Dynamic Workload Scaling

Last Updated: July 11, 2011

About the Authors

Brian Howard

Brian Howard, Software Q/A Engineer, Systems Architecture and Strategy Unit (SASU), Cisco Systems

Brian is a Software Quality Assurance testing engineering in SASU focusing on new and innovative Data Center Interconnect (DCI) technologies. Recent DCI design and test efforts include OTV, A-VPLS, Nexus 1000v and the Virtual Security Gateway, all components utilized in the DCI systems. Brian has been providing quality initiatives and testing in Advanced Services and the Cisco Corporate Development Office for 11 years, focusing primarily on routing and switching, and most recently, in Data Center virtualization using DCI.

Soumen Chatterjee

Soumen Chatterjee, Technical Marketing Engineer, Data Center Application Network Services (DCANS), Cisco Systems

Soumen is a Technical Marketing Engineer (TME) focusing on advanced Data Center technologies and architectures. His primary focus utilizes the Application Control Engine (ACE) to accelerate and optimize application delivery and server load balancing within the Data Center. Soumen has recently completed an MBA from the Anderson Business school at the University of California, Los Angeles (UCLA).

Dynamic Workload Scaling

Dynamic Workload Scaling

Today's data center designs are more focused on the ability to share computing resources between data centers. In the past, data centers were typically linked together using a Layer 3 transport technology into a service provider cloud, which kept data centers isolated. Network events happening in one data center would not affect users in other data centers. However, with the rapid development of server virtualization technology, extending Layer 2 functionality between data centers has become the norm. Being able to extend L2 networks between data centers provides network administrators the ability to dynamically allocate computing resources from one data center to another without disruption to the existing network traffic. Cisco's Dynamic Workload Scaling (DWS) is one way to achieve this essential and flexible resource allocation.

The Systems Architecture and Strategy Unit (SASU) collaborated with the Data Center Application Networking Services (DCANS) team to develop and execute a series of test cases that introduce DWS technology and how it can be utilized over a Data Center Interconnect (DCI) enabled system.

DWS Components

The Cisco Application Control Engine (ACE) along with Cisco Nexus 7000 series switches and VMware vCenter provides data center administrators with a unique set of tools that allow them to dynamically allocate computing resources between data centers.

The components of DWS need to work together to

- Present an "extended" server farm consisting of multiple VMs between two or more data centers
- Establish and monitor the locality of those real servers in the extended server farm; i.e. are the real servers "local" or "remote" from the perspective of the load balancer
- Extend the Layer 2 domain between the data centers to enable seamless connectivity between the VIP and the servers in the server farm
- Establish and monitor certain resource utilization thresholds which are used as decision points by DWS for sending incoming client requests only to real servers in the "local" data center or sending them also to real servers in the "remote" data centers

Cisco ACE-30

The ACE-30 load balancer is the backbone of Cisco's DWS, allows the network administrator to dynamically monitor CPU loads and memory resources on the VMware Virtual Machines, through ties into vCenter. The network administrator has the ability to configure predefined thresholds on the ACE-30 to monitor CPU and memory utilization of the local server farms. If those thresholds are met or exceeded, the ACE-30 will forward new client-to-server traffic to server farm computing resources in the remote datacenter. The ACE-30 learns the "local" or "remote" proximity of these server farm resources from the Nexus 7000.

Nexus 7000

The Nexus 7000 data center switches provide the infrastructure and technology to extend the Layer 2 domain between the two (or more) data centers using Overlay Transport Virtualization (OTV) technology. OTV maintains awareness of the relative "location" of host MAC addresses within the overlay. Thus, the ACE-30 is able to use SNMP to poll the Nexus 7000 OTV subsystem and determine whether the particular server IPs are "local" or "remote."

VMware vCenter

As mentioned above, vCenter communicates the CPU and memory status of the VM real servers in the server farm to the ACE-30, which then makes the decision to dynamically send incoming connections to "remote" servers in the server farm. vCenter also provides the administrative environment through which a VM can be moved, using VMotion, from one location or data center to another.

Focus of This Paper

The goal of this document is to show a few examples of how DWS might be implemented by network administrators to make more efficient use of data center resources. This paper is intended to not only present DWS, but also to review several examples of configuration and operation.

