TECHNICAL WHITE PAPER – **DECEMBER 2016**

CONFIGURING VMWARE IDENTITY MANAGER FOR MULTIPLE DATA CENTERS

VMware Identity Manager 2.8

vmware®

Table of Contents

Introduction

VMware Identity Manager™ supports the configuration of multiple data centers in active / hot standby mode for efficient failover functionality. This paper includes detailed instructions for configuring VMware Identity Manager across multiple data centers to provide a high availability solution.

Audience

This paper is written for IT architects, consultants, and administrators who are looking for a high availability solution across multiple data centers. The reader should have some familiarity with

- VMware Identity Manager single instance and clustered environment configuration
- Databases supported by VMware Identity Manager
- SQL Server Always On Availability Groups
- Data center and high availability concepts

These topics are discussed in the VMware Identity Manager documentation and in VMware Identity Manager On-Premises Deployment Considerations.

What Is VMware Identity Manager?

VMware Identity Manager is an authentication and application access portal that provides a single point of application provisioning and entitlement for enterprise desktop and mobile users as part of VMware Workspace™ ONE™.

In particular, VMware Identity Manager provides

- Enterprise single sign-on
- Self-service app store
- Identity management with adaptive access
- Enterprise-grade hybrid cloud infrastructure

Although this paper describes on-premises deployment, VMware Identity Manager is also offered for Software-as-a-Service (SaaS) deployment with high availability features built-in.

Database Considerations

VMware Identity Manager can be set up with an internal or external database to store and organize server data. An internal PostgreSQL database is embedded in the VMware Identity Manager virtual appliance, and is the default. This embedded PostgreSQL database is useful for small deployments and does not require any additional configuration outside VMware Identity Manager. However, it is not meant to be used in large-scale or multi-site configurations, such as those described in this paper.

Multi-Data-Center Mode

You can configure VMware Identity Manager across multiple data centers (multi-data-center mode) for high availability to reduce downtime when a failover occurs.

In multi-data-center mode, VMware Identity Manager is configured as active / hot standby, as opposed to active / active. Multi-data-center mode provides a hot standby in the second data center, which becomes active in a failover scenario.

VMware Identity Manager does not perform read and write operations simultaneously in both data centers, so administrators should not allow requests to be sent to both data centers simultaneously.

In this paper, individual VMware Identity Manager virtual appliances are called *nodes*. A *cluster* consists of several nodes. Although it is possible for a data center to contain multiple clusters, one cluster per data center is the norm. Consequently, the terms cluster and data center are often used interchangeably.

Use Cases

The use cases considered in this paper all involve deployment and configuration of VMware Identity Manager in multi-data-center mode. The key differences between them involve the use of databases from different vendors: Microsoft, Postgres, and Oracle.

The following sections provide detailed instructions for these use cases:

- Multi-Data-Center Configuration with Microsoft SQL Server Database
- Multi-Data-Center Configuration with External PostgreSQL or Oracle Database

Multi-Data-Center Configuration with Microsoft SQL Server Database

Many organizations are moving toward Microsoft SQL Server database because of its high availability and disaster recovery capabilities, in particular its Always On Availability Groups feature.

Prerequisites

To configure VMware Identity Manager in multi-data-center mode with Microsoft SQL Server databases, the organization must have the following:

- Multiple data centers with low network latency between them
- SQL Server 2014 or later
- SQL Server Always On Availability Groups feature enabled
- SQL Server Always On Listener configured correctly

For information on proper configuration of these Microsoft products and features, see the Microsoft documentation.

Each cluster must be configured so that its nodes are connected with the following components:

- Elasticsearch for auditing, reports, and sync logs
- RabbitMQ for use as the messaging service
- Ehcache to cache database objects

See Figure 1 for a graphic depiction of this setup.

Note: The steps below are designed to make sure that these components are configured correctly, and that they do not accidently become connected to a node or nodes in a different cluster.

Architecture

Figure 1 shows how VMware Identity Manager can be configured to be always available with zero downtime. To take advantage of this capability, use the SQL Server Always On Listener hostname / IP address instead of database server hostname / IP address when you configure the VMware Identity Manager virtual appliance.

