

White Paper

Ensuring High Availability - XenApp and XenDesktop deployments on Citrix Cloud

Setup highly available XenApp and XenDesktop deployments on-prem or on public clouds

This whitepaper explains how you can use the XenApp and XenDesktop service on Citrix Cloud to deliver highly available apps and desktops on-premises or on public cloud infrastructure. It also covers disaster recovery measures taken by Citrix and how you can be notified in case of outages and planned updates

Citrix Cloud is hosted on reliable, major public cloud infrastructure

Citrix Cloud and the XenApp and XenDesktop service run on major public cloud infrastructure.

You can always give us feedback directly from within the Citrix Cloud console.

Citrix performs a nightly backup of all data in the Citrix Cloud including the XenApp and XenDesktop service. The nightly backup ensures that in the case of an outage we are able to quickly restore the environment. We also regularly test our ability to restore the service from the ground up. In addition to this, we monitor the environment 24x7 with a well defined escalation path in the event of any issues. This way, the service is built on reliable infrastructure and also is prepared to quickly recover from any outages.

As a customer, you can sign up to be notified if any outage is detected at our end and updated on the status of the outage till it resolution. We track and publish the current status of the uptime of all services. Please visit the [Citrix Cloud status page](#) to see more details and subscribe for alerts with your email address.

The next sections will go over ways you can setup your on-premises deployments or deployments in Azure or in Amazon so that they are also highly available.

Citrix Cloud and XenApp and XenDesktop service

The XenApp and XenDesktop service is built by Citrix on cloud infrastructure leveraging the public cloud in order to scale the service, scale the Desktop Delivery Controllers and all of the infrastructure, including the databases. This gives us the ability to create a highly available service. Any data, such as configuration data uploaded by the on-premise components, monitoring data that Citrix Director collects, and so on are hosted by Citrix.