Use Cases Overview

There are three DWS use cases that are covered in this white paper. They are:

- DWS Using ACE-30 Roundrobin Load Balance Predictor
- DWS Using ACE-30 Leastconns Load Balance Predictor
- VMware Vmotion with DWS

Basic Configuration

All of the use cases covered in this paper will utilize the topology depicted in the diagram below. It consists of two data centers, separated by a distance of 100 km (simulated using a delay generator). A single Nexus 7000 switch provided the OTV-based connectivity between the data centers over a virtual

private cloud. There is one ACE-30 (a module living in a Catalyst 6500) that controls traffic to the server farm that is essentially extended between the two data centers. The vCenter server lives in the same data center as the ACE-30.

Figure 1 DWS Topology Overview

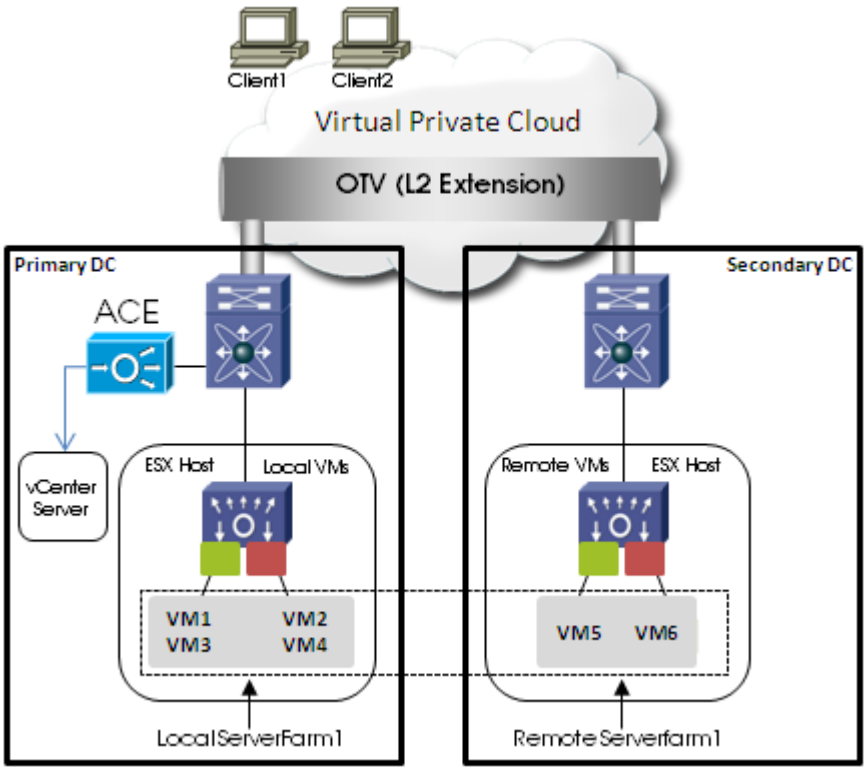

ACE-OTV

OTV extends the layer 2 VLANs between the data centers through the Cloud.

Note Cisco OTV is only supported on the Cisco Nexus 7000.

The servers that are configured in the ACE-30 server farm can exist in any VLAN that is being extended by OTV. These same server VLANs also must be trunked to the ACE-30 module residing in the Catalyst 6500. The Cisco ACE-30 polls the OTV information from the Nexus 7000 every 60 seconds to determine which VM real servers in the server farm are local and which are remote.

In the validation set up, the ACE-30 determined that rservers VM1-VM4 are in the local data center (the "[L]" flag) and VM5 and VM6 are in the remote data center (the "[R]" flag). This information was gathered by determining through SNMP polling of the Nexus 7000 which server MAC addresses resided locally or were learned from OTV's overlay interface. In the example below, rserver VM1, VM2, VM5 and VM6 are configured in Vlan2501. VM3 and VM4 reside in Vlan2502. Vlan 2501 and Vlan2502 are also trunked to the ACE-30, as well as being extended between the data centers via OTV.

The following was configured in the admin context on the ACE-30, allowing SNMP polling of the necessary MIBs on the Nexus 7000.

```
nexus-device  dc1c-agg-7k1
  ip-address 10.0.189.37
  credentials admin Cisco
```