- The global load balancer refers to a DNS record that points to the primary load balancer. This DNS record must also be the Fully Qualified Domain Name (FQDN) for all VMware Identity Manager appliances.
- The primary load balancer manages the active cluster, where requests should be sent. The load balancer in Cluster 1 is the primary load balancer.
- The dotted arrow from the global load balancer to the load balancer in Cluster 2 indicates that upon failover, traffic is switched to the load balancer in Cluster 2.
- The dotted arrow from the SQL Server Always On Listener to the Secondary SQL Server database indicates that upon failover, database queries are switched to the Secondary SQL Server database.

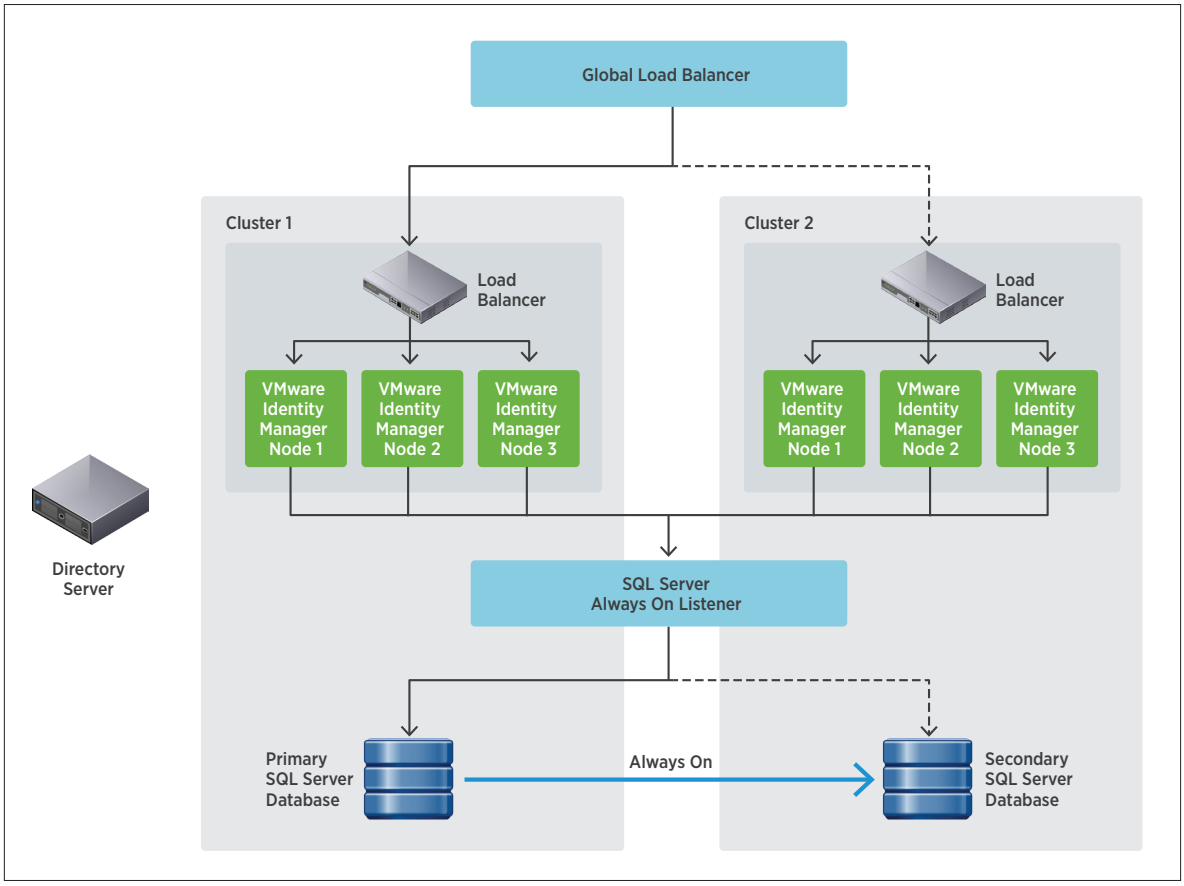

Figure 1: Multi-Site Configuration Using SQL Server

Configuring the First Cluster

Follow the instructions in VMware Identity Manager Installation and Configuration to create the first cluster, and use the following steps to configure the first cluster.