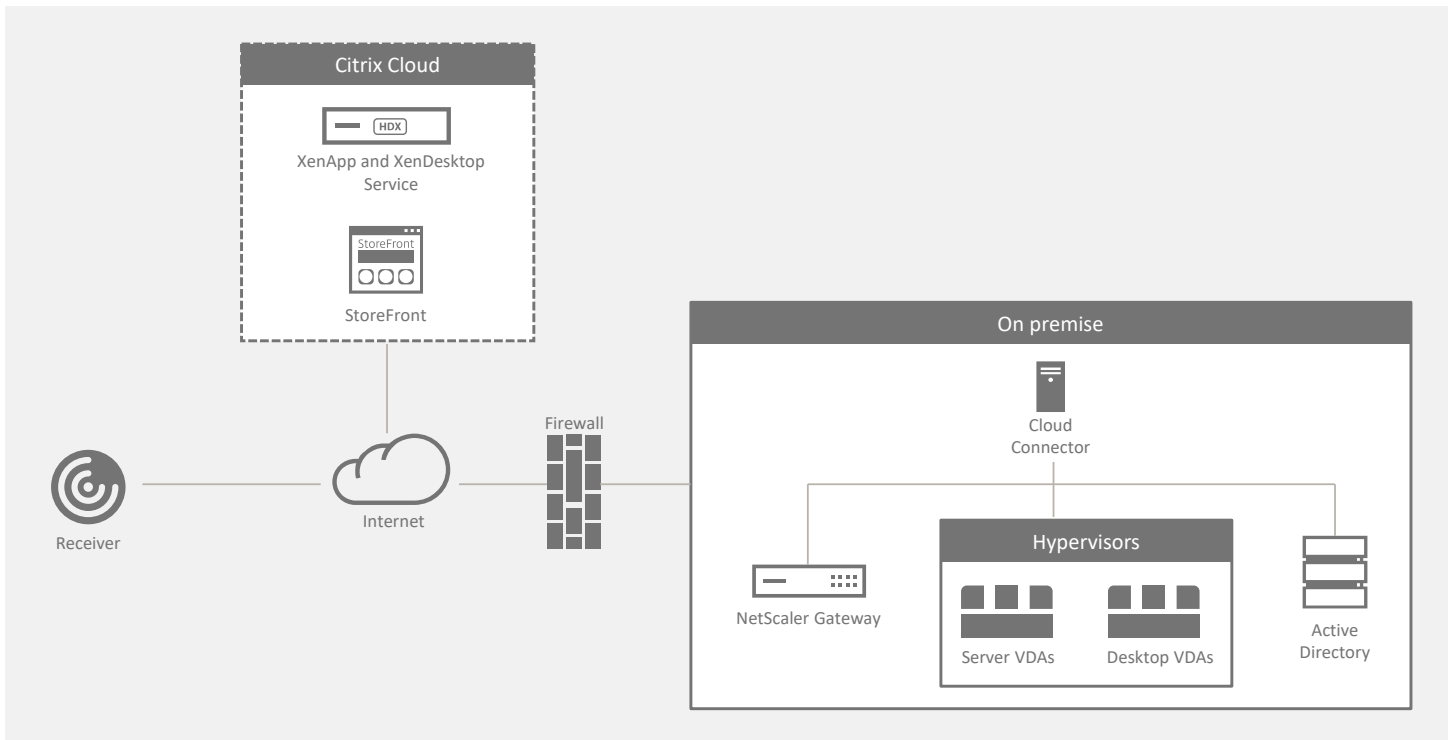


Fig 1: Sample deployment using XenApp and XenDesktop service on Citrix Cloud.

The diagram above shows a sample deployment of apps and desktops using XenApp and XenDesktop service. It uses the XenApp and XenDesktop service running on Citrix Cloud and the Virtual Delivery Agents (VDA) hosted on-premises in a Resource Location.

There may be additional configurations and ways to deploy your applications and desktops, but they all follow the same

principles outlined here. Additionally, we will cover ways to ensure high availability if you host your VDAs on public cloud infrastructure.

For more information on the service architecture and introduction to Citrix Cloud please visit the [Citrix Cloud page](#).

Setting up your resource location for high availability - on-premise, on Azure or on Amazon

You need to ensure that the Resource Locations hosting the VDAs, hypervisors and Active Directory are designed for high availability. You can setup your Resource Location on premises, in Azure or Amazon.

In general, we recommend that everything is deployed in either pairs or N+1 configuration so that there is no single point of failure. This is especially important for connectors.

The reason is that we auto-update them and want to ensure that you don't have any downtime when an update occurs. The updates are staggered amongst those connectors, so that there are never two connectors updating simultaneously. In fact, some time is given between connector updates so that we can monitor it. This will ensure that, if there is an issue, it will get rolled out to one connector, be detected and monitored before it would ever impact the other connector.

Next is the Internet connectivity. Since it is a cloud service, the Resource Locations need an Internet connection in order to use the cloud service. However, it is only needed for the initial brokering. So when the user launches their app via StoreFront or their desktop, or does a reconnection, the connector needs to be able to reach the cloud service. If the Resource Location loses connectivity, the existing sessions are not going to be impacted.

It's also important to know that, when it comes to Internet connectivity, it is not just the ISP. We also need to look at any firewalls or network filtering that may exist between the connector and the cloud. Since the service is built on the public cloud, it may use many different IP addresses, and these IP addresses may change over time since the service uses load balancers in the cloud to run the cloud service. So, we need to make sure that the connectors in the Resource Locations can reach out over port 443 to the Internet.

And finally, the last piece is to follow the recommendations of the cloud vendor that

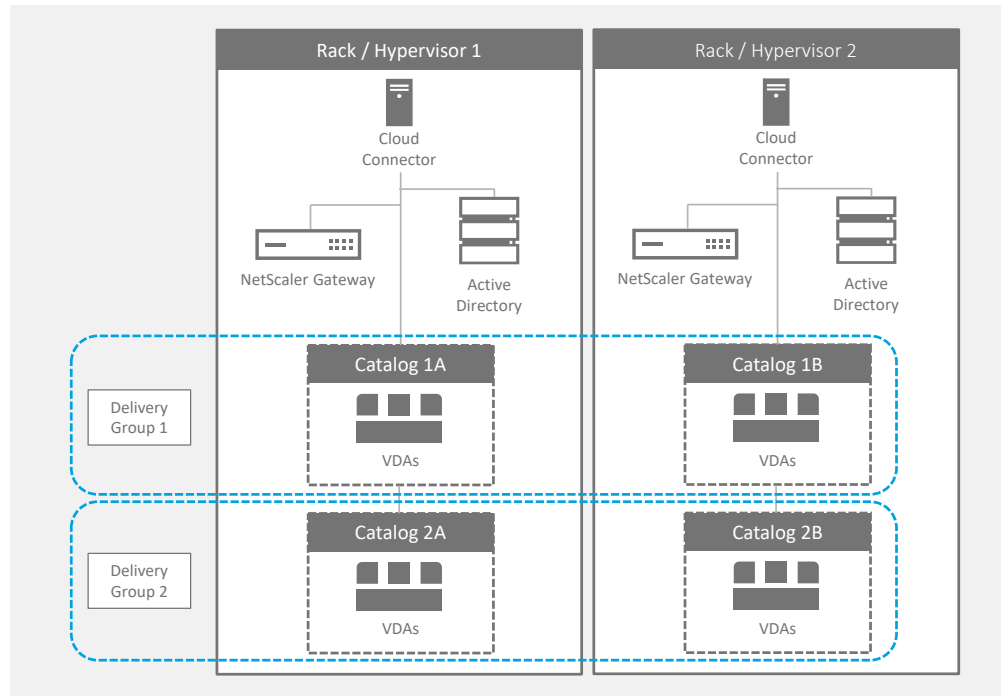


Fig 2: On-premise deployment designed for high availability.

you used to deploy the resource location, if that was the case. Let us go a little deeper into each case and explore an example deployment.