Below is the configuration captured from user-defined context of the ACE.

```
dc1c-ace-s9/cap-exp# show  serverfarm SRV1

   Codes: L - local,   R - remote

 serverfarm     : SRV1, type: HOST
 total rservers : 6
 state          : ACTIVE
 DWS state      : ENABLED_REMOTE_LB(Bursting traffic to local and remote VMs)
----------------------------------------------------------connections-----------
       real               weight     state    current     total      failures
---+--------------------+-----+------------+----------+----------+---------------
   rserver: VM1
       10.25.1.11:0          8   OPERATIONAL [L] 0          0          0
   rserver: VM2
       10.25.1.12:0          8   OPERATIONAL [L] 0          0          0
   rserver: VM3
       10.25.2.11:0          8   OPERATIONAL [L] 0          0          0
   rserver: VM4
       10.25.2.12:0          8   OPERATIONAL [L] 0          0          0
   rserver: VM5
       10.25.1.21:0          8   OPERATIONAL [R] 0          0          0
   rserver: VM6
      10.25.1.22:0           8   OPERATIONAL [R] 0          0          0
```

ACE-vCenter

The ACE-30 also polls the VMServer MIBs from vCenter to determine each VM's CPU and memory utilization. The ACE-30 then performs an aggregate calculation of all the servers in the server farm configured for DWS and compares it to the configured threshold limits. If those limits are met or exceeded, any new client connections coming into the data center will be load balanced to the servers in the remote data center, based on the configured predictor.

The configuration for accessing vCenter information from the ACE-30 is below.

```
vm-controller VCENTER-DC1
  credentials administrator Cisco
  url https://10.0.189.1/sdk (VCenter's IP address)
```

Once the ACE knows how to reach the vCenter machine, then configuration and application of the probe to the server farm is necessary (Note: this configuration is done from user-defined context of the ACE).

```
probe vm VCENTER-DC1
  load cpu burst-threshold max 50 min 25
  vm-controller VCENTER-DC1

serverfarm host SRV1
  dws burst probe VCENTER-DC1
  rserver VM1
    inservice
<snip>
  rserver VM6
    inservice
```

The above probe configuration sets the minimum CPU threshold to 25% and maximum CPU threshold to 50%. Once the aggregate of all local servers in server farm SRV1 exceed or meet the CPU 50%, the ACE-30 will start to load balance new connections to the remote servers VM5 and VM6. The aggregate CPU calculation is performed by the ACE-30 and is determined by polling VMware vCenter SNMP performance MIBs for each of the virtual machines. Also, note that the minimum threshold is configured so that once the CPU utilization drops below 25%, any new client connections would again be load balanced only to the local data center. The minimum threshold only kicks in once the maximum threshold is met or exceeded.

ACE Load Balancing

Depending on the load-balancing algorithm—or *predictor*—that is configured, the ACE performs a series of checks and calculations to determine which server can best service each client request. The ACE is able to base server selection on several factors including the source or destination address, cookies, URLs, HTTP headers, or the server with the fewest connections with respect to load.

The following **serverfarm** configuration is done from the user-defined context on the ACE and is used to define which load balancing predictor will be used for that server farm.

```
dc1c-ace-s9/cap-exp(config-sfarm-host)# predictor ?
  hash             Configure 'hash' Predictor algorithms
  least-bandwidth  Configure 'least bandwidth' Predictor algorithm
  least-loaded     Configure 'least loaded' predictor algorithm
  Leastconns       Configure 'least conns' Predictor algorithm
  response         Configure 'response' Predictor algorithm
  roundrobin       Configure 'RoundRobin' Predictor algorithm (default)
```

Note By default, **roundrobin** is the default load balancing algorithm on the ACE-30 and is not shown in the running configuration unless manually configured. For more information on Predictor Load Balancing definitions:

http://www.cisco.com/en/US/docs/interfaces_modules/services_modules/ace/vA4_2_0/configuration/slb/guide/overview.html#wp1000976

Two predictors will be in focus for the first two use cases presented in this paper. The **roundrobin** predictor sends all incoming connections to each real server in the server farm, one by one in succession, irrespective of how many connections may already exist on any one of the real servers. The **Leastconns** predictor will send incoming connections to those servers that have the least number of existing connections, until all servers have an equal number of connections.

Tip The assumed goal of DWS is to ease the load on “local” real servers by having “remote” real servers available in another data center. As will be demonstrated below, the **Leastconns** predictor is best at delivering on this goal at this time. With **Leastconns**, after the “remote” real servers in the extended server farm have been activated, any new incoming connection will be sent to the “remote” servers, as they have significantly less existing connections than the “local” servers. With the **roundrobin** predictor, the already taxed “local” servers will continue to receive connections, burdening them even more.

Use Cases in Practice

As mentioned above, the following use cases will show how DWS functions with the **roundrobin** and **Leastconns** load balancing predictors, as well as how DWS functions when a VM is moved (using VMotion) from one data center to another. A combination of narrative, diagrams and device (CLI) output is used to walk through each of these use cases.