Note: Select the appropriate VMware Identity Manager version number from the drop-down menu.

Note: When you configure the database, be sure to add the hostname or IP address of the SQL server listener, for instance:

```
jdbc:sqlserver://<listener_hostname>;DatabaseName=saas
```

Configure Elasticsearch, RabbitMQ, and Ehcache

To use the following steps to configure Elasticsearch, RabbitMQ, and Ehcache on all nodes in the first cluster, first SSH to the VMware Identity Manager virtual appliance and log in with root privileges.

1. Configure Elasticsearch.
   a. Disable the cron job for Elasticsearch:
      i. Open the `hznelasticsearchsync` file:
      `vi /etc/cron.d/hznelasticsearchsync`
      ii. Comment out the cron job:
      `#*/1 * * * * root /usr/local/horizon/scripts/elasticsearchnodes.hzn`
   b. Add the IP addresses of all nodes in the cluster (that is, in the same data center):
      i. Open the `elasticsearch` file:
      `vi /etc/sysconfig/elasticsearch`
      ii. Add IP addresses of all nodes in the cluster:
      `ES_UNICAST_HOSTS=<ip_address_1>,<ip_address_2>,<ip_address_3>`
      Note: If this file already contains IP addresses, you must verify that they belong to same cluster.
   c. For Elasticsearch replication across clusters, add the FQDN of the second cluster's load balancer:
      i. Open the `runtime-config.properties` file:
      `vi /usr/local/horizon/conf/runtime-config.properties`
      ii. Add the load balancer of the second cluster:
      `analytics.replication.peers=<LB_FQDN_of_other_cluster>`
2. Configure RabbitMQ.
   a. Disable the cron job for RabbitMQ:
      i. Open the `hznrabbitmqsyn` file:
      `vi /etc/cron.d/hznrabbitmqsyn`
      ii. Comment out the cron job:
      `#*/1 * * * * root /usr/local/horizon/scripts/rabbitmqnodes.hzn`
   b. Add the node hostnames of all the nodes in the cluster (that is, in the same data center):
      `vi /usr/local/horizon/scripts/rabbitmqnodes.hzn:`

c. Comment out the following lines:

```
#if test $(curl -X GET -k
https://localhost/SAAS/API/1.0/REST/system/health/allOk -sL -w "%{http_code}\\n" -o /
dev/null) -ne 200 ; then
  #   echo SAAS not running, aborting
  #   exit 0
#fi
```

d. Comment out the following line:

```
#nodes=$(uniqList true $(enumeratenodenames))
```

e. Add the node hostnames, using only the prefix (not the FQDN), space-delimited:

```
nodes="node1 node2 node3"
```

f. Add the IP address and hostname mapping of the other nodes in `/etc/hosts`:

```
<Ip_address> node2.hs.trcint.com node2
<Ip_address> node3.hs.trcint.com node3
```

Note: This step is necessary only if there is no DNS that can resolve the PQDN and FQDN.

g. Run the following script to build the RabbitMQ cluster:

```
/usr/local/horizon/scripts/rabbitmqnodes.hzn
```

3. Configure Ehcache.

   Add the FQDN of peer nodes in the cluster, but do not add the FQDN of the current node. FQDNs are colon-delimited, for instance:

   ```
   server1.example.com:server2.example.com
   ```

   a. Open the `runtime-config.properties` file:

   ```
   vi /usr/local/horizon/conf/runtime-config.properties
   ```

   b. Add the FQDN:

   ```
   ehcache.replication.rmi.servers=node2.hs.trcint.com: node3.hs.trcint.com
   ```

4. Restart the VMware Identity Manager service:

   ```
   service horizon-workspace restart
   ```

Verify Elasticsearch, RabbitMQ, and Ehcache

Use the following steps to verify the health of Elasticsearch, RabbitMQ, and Ehcache clusters:

