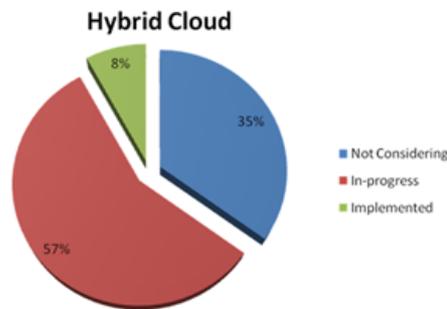


Cloud Infrastructure Integration Model: Virtualization



Lori MacVittie, 2011-19-10

Examining architectures on which hybrid clouds are based...



Source: F5 survey of 538 Enterprise IT professionals, July 2011

IT professionals, in general, appear to consider themselves well along the path toward IT as a Service with a significant plurality of them engaged in implementing many of the building blocks necessary to support the effort. IaaS, PaaS, and hybrid cloud computing models are essential for IT to realize an environment in which (manageable) IT as a Service can become reality.

That IT professionals –65% of them to be exact – note their organization is in-progress or already completed with a hybrid cloud implementation is telling, as it indicates a desire to leverage resources

from a public cloud provider.

What the simple “hybrid cloud” moniker doesn’t illuminate is *how* IT organizations are implementing such a beast. To be sure, integration is always a rough road and integrating not just resources but its supporting infrastructure must certainly be a non-trivial task. That’s especially true given that there exists no “standard” or even “best practices” means of integrating the infrastructure between a cloud and a corporate data center.

Existing standards and best practices with respect to network and site-level virtualization provide an alternative to a bridged integration model.

Without diving into the mechanism – standards-based or product solution – we can still examine the integration model from the perspective of its architectural goals, its advantages and disadvantages.

THE VIRTUALIZATION CLOUD INTEGRATION ARCHITECTURE

The basic premise of a virtualization-based cloud integration architecture is to transparently enable communication with and use of cloud-deployed resources. While the most common type of resources to be integrated will be applications, it is also the case that these resources may be storage or even solution focused. A virtualization-based cloud integration architecture provides for transparent run-time utilization of those resources as a means to enable on-demand scalability and/or improve performance for a highly dispersed end-user base.

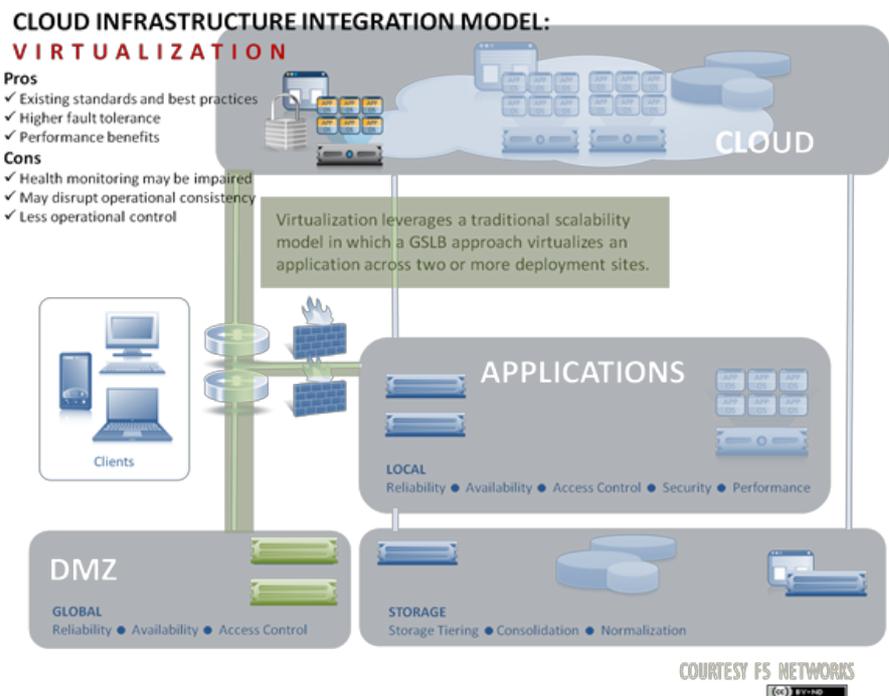
Sometimes referred to as cloud-bursting, a virtualized cloud integration architecture presents a single view of an application or site regardless of how many physical implementations there may be. This model is based on existing GSLB (Global Server load balancing) concepts and leverages existing best practices around those concepts to integrate physically disparate resources into a single application “instance”.

This allows organizations to leverage commoditized compute in cloud computing environments either to provide greater performance – by moving the application closer to both the Internet backbone and the end-user – or to enhance scalability by extending resources available to the application into external, potentially temporary, environments.

A global application delivery service is responsible for monitoring the overall availability and performance of the application and directing end-users to the appropriate location based on configurable variables such as location, performance, costs, and capacity. This model has the added benefit of providing a higher level of fault tolerance because should either site fail, the global application delivery service simply directs end-users to the available instance. Redundancy is an integral component of fault tolerant architectures, and two or more sites fulfill that need. Performance is generally improved by leveraging the ability of global application delivery services to compare end-user location, network conditions and application performance and determine which site will provide the best performance for the given user.

Because this model does not rely upon a WAN or tunnel, as with a bridged model, performance is also improved because it eliminates much of the overhead inherent in intra-environment communications on the back-end.

There are negatives, however, that can imperil these benefits from being realized. Inconsistent architectural components may inhibit accurate monitoring that impedes some routing decisions. Best practice models for global application delivery imply a local application delivery service responsible for load balancing. If a heterogeneous model of local application delivery is used (two different load balancing services) then it may be the case that monitoring and measurements are not consistently available across the disparate sites. This may result in decisions being made by the global application delivery service that are not as able to meet service-level requirements as would be the case when using operationally consistent architectural components.



This lack of architectural consistency can also result in a reduced security posture if access and control policies cannot be replicated in the cloud-hosted environment. This is particularly troubling in a model in which application data from the cloud-hosted instances may be reintroduced into corporate data stores. If data in the cloud is corrupted, it can be introduced into the corporate data store and potentially wreak havoc on applications, systems, and end-users that later access that tainted data.

Because of the level of reliance on architectural parity across environments, this model requires more preparation to ensure consistency in security policy enforcement as well as ensuring the proper variables can be leveraged to make the best-fit decision with respect to end-user access.

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