The aggregate CPU threshold for the server farm probe was configured with 50% as a maximum and 25% as a minimum. Again, this means that after 50% aggregate threshold is reached on the local servers, DWS would kick in and the ACE would begin to include remote servers in the load balancing of incoming connections. If the aggregate threshold drops below 25%, the ACE would stop sending to remote servers and send new connections only to local servers again. These maximum and minimum thresholds are seen in the following ACE output, showing server farm probe detail:

```
dc1c-ace-s9/cap-exp# show probe VCENTER-DC1 detail

 probe       : VCENTER-DC1
 type        : VM
 state       : ACTIVE
 description :
---------------------------------------------
   interval          :  10      vm-controller :  VCENTER-DC1
   cpu-load:
      burst-threshold:
       max threshold :  50      min threshold :  25
   mem-load:
      burst-threshold:
       max threshold :  99      min threshold :  99
                 ------------------ probe results ------------------
   associations   ip-address     cpu-load  mem-load   health
   ------------ ---------------+---------+---------+----------------------
   serverfarm  : SRV1
                aggregate-stats    50       53       BURST_REMOTE

   No. Passed probe    :      60874         No. Failed probe  : 0
   No. Probes skipped  :          0         Last status code  : 0
   Last probe time     : Wed Jun  1 14:57:29 2011
   Last fail time      : Never

     real      : VM1[0]
                 10.25.1.11         50       52       SUCCESS

     real      : VM2[0]
                 10.25.1.12         52       54       SUCCESS

     real      : VM3[0]
                 10.25.2.11         51       55       SUCCESS

     real      : VM4[0]
                 10.25.2.12         49       54       SUCCESS
```

DWS Using Roundrobin Predictor

In this use case, the **roundrobin** predictor was configured on the ACE-30.

```
dc1c-ace-s9/cap-exp# show serverfarm SRV1 detail

   Codes: L - local,   R - remote

 serverfarm     : SRV1, type: HOST
```

```
 total rservers : 6
 state          : INACTIVE
 DWS state      : ENABLED_REMOTE_LB(Bursting traffic to local and remote VMs)
 active rservers: 6
 description    : -
 predictor      : ROUNDROBIN
 failaction     : -
 back-inservice    : 0
 partial-threshold : 0
 num times failover       : 0
 num times back inservice : 0
 total conn-dropcount : 0

<snip>
```

Two clients established connections to four servers in the local data center. All of the two clients' connections were load balanced between the four servers in the local data center. The aggregate CPU utilization was approximately 50%.

```
 dc1c-ace-s9/cap-exp# show probe VCENTER-DC1 detail

 probe       : VCENTER-DC1
 type        : VM
 state       : ACTIVE
 description :
----------------------------------------------
   interval          :  10      vm-controller :  VCENTER-DC1
   cpu-load:
      burst-threshold:
       max threshold :  50      min threshold :  25
   mem-load:
      burst-threshold:
       max threshold :  99      min threshold :  99
                 ------------------ probe results ------------------
   associations   ip-address     cpu-load  mem-load   health
   ------------ ---------------+---------+---------+-----------------------
   serverfarm  : SRV1
               aggregate-stats    50       53       BURST_REMOTE

   No. Passed probe    :      60874         No. Failed probe  : 0
   No. Probes skipped  :          0         Last status code  : 0
   Last probe time     : Wed Jun  1 14:57:29 2011
   Last fail time      : Never

     real      : VM1[0]
               10.25.1.11         50       52       SUCCESS

     real      : VM2[0]
               10.25.1.12         52       54       SUCCESS

     real      : VM3[0]
               10.25.2.11         51       55       SUCCESS

     real      : VM4[0]
               10.25.2.12         49       54       SUCCESS
```

A look at the real servers in server farm SRV1 shows that the four local servers are currently handling client connections while the two remote servers are not yet.

```
dc1c-ace-s9/cap-exp# show serverfarm SRV1

   Codes: L - local,   R - remote

 serverfarm     : SRV1, type: HOST
```

```
total rservers : 6
state          : ACTIVE
DWS state      : ENABLED_REMOTE_LB(Bursting traffic to local VMs)

                                                  ----------connections-----------
      real                  weight state        current    total     failures
---+--------------------+-----+------------+----------+----------+---------
  rserver: VM5
      10.25.1.21:0          8   OPERATIONAL [R] 0          0          0
  rserver: VM6
      10.25.1.22:0          8   OPERATIONAL [R] 0          0          0
```

Figure 2 represents the state of the DWS topology at this point, prior to additional client traffic coming in.