1. For Elasticsearch, run the following command on all nodes

   ```
   curl 'http://localhost:9200/_cluster/health?pretty'
   ```

   The result should be:

   ```
   {
     "cluster_name" : "horizon",
     "status" : "green",
     "timed_out" : false,
     "number_of_nodes" : 3,
     "number_of_data_nodes" : 3,
     "active_primary_shards" : 20,
     "active_shards" : 40,
     "relocating_shards" : 0,
     "initializing_shards" : 0,
     "unassigned_shards" : 0,
     "delayed_unassigned_shards" : 0,
     "number_of_pending_tasks" : 0,
     "number_of_in_flight_fetch" : 0
   }
   ```

2. For RabbitMQ, run the following command on all nodes:

   ```
   rabbitmqctl cluster_status
   ```

   The result should be:

   ```
   Cluster status of node 'rabbitmq@node3' ...
   [{nodes,[{disc,[ 'rabbitmq@node1','rabbitmq@node2','rabbitmq@node3']}]},
    {running_nodes,['rabbitmq@node3']},
     {cluster_name,<<"rabbitmq@node2.hs.trcint.com">>},
     {partitions,[]},
     {alarms,[{'rabbitmq@node3',[]}]}]
   ```

3. For Ehcache, find the following line in `horizon.log`:

   ```
   Added ehcache replication peer: //node3.hs.trcint.com:40002
   ```

   The hostname should match the other nodes in the cluster.

Configuring the Second Cluster

Use the following steps to configure all nodes in the second cluster:

1. Build the nodes in the second cluster by cloning them from the primary VMware Identity Manager appliance.
2. Update the networking properties.

   Note: On the first boot of the appliances, the clusters have the same configuration as primary node.

3. Update `iptables`:
   a. Verify that the `/usr/local/horizon/conf/flags/enable.rabbitmq` file is present.
      `touch /usr/local/horizon/conf/flags/enable.rabbitmq`
   b. Update the IP addresses of all nodes in second cluster by finding and replacing ALL_IPS.
      `vi /usr/local/horizon/scripts/updateiptables.hzn`
      `ALL_IPS="<ip_address_1> <ip_address_2> <ip_address_3>"`
   c. Open ports by running this script:
      `/usr/local/horizon/scripts/updateiptables.hzn`
4. Repeat the configuration and verification steps used on the first cluster.
   Note: The cron job is already disabled.

Failover and Failback

When a failure occurs in the first data center, the second data center takes over. The SQL Server Always On Listener promotes the secondary database to primary; however, you must point the global load balancer to the load balancer in the second cluster. VMware Identity Manager connects to the databases through the SQL Server Always On Listener to make a seamless transition to the second cluster.

You can fail back to the primary cluster when it becomes available again. In this instance, however, you must point the global load balancer or the DNS record to the load balancer in the primary cluster.

Clear the cache in the secondary cluster to avoid inconsistencies.

To use REST APIs to clear the cache:

```
PATH: /SAAS/jersey/manager/api/removeAllCaches
Method: POST
Roles Allowed: OPERATOR only
```

Multi-Data-Center Configuration with External PostgreSQL or Oracle Database

Some organizations prefer PostgreSQL, Oracle, or other databases. When VMware Identity Manager is configured across multiple data centers for high availability with PostgreSQL or Oracle databases, the VMware Identity Manager appliances in the secondary data center are configured in read-only mode. Upon failover, users can log in and launch applications from the VMware Identity Manager portal. However, because VMware Identity Manager appliances in the secondary data center are in read-only mode, most administrative operations, such as adding users or apps, or entitling users, no longer work after a failover to that data center.

To make the secondary data center fully functional in read / write mode, you must restart the VMware Identity Manager service after a failover event.

Note: Because requests cannot be sent to nodes in both clusters, VMware Identity Manager does not provide an active / active solution.

