IT is ultimately responsible for maintaining the consistent levels of performance expected by end-users – whether internal or external.

Virtualization and [cloud computing](#) introduce a variety of challenges for operations whose primary focus is performance. From lack of visibility to lack of control, dealing with performance issues is getting more and more difficult.



The situation is one of which IT is acutely aware. [A ServicePilot Technologies survey](#) (2011) indicates virtualization, the pace of emerging technology, lack of visibility and inconsistent service models as challenges to discovering the root cause of application performance issues. Visibility, unsurprisingly, was cited as the biggest challenge, with 74% of respondents checking it off.

These challenges are not unrelated. Virtualization's tendency toward east-west traffic patterns can inhibit visibility, with few solutions available to monitor intra-virtual

machines deployed on the same physical machine. Cloud computing – highly virtual in both form factor and in model – contributes to the lack of visibility as well as challenges associated with disconnected service models as enterprise and cloud computing providers rarely leverage the same monitoring systems.

Most disturbing, all these challenges contribute to an expanding gap between performance expectations (SLA) and the ability of IT to address application performance issues, especially in the cloud.

## YES, YET ANOTHER GAP

There are many "gaps" associated with virtualization and cloud computing: the gap between dev and ops, the gap between ops and the network, the gap between scalability of operations and the volatility of the network. The gap between application performance expectations and the ability to affect it is just another example of how technology designed to solve one problem can often illuminate or even create another.

Unfortunately for operations, application performance is critical. Degrading performance impacts reputation, productivity, and ultimately the bottom line. It increases IT costs as end-users phone the help desk, redirects resources from other just as important tasks toward solving the problem and ultimately delaying other projects.

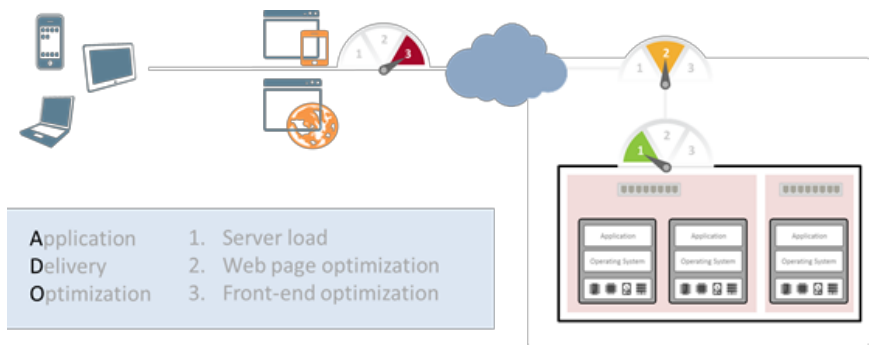
This gap is not one that can be ignored or put off or dismissed with a "we'll get to that". Application performance always has been – and will continue to be – a primary focus for IT operations. An even bigger challenge than knowing there's a performance problem is what to do about it – particularly in a cloud computing environment where tweaking QoS policies just isn't an option.

What IT needs – both in the data center and in the cloud – is a single, strategic point of control at which to apply services designed to improve performance at three critical points in the delivery chain: the front, middle, and back-end.

## FILLING THE GAP IN THE CLOUD

Such a combined performance solution is known as ADO – Application Delivery Optimization – and it uses a variety of acceleration and optimization techniques to fill the gap between SLA expectations and the lack of control in cloud computing environments.

A single, strategic implementation and enforcement point for such policies is necessary in cloud computing (and highly volatile virtualized) environments because of the topological challenges created by the core model. Not only is the reality of application instances (virtual machines) popping up and moving



around problematic, but the same occurs with virtualized network appliances and services designed to address specific pain points involving performance. The challenge of dealing with a topologically mobile architecture – particularly in public cloud computing environments – is likely to prove more trouble than it's worth. A single, unified ADO solution, however, provides a single control plane through which optimizations and enhancements can be applied across all three critical points in the delivery chain – without the topological obstacles.

By leveraging a single, strategic point of control, operations is able to leverage the power of dynamism and context to ensure that the appropriate performance-related services are applied intelligently. That means not applying compression to already compressed content (such as JPEG images) and recognizing the unique quirks of browsers when used on different devices.

ADO further enhances [load balancing](#) services by providing performance-aware algorithms and network-related optimizations that can dramatically impact the load and thus performance of applications.

What's needed to fill the gap between user-expectations and actual performance in the cloud is the ability of operations to apply appropriate services with alacrity. Operations needs a simple yet powerful means by which performance-related concerns can be addressed in an environment where visibility into the root cause is likely extremely limited. A single service solution that can simultaneously address all three delivery chain pain points is the best way to accomplish that and fill the gap between expectations and reality.

- 
- [The Need for \(HTML5\) Speed](#)
  - [Mobile versus Mobile: An Identity Crisis](#)
  - [SPDY versus HTML5 WebSockets](#)
  - [WILS: WPO versus FEO](#)
  - [WILS: Application Acceleration versus Optimization](#)
  - [F5 Friday: F5 Application Delivery Optimization \(ADO\)](#)
  - [The Four V's of Big Data](#)
  - [The "All of the Above" Approach to Improving Application Performance](#)
  - [Y U No Support SPDY Yet?](#)
- 

F5 Networks, Inc. | 401 Elliot Avenue West, Seattle, WA 98119 | 888-882-4447 | [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)