Figure 2 DWS With Roundrobin Predictor, Pre-Burst

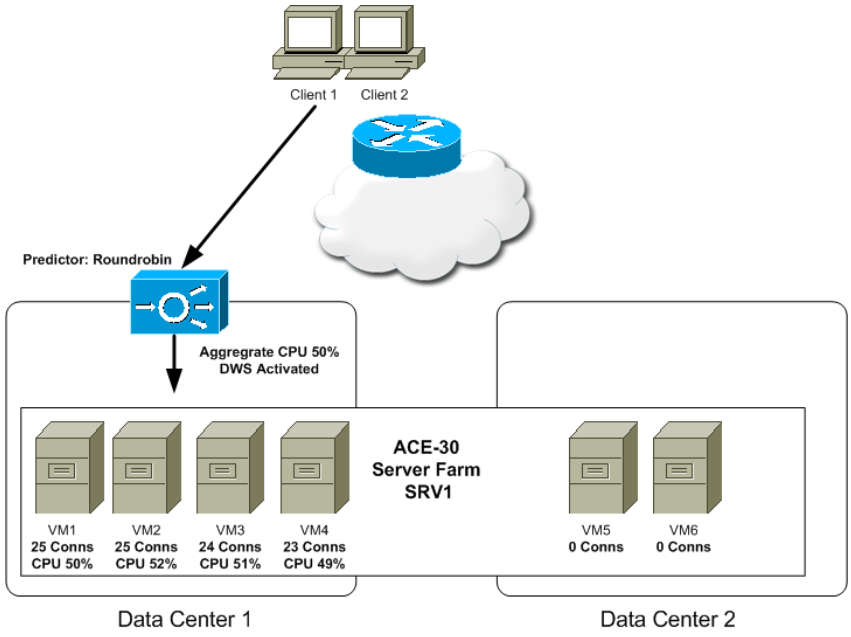

Now, Client 3 begins a new session and creates new connections into the data center and to the server farm. The new connections are load balanced equally to servers in both the local data center and the remote data center due to the **roundrobin** predictor. Note that the aggregate CPU utilization of the server in the local data center increases, as new connections are load balanced between both the local and remote data center resources.

```
dc1c-ace-s9/cap-exp# show probe VCENTER-DC1 detail

 probe       : VCENTER-DC1
 type        : VM
 state       : ACTIVE
 description :
----------------------------------------------
   interval          :  10      vm-controller :  VCENTER-DC1
   cpu-load:
      burst-threshold:
       max threshold :  50      min threshold :  25
   mem-load:
      burst-threshold:
```

```
    max threshold :  99      min threshold :  99
              ------------------ probe results ------------------
  associations   ip-address     cpu-load  mem-load   health
  ------------ ---------------+---------+---------+-----------------------
  serverfarm  : SRV1
              aggregate-stats    61       66       BURST_REMOTE

  No. Passed probe    :      60909         No. Failed probe  : 0
  No. Probes skipped  :          0         Last status code  : 0
  Last probe time     : Wed Jun  1 15:03:20 2011
  Last fail time      : Never

    real      : VM1[0]
              10.25.1.11         63       72       SUCCESS

    real      : VM2[0]
              10.25.1.12         64       63       SUCCESS

    real      : VM3[0]
              10.25.2.11         62       67       SUCCESS

    real      : VM4[0]
              10.25.2.12         58       65       SUCCESS
```

The ACE shows that now all real servers, local and remote, are receiving these new client connections. This is one of the disadvantages of the **roundrobin** predictor, compared with the **Leastconns** predictor: Servers already taxed with existing connections are asked to take on even more load.

```
dc1c-ace-s9/cap-exp# show serverfarm SRV1

   Codes: L - local,   R - remote

 serverfarm     : SRV1, type: HOST
 total rservers : 6
 state          : ACTIVE
 DWS state      : ENABLED_REMOTE_LB(Bursting traffic to local and remote VMs)
                                                ----------connections-----------
      real                  weight state        current    total      failures
 ---+--------------------+-----+------------+----------+----------+---------
   rserver: VM1
       10.25.1.11:0          8   OPERATIONAL [L] 33         82         0
   rserver: VM2
       10.25.1.12:0          8   OPERATIONAL [L] 34         81         0
   rserver: VM3
       10.25.2.11:0          8   OPERATIONAL [L] 32         79         0
   rserver: VM4
       10.25.2.12:0          8   OPERATIONAL [L] 31         79         0
   rserver: VM5
       10.25.1.21:0          8   OPERATIONAL [R] 8          8          0
   rserver: VM6
       10.25.1.22:0          8   OPERATIONAL [R] 8          8          0
```

Figure 3 represents the state of the DWS topology at this point, with the additional client traffic coming in.