Prerequisites

To configure VMware Identity Manager in multi-data center mode with external PostgreSQL or Oracle databases, the organization must have the following items:

- External PostgreSQL or external Oracle databases
- A master-slave database configuration
- Multiple data centers with low network latency between them

In addition, the following components must be configured on all nodes in each data center:

- Elasticsearch for auditing, reports, and sync logs
- RabbitMQ for use as the messaging service
- Ehcache to cache database objects

See Figure 2 for a graphic depiction of this setup.

Note: The steps below are designed to make sure that these components are configured correctly, and that they do not accidently become connected to a node or nodes in a different cluster.

Architecture

Figure 2 shows how VMware Identity Manager can be configured with external PostgreSQL or Oracle databases.

- The global load balancer refers to a DNS record that points to the primary load balancer. This DNS record must also be the Fully Qualified Domain Name (FQDN) for all VMware Identity Manager appliances.
- The primary load balancer manages the active cluster, where requests should be sent. The load balancer in Cluster 1 is the primary load balancer.
- The dotted arrow from the global load balancer to the load balancer in Cluster 2 indicates that upon failover, traffic is switched to the load balancer in Cluster 2.

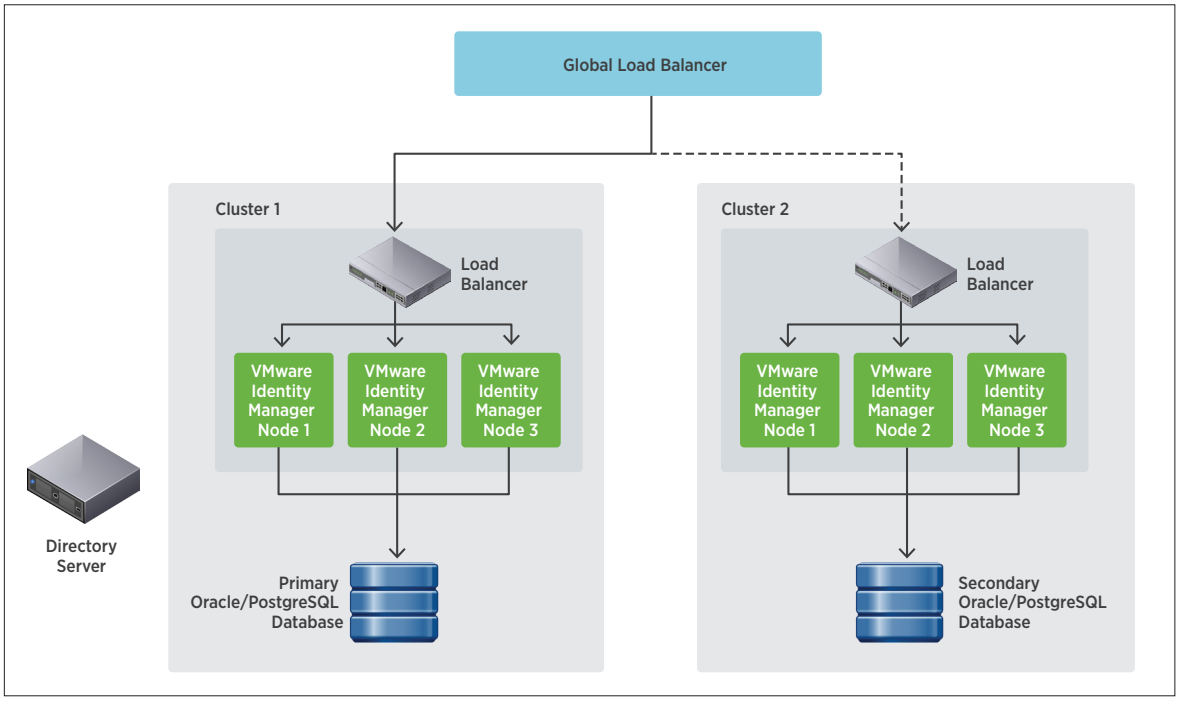

Figure 2: Multi-Site Configuration Using External PostgreSQL or Oracle Databases

The steps to configure these three components are exactly the same as those for configuring VMware Identity Manager AlwaysOn with Microsoft SQL Server database, except that nodes in the second cluster are in read-only mode. See Configuring the First Cluster and Configuring the Second Cluster.

