

Three data center layout

Road map for incorporating two on premise data centers with a third data center off site and metered.

The Mission: Rely on scaling a third data center automatically if and when data center one and two become unavailable.

Citrix GSLB

Load balancing citrix.lcmchealth.org across three data centers hosting their own gateways and internal Storefront urls. Each Storefront internal url is of the set of servers physically located in that data center.

public url	citrix.lcmchealth.org		
data center	one	two	three
storefront vip	storefront-dc1	storefront-dc2	storefront-dc3
storefront servers	storefront-dc1-01 storefront-dc1-02 storefront-dc1-03	storefront-dc2-01 storefront-dc2-02 storefront-dc2-03	storefront-dc3-01 storefront-dc3-02 storefront-dc3-03

StoreFront servers set

A single group across the three data centers.

data center	one	two	three
storefront servers	storefront-dc1-01 storefront-dc1-02 storefront-dc1-03	storefront-dc2-01 storefront-dc2-02 storefront-dc2-03	storefront-dc3-01 storefront-dc3-02 storefront-dc3-03

Storefront store

A core functionality of StoreFront is the ability to aggregate and de-duplicate “common” application and desktop resources from multiple Citrix Virtual Apps and Desktops (CVAD) Sites. This functionality is commonly referred to as multi-site aggregation.

Single store will provide all apps and desktops across the three data centers. The store will have two citrix sites using site aggregation. Site one will have priority over site two.

store	one
site(s)	one, two
aggregation priority	one, then two

Citrix sites

There will be two citrix sites. One providing access to vda's found on data center one and data center two. The second site providing acces to vda's on data center 3 with automatical scaling when capacity is in need.

Site: one

Designed to provide access to datacenter one and two vda's during business as usual.

site	one		
data center	one	two	three
delivery controllers or cloud connectors	site-dc1-01 site-dc1-02	site-dc2-01 site-dc2-02	
vda's providing apps and desktops	vda01, vda02, etc.	vda01, vda02, etc.	

Site: two

With this site having con's across all three data centers the site will be available at all times. Yet its applications and desktops will only be available, primarily, if site one becomes unavailable. Once the Storefront store checks against site one and two, if one is not found and all of two's apps and desktops will be available.

As users take up connections on vda's in site two other settings will providing for production additional vda's to maintain a good user experience in real time.

Continues confirmation data center three is operating

By having con's in all three data centers we can have daily users working out of data center three vda's to understand and know with confidence data center three is fully functioning. This can be accomplished a number of ways. For example AD Group: DataCenter3Users. A small group of people can be added to this AD Group, which explicity blocks access to App1 from site one. The result being they see App1 on site two.

site	two		
data center	one	two	three
delivery controllers or cloud connectors	site-dc1-01 site-dc1-02	site-dc2-01 site-dc2-02	site-dc3-01 site-dc3-02
vda's providing apps and desktops (using automatic load scaling)			vda01, vda02, etc.

Example one: business as usual

In this example all three data centers are healthy and available. User signs in citrix.lcmchealth.org and launches application: hyperspace. Results are application is launched from site one, through vda in data center one or two.

Business as usual

User is not presented with hyperspace from site two since site one is available. User has no possible way to reach site two's hyperspace application while site one is functioning.

example		
site	one	two
application	hyperspace	hyperspace
aggregation priority	one	two
datacenter of vda(s)	one, two	three

Example two: data center one and two are unavailable

In this example all data centers one and two have become unhealthy and unavailable. User signs in citrix.lcmchealth.org and launches application: hyperspace. Results are application is launched from site two, through vda in data center three.

Disaster recovery fail

User is not presented with hyperspace from site one since site one is unavailable. User has no possible way to reach site one's hyperspace application while site one is not functioning.

Through automatical scaling additonal vda's will come online as needed to provide capacity for all users.

example		
site	one	two
application	hyperspace	hyperspace
aggregation priority	one	two
datacenter of vda(s)	one, two	three