Figure 3 DWS With Roundrobin Predictor, Post-Burst

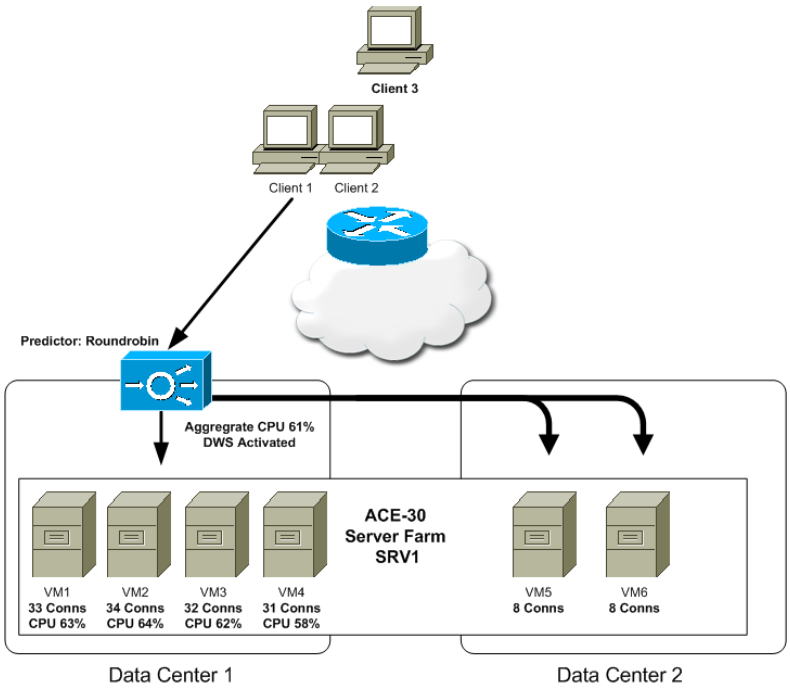

DWS Using Leastconns Predictor

In this use case, the **leastconns** predictor was configured on the ACE-30.

```
dc1c-ace-s9/cap-exp# show serverfarm SRV1 detail

   Codes: L - local,   R - remote

 serverfarm     : SRV1, type: HOST
 total rservers : 6
 state          : ACTIVE
 DWS state      : ENABLED_REMOTE_LB(Bursting traffic to local VMs)
 active rservers: 6
 description    : -
 predictor      : LEASTCONNS
   slowstart    : 0 secs
 failaction     : -
 back-inservice    : 0
 partial-threshold : 0
 num times failover       : 0
 num times back inservice : 0
 total conn-dropcount : 0
<snip>
```

Two clients established connections to four servers in the local data center. All of the two clients' connections were load balanced between the four servers in the local data center. The aggregate CPU utilization was approximately 50%.

```
dc1c-ace-s9/cap-exp# show probe VCENTER-DC1 detail

 probe       : VCENTER-DC1
 type        : VM
 state       : ACTIVE
 description :
----------------------------------------------
   interval          :  10      vm-controller :  VCENTER-DC1
   cpu-load:
      burst-threshold:
       max threshold :  50      min threshold :  25
   mem-load:
      burst-threshold:
       max threshold :  99      min threshold :  99
                ------------------ probe results ------------------
   associations   ip-address     cpu-load  mem-load   health
   ------------ ---------------+---------+---------+---------------
   serverfarm  : SRV1
                aggregate-stats    50       53       BURST_REMOTE

   No. Passed probe    :          89         No. Failed probe  : 0
   No. Probes skipped  :          0          Last status code  : 0
   Last probe time     : Wed Jun  1 16:07:35 2011
   Last fail time      : Never

     real      : VM1[0]
                 10.25.1.11         50       52       SUCCESS

     real      : VM2[0]
                 10.25.1.12         52       54       SUCCESS

     real      : VM3[0]
                 10.25.2.11         51       55       SUCCESS

     real      : VM4[0]
                 10.25.2.12         49       54       SUCCESS
```