First, let us take the case of an on-premises deployment. This means that the VDAs are deployed on premises and managed from Citrix Cloud.

We recommend that you deploy in pairs or N+1. So you would have two connectors, two NetScalers, two domain controllers and for every delivery group you would have two VDAs. Now, to ensure high availability, make sure that those pairs of each component are deployed on different hypervisors. This is to eliminate a single point of failure, such as, if

one physical machine fails, or you need to update your hypervisor and you have to take it offline. If you have the components deployed on different hypervisors, you can do maintenance on one and the other one will take over.

You may go a step further and want more than just hypervisors to be isolated. We can configure the two connectors to be on totally different racks. This way, they can be on different UPSes, so they have redundant power. They can be on different network switches, so the switch isn't a single point of failure. Or go even further, and spread them out to different datacenters. This is how you can deploy a highly available, on-premises environment.

We recommend that everything is deployed in either pairs or N+1 configuration, especially the connectors

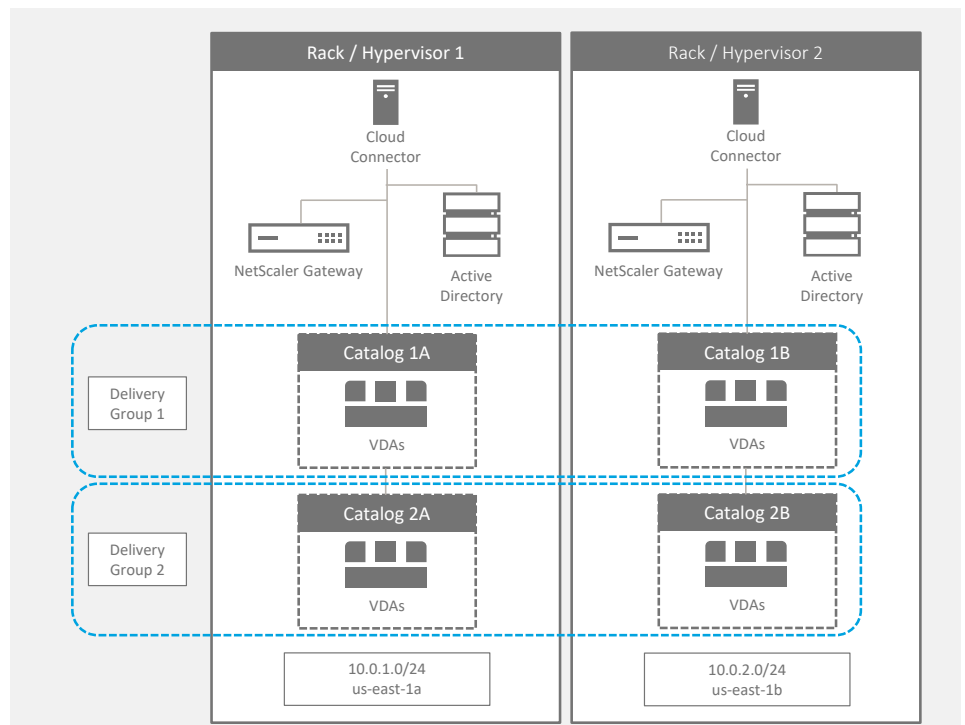


Fig 3: Deployment on Amazon designed for high availability.

Now, let us look at how we can deploy our VDAs in Amazon and manage them from Citrix Cloud, designed for high availability.

Amazon uses a concept called an "availability zone." In the example above (U.S. East), you'll see that there are two availability zones, U.S. East 1A and U.S. East 1B. That letter at the end of the region shows the availability zone. When deploying a resource location in AWS, the first thing you have to do is to pick the region. We recommend that you find the region that is closest to your users, for the best performance. Then pick two availability zones within that region in order to deploy your resource location.

In AWS, a subnet cannot span availability zones. So you would need to create a subnet in both of those availability zones and make

sure that you can route between those two subnets. The reason for this is that the connector in one zone may need to talk to one of the VDAs in the other zone. Similarly, for the gateway. A connection might come in through the gateway in one zone, and it needs to route to the VDA in the other zone. So please make sure that you have set up the routing tables to allow traffic to flow between these two subnets and availability zones.

