Introduction

Introduction

HyperFlex Stretch Cluster enables you to deploy an Active-Active disaster avoidance solution for mission critical workloads requiring high uptime (near zero Recovery Time Objective) and no data loss (zero Recovery Point Objective). A stretch cluster is a single cluster with geographically distributed nodes. Both sides of the cluster act as primary for certain user VMs. The data for these VMs is replicated synchronously on the other site. Stretch clusters enable you to access the entire cluster even if one of the sites were to completely go down. Typically these sites are connected with a low latency, dedicated, high-speed link between them.

Beginning with HyperFlex 5.5(1a) there are two Stretch Cluster models. Witness VM and Invisible Cloud Witness. The Invisible Cloud Witness also introduces the Preferred Site feature, designed to serve requests in the event that the connectivity between sites is down.

.

Stretch Cluster Invisible Cloud Witness Model

Stretch Clusters installed using HXDP 5.5(1a) will auto-configure an Invisible Cloud Witness for site arbitration, without the need for of an external witness service in a third site. Invisible Cloud Witness automatically runs the latest Arbitrator version, user maintenance of this component is not required.
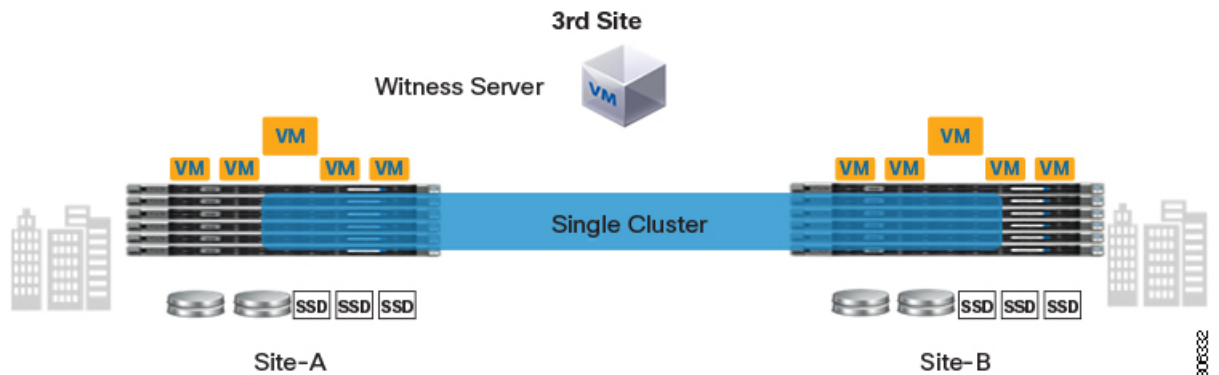

Stretch Cluster Witness VM Model

If you have an existing Stretch Cluster (HXDP Release 5.0(x) and earlier), you will continue to use the VM based witness as described in the Cisco HyperFlex Systems Stretch Cluster Guide, Release 5.0 and the Cisco HyperFlex Systems Upgrade Guide for VMware ESXi, Release 5.0.

⚠

Attention   This model only supports existing Stretch clusters deployed on HX Release 5.0(x) and earlier. All new Stretch Clusters will use Invisible Cloud Witness Model.

This chapter focuses on new Invisible Cloud Witness stretch cluster installations, provides information about monitoring stretch clusters, and discusses guidelines to follow when deploying Invisible Cloud Witness stretch clusters.