A look at the real servers in server farm SRV1 shows that the four local servers are currently handling client connections while the two remote servers are not yet.

```
dc1c-ace-s9/cap-exp# show serverfarm SRV1

   Codes: L - local,   R - remote
 serverfarm     : SRV1, type: HOST
 total rservers : 6
 state          : ACTIVE
 DWS state      : ENABLED_REMOTE_LB(Bursting traffic to local VMs)


----------connections-----------
       real              weight state        current    total      failures
   ---+--------------------+-----+------------+----------+----------+-------
   rserver: VM1
       10.25.1.11:0          8   OPERATIONAL [L] 25         108        0
   rserver: VM2
       10.25.1.12:0          8   OPERATIONAL [L] 24         106        0
   rserver: VM3
       10.25.2.11:0          8   OPERATIONAL [L] 24         105        0
   rserver: VM4
       10.25.2.12:0          8   OPERATIONAL [L] 24         104        0
   rserver: VM5
       10.25.1.21:0          8   OPERATIONAL [R] 0           0         0
   rserver: VM6
       10.25.1.22:0          8   OPERATIONAL [R] 0           0         0
```

Figure 4 represents the state of the DWS topology at this point, prior to additional client traffic coming in.

Figure 4 DWS With Leastconns Predictor, Pre-Burst

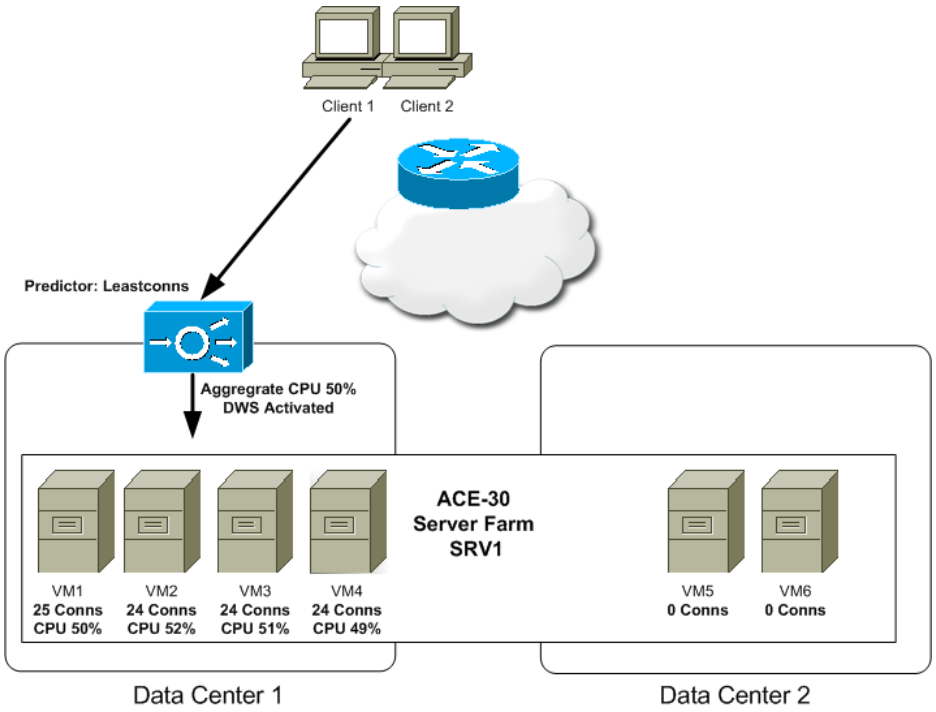

Now, Client 3 begins a new session and creates new connections into the data center and to the server farm. The new connections are load balanced equally to those servers with the least number of existing connections. In the current state, those servers reside in the remote data center. Note that the aggregate CPU utilization of the servers in the local data center remains unchanged, as new connections are load balanced only to those servers in the remote data center.

```
dc1c-ace-s9/cap-exp# show probe VCENTER-DC1 detail

 probe       : VCENTER-DC1
 type        : VM
 state       : ACTIVE
 description :
---------------------------------------------
   interval          :  10      vm-controller :  VCENTER-DC1
   cpu-load:
      burst-threshold:
       max threshold :  50      min threshold :  25
   mem-load:
      burst-threshold:
       max threshold :  99      min threshold :  99
                ------------------ probe results ------------------
   associations   ip-address     cpu-load  mem-load   health
   ------------ ---------------+---------+---------+---------------
   serverfarm  : SRV1
                 aggregate-stats    50        53        BURST_REMOTE

   No. Passed probe    :         142          No. Failed probe   : 0
   No. Probes skipped  :           0          Last status code   : 0
   Last probe time     : Wed Jun  1 16:16:25 2011
```

```
   Last fail time      : Never

     real      : VM1[0]
                 10.25.1.11         51       52       SUCCESS

     real      : VM2[0]
                 10.25.1.12         52       54       SUCCESS

     real      : VM3[0]
                 10.25.2.11         51       55       SUCCESS

     real      : VM4[0]
                 10.25.2.12         50       54       SUCCESS
```