To configure the nodes in the second cluster in read-only mode:

1. Edit the `runtime-config.properties` file:
   `vi /usr/local/horizon/conf/runtime-config.properties`
2. Add the following line:
   `read.only.service=true`

Failover and Failback

When a failure occurs in the first data center, the second cluster (in the second data center) takes over; however, the second cluster is in read-only mode at this point. You must point the global load balancer to the load balancer in the second cluster.

You can continue to use VMware Identity Manager in read-only mode if the failover is temporary. To enable read-write mode on the failed-over cluster, take the following steps:

1. Update `runtime-config.properties`:
   a. Edit the `runtime-config.properties` file.
      `vi /usr/local/horizon/conf/runtime-config.properties to`
   b. Set the following value:
      `read.only.service=false`
2. Restart the VMware Identity Manager service:
   `service horizon-workspace restart`

Failback procedures are the same for all configurations described in this paper.

You can fail back to the primary cluster when it becomes available again after a failover event. In this instance, however, you must point the global load balancer or the DNS record to the load balancer in the primary cluster.

To avoid inconsistencies, clear the cache in the secondary cluster.

To use REST APIs to clear the cache, set the following values:

```
PATH: /SAAS/jersey/manager/api/removeAllCaches
Method: POST
Roles Allowed: OPERATOR only
```

Troubleshooting

The following practices and procedures apply to all configurations discussed in this paper.

RabbitMQ

If RabbitMQ does not start, or if the health URL (`https://hostname/SAAS/API/1.0/REST/system/health/`) shows `"MessagingConnectionOk":"false":`

1. Verify that ports 4369, 5700, and 25672 are open.
   a. To open ports, run the following command:
      `touch /usr/local/horizon/conf/flags/enable.rabbitmq`
   b. Run the following script:
      `/usr/local/horizon/scripts/updateiptables.hzn`
2. Restart RabbitMQ:
   a. Kill existing RabbitMQ processes, if any.
      `rabbitmqctl stop`
      `rabbitmq-server –detached`
   b. You might also need to restart the VMware Identity Manager service if RabbitMQ does not come up gracefully:
      `service horizon-workspace restart`

Elasticsearch

If Elasticsearch does not start correctly and status is red:

1. Verify that port 9300 is open.
2. Restart Elasticsearch on all nodes in the cluster:
   `service elasticsearch restart`
3. Check logs for more details:
   `cd /opt/vmware/elasticsearch/logs`
   `tail -f horizon.log`

Summary

This paper offers concise instructions for setting up a high availability solution using VMware Identity Manager. Because there are differences among major database vendors, separate instructions are provided for the most commonly used configurations.

Although VMware Identity Manager contains an internal PostgreSQL database, no internal database is suitable for failover scenarios. External PostgreSQL and Oracle databases, among others, provide this functionality. Microsoft SQL Server with Always On Availability Groups provides additional functionality that enables automatic failover.

Additional Resources

VMware Identity Manager documentation

VMware Identity Manager On-Premises Deployment Considerations

VMware Knowledge Base article: Using embedded vPostgres in Production for VMware Workspace Portal VA 2.1 (and VMware Identity Manager 2.4) (2094258)

Microsoft documentation

About the Author and Contributors

This paper was written by Pitambari Parekh, Software Engineer, VMware End-User Computing, who also devised and tested the procedures. She wishes to thank the following contributors and reviewers:

- Sridevi Ravuri, Senior Director, Research and Development, VMware End-User Computing
- Ravi Chayanam, Senior Manager, Research and Development, VMware End-User Computing
- Karen Zelenko, Staff Program Manager, VMware End-User Computing
- Gregory Armanini, Senior Product Line Manager, VMware End-User Computing
- Fanny Strudel, Senior Manager, Research and Development, VMware End-User Computing
- Michael Almond, Staff Engineer, VMware End-User Computing
- Gary Sloane, VMware Consulting Editor

To comment on this paper, contact VMware End-User-Computing Technical Marketing at euc_tech_content_feedback@vmware.com.

VMware, Inc. 3401 Hillview Avenue Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001 www.vmware.com