We want to make sure that your VDAs are also highly available. So for every delivery group, you need to put some VDAs in Availability Zone A; and put some VDAs in Availability Zone B. The easiest way to do this is to create two catalogs. You can have both catalogs in the same delivery group. This way if an availability zone fails, you still have VDAs in the other catalog that will take over, and users are still able to get to their applications.

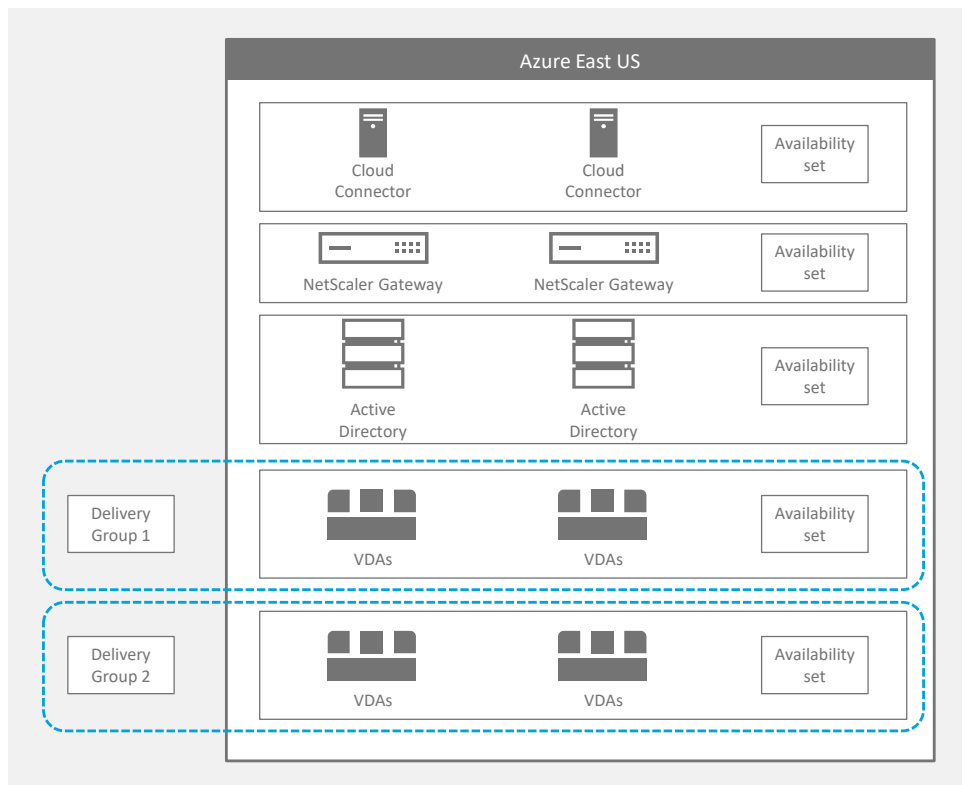


Fig 4: Deployment on Azure designed for high availability.

Finally, let us look at how we can deploy our VDAs in Azure and manage them from Citrix Cloud, designed for high availability.

Azure has a concept called "update domains" and "fault domains". Update domains are when Azure does updates to their infrastructure, such as the Hyper-V servers that are running Azure. Azure has five different update domains. So when they do updates, if you have five different VMs, and you've set up your high availability across those domains, only one will be impacted by any updates at one time.

Likewise, every region in Azure has at least three fault domains. They are for things like power, networking, cooling, etc. or the physical

infrastructure that is running the cloud is broken up into three groups in every region. So you should spread your VMs across those fault domains.

The update domains and fault domains are not exposed to you as a customer. You do not get to pick where a VM is created. Instead, Azure has a simple high-level concept called an "availability set". So if you create an availability set and add VMs in it, it tells Azure that this is a group of identical VMs. Azure will then spread them out as much as possible amongst the update domains and fault domains.

For more information, please visit the [Citrix Cloud page](#).



Enterprise Sales

North America | 800-424-8749
Worldwide | +1 408-790-8000

Locations

Corporate Headquarters | 851 Cypress Creek Road Fort Lauderdale, FL 33309 United States
Silicon Valley | 4988 Great America Parkway Santa Clara, CA 95054 United States

Copyright© 2016 Inc. All rights reserved. Citrix, the Citrix logo, and other marks appearing herein are property of Citrix Systems, Inc. and/or one or more of its subsidiaries, and may be registered with the U.S. Patent and Trademark Office and in other countries. All other marks are the property of their respective owner/s.