Highly Available Hybrid

Peter Silva, 2014-12-08

Achieving the ultimate 'Five Nines' of web site availability (around 5 minutes of downtime a year) has been a goal of many organizations since the beginning of the internet era. There are several ways to accomplish this but essentially a few principles apply.

- Eliminate single points of failure by adding redundancy so if one component fails, the entire system still works.
- Have reliable crossover to the duplicate systems so they are ready when needed.
- And have the ability to detect failures as they occur so proper action can be taken.

If the first two are in place, hopefully you never see a failure but maintenance is a must.

BIG-IP high availability (HA) functionality, such as connection mirroring, configuration synchronization, and network failover, allow core system services to be available for BIG-IP to manage in the event that a particular application instance becomes unavailable. Organizations can synchronize BIG-IP configurations across data centers to ensure the most up to date policy is being enforced throughout the entire infrastructure. In addition, BIG-IP itself can be deployed as a redundant system either in active/standby or active/active mode.

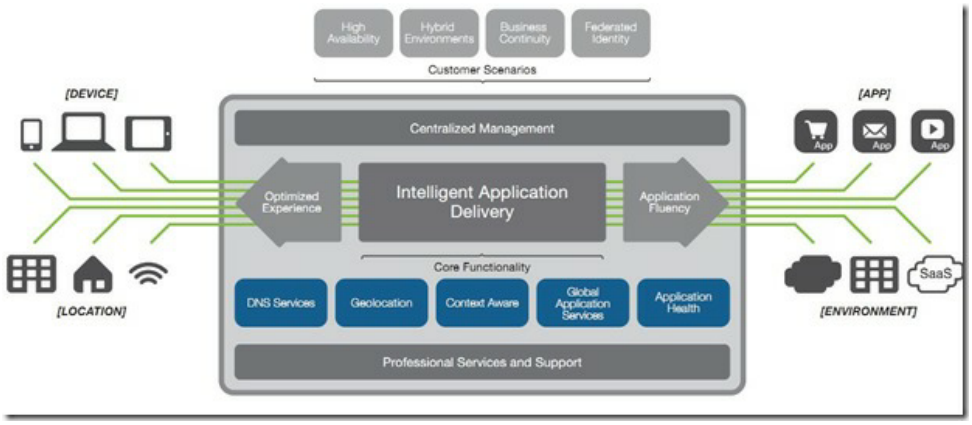

Web applications come in all shapes and sizes from static to dynamic, from simple to complex from specific to general. No matter the size, availability is important to support the customers and the business. The most basic high-availability architecture is the typical 3-tier design. A pair of ADCs in the DMZ terminates the connection; they in turn intelligently distribute the client request to a pool (multiple) of application servers which then query the database servers for the appropriate content. Each tier has redundant servers so in the event of a server outage, the others take the load and the system stays available.

This is a tried and true design for most operations and provides resilient application availability within a typical data center. But fault tolerance between two data centers is even more reliable than multiple servers in a single location, simply because that one data center is a single point of failure.

A hybrid data center approach allows organizations to not only distribute their applications when it makes sense but can also provide global fault tolerance to the system overall. Depending on how an organization's disaster recovery infrastructure is designed, this can be an active site, a hot-standby, some leased hosting space, a cloud provider or some other contained compute location. As soon as that server, application, or even location starts to have trouble, organizations can seamlessly maneuver around the issue and continue to deliver their applications.

Driven by applications and workloads, a hybrid data center is really a technology strategy of the entire infrastructure mix of on premise and off-premise data compute resources. IT workloads reside in conventional enterprise IT (legacy systems), an on premise private cloud (mission critical apps), at a third-party off-premise location (managed, hosting or cloud provider) or a combination of all three.

The various combinations of hybrid data center types can be as diverse as the industries that use them. Enterprises probably already have some level of hybrid, even if it is a mix of owned space plus SaaS. Enterprises typically like to keep sensitive assets in house but have started to migrate workloads to hybrid data centers. Financial industries might have different requirements than retail. Startups might start completely with a cloud based service and then begin to build their own facility if needed. Mobile app developers, particularly games, often use the cloud for development and then bring it in-house once it is released. Enterprises, on the other hand, have (historically) developed in house and then pushed out to a data center when ready. The variants of industries, situations and challenges the hybrid approach can address is vast.

Manage services rather than boxes.

ps

Related

- Hybrid DDoS Needs Hybrid Defense
- The Conspecific Hybrid Cloud
- The future of cloud is hybrid ... and seamless
- Hybrid–The New Normal
- Hybrid Architectures Do Not Require Private Cloud

Technorati Tags: f5,hybrid,cloud,datacenter,applications,availability,silva

Connect with Peter: Connect with F5:

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