Horizon Workspace Data Protection and Disaster Recovery Best Practices

The Horizon Workspace software stores persistent data in multiple locations. Best practices guidelines for data protection and disaster recovery are described in this document.

Data protection is the practice of providing sufficient data durability and consistency so that runtime software or hardware failures do not cause data loss.

Disaster recovery addresses the catastrophic failure of primary disk groups or an entire data center where the software or hardware environment must be restored from the last backup. Disaster recovery typically implies limited data loss after the last backup. The purpose of disaster recovery planning is to minimize the amount of data loss and to guarantee data consistency after the data is restored.

NOTE This article assumes that you understand virtual machine storage management.

Data Stores and Data Protection

In a Horizon Workspace production deployment, three types of data are persistently stored:

- Data stored in an external Postgres database
- Data stored in virtual machine disk format (VMDK)
- Horizon Data blobs stored on shared NAS volumes or on VMDKs

NOTE The Postgres database is considered an external store because its performance, availability and reliability are managed outside of the Horizon Workspace software. The external Postgres database must have its own availability and reliability design, and documenting that process is outside the scope of this document.

Data Stored in VMDKs

For Horizon Workspace, all virtual machines are persisted in a VMDK. The VMDKs included are as follows:

Configurator	configurator-va.vmdk
Connector	connector-va-000001.vmdk

Data	data-va-000001.vmdk data-va_1-000001.vmdk data-va_2-000001.vmdk data-va_3-000001.vmdk data-va_4-000001.vmdk data-va_5-000001.vmdk data-va_6-000001.vmdk data-va_7-000001.vmdk data-va_8-000001.vmdk
Gateway	gateway-va-000001.vmdk
Service	service-va-000001.vmdk service-va_1-000001.vmdk

Initially all VMDKs are located in a single vSphere data-store. For performance and disaster recovery purposes, more than one data-store must be created on physically isolated backup volumes. The VMDKs for static OS and binary images can be located on slower-disk volumes. Horizon Data database, index, logs, directory and temp should be located on faster-disk volumes.

(Optional). For added protection, the backup VMDKs can reside on a data-store that is on a physically isolated volume from the other Horizon Data VMDKs.

All data-store volumes must be supported by underlying RAIDed disk groups in order to guarantee sufficient durability.

Storage for Horizon Data Blobs

The Horizon Data blob store, whether on shared NAS volumes or VMDKs, must be supported by underlying RAIDed disk groups in order to guarantee sufficient durability.

NOTE By default a VMDK is configured to store the Data blobs. The storage option you choose depends on your storage requirements. See the Horizon Workspace Installation Guide, Add Storage to the Data Virtual appliance.

Backing Up Your Data

For complete disaster recovery, the three Horizon Workspace data-store types must be backed up.

The backup method for the externally managed Postgres database is outside the scope of this document. But the overall Horizon Workspace disaster recovery plan must include the backup of the external Postgres database.

A backup set includes backups of the following data-store types:

- External Postgres database
- Point-in-time backup of the Horizon Workspace vApp for each of the virtual machines
- Point-in-time backup of the Horizon Data blob-store

In the case of catastrophic failure, the restore must be done from a single backup set to achieve consistency with the recovered systems.

Backup Sequence

To create a backup set you must backup the data in the following order:

1. Create a point-in-time backup of the externally managed Postgres database.
2. Take a point-in-time backup of all virtual machines within the Horizon Workspace vApp.
3. If using NAS for blobs, take a point-in-time backup of the Data blob-stores.

If using VMDK, Data blob-store backup is part of the backup taken in step 2.

Disaster Recovery Restoration Sequence

The single backup set takes multiple steps across physically separate components. In the case of disaster recovery, data is restored up to the point-in-time that you started the backup of the Postgres database.

System recovery after a catastrophic failure must be in the following order using the last known good backup set:

1. If using NAS, restore blobs from the blob backup. If using VMDKs, go to step 2.
2. Restore the Horizon Workspace virtual machines from the virtual machine backup.
3. Restore externally managed Postgres database.
4. Power on the Horizon Workspace vApp and the Microsoft Preview server.

NOTE The Microsoft Preview server can be recovered from a backup of the preview server virtual machine or the server can be reinstalled.

Horizon Workspace should be back to active mode.

March 2013