The ACE shows that the number of connections to the local real servers has not increased. Instead, only the remote real servers have received new connections from the ACE, due to the **Leastconns** predictor.

```
dc1c-ace-s9/cap-exp# show serverfarm SRV1

   Codes: L - local,   R - remote

 serverfarm     : SRV1, type: HOST
 total rservers : 6
 state          : ACTIVE
 DWS state      : ENABLED_REMOTE_LB(Bursting traffic to local and remote VMs)

                                              ----------connections-----------
     real              weight state        current    total      failures
   ---+--------------------+-----+------------+----------+----------+-------
   rserver: VM1
       10.25.1.11:0          8   OPERATIONAL [L] 25         109        0
   rserver: VM2
       10.25.1.12:0          8   OPERATIONAL [L] 24         107        0
   rserver: VM3
       10.25.2.11:0          8   OPERATIONAL [L] 24         105        0
   rserver: VM4
       10.25.2.12:0          8   OPERATIONAL [L] 24         105        0
   rserver: VM5
       10.25.1.21:0          8   OPERATIONAL [R] 24         24         0
   rserver: VM6
       10.25.1.22:0          8   OPERATIONAL [R] 23         23         0
```

Figure 5 represents the state of the DWS topology at this point, with the additional client traffic coming in.

Figure 5 DWS With Leastconns Predictor, Post-Burst

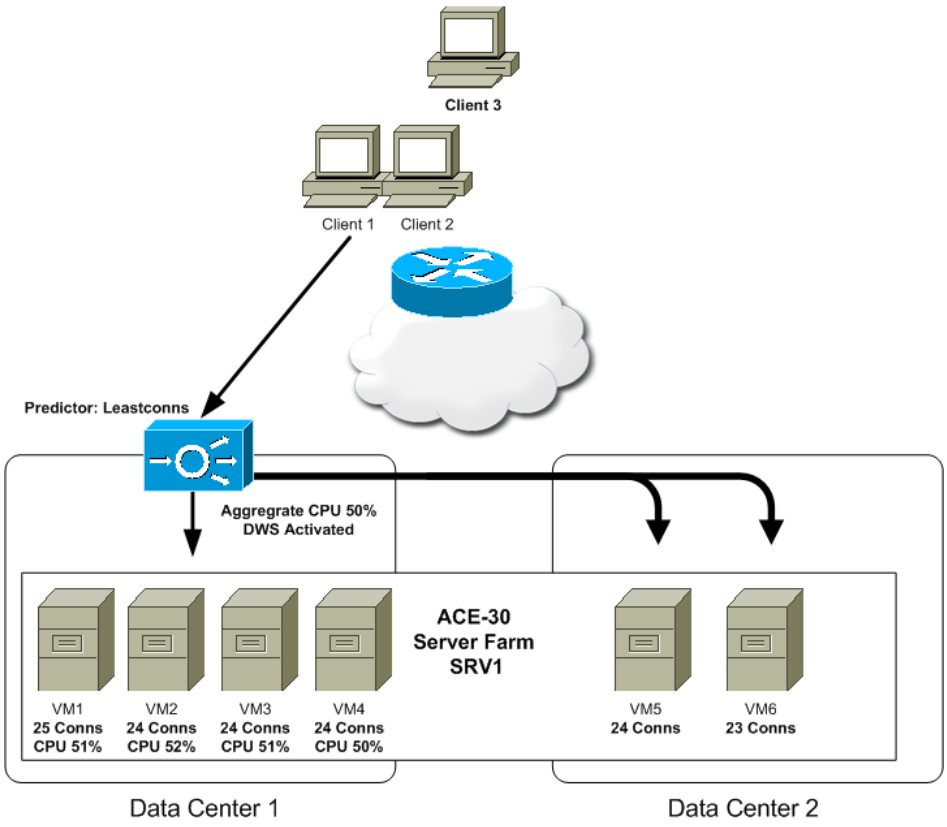

DWS and VMotion

DWS, in conjunction with OTV to extend Layer 2 between data centers, allows the network administrator to monitor server utilization and utilize VMware's VMotion mobility mechanism to move server resources from one data center to another with no disruption to the current network or its existing connections. Once the server has been relocated, the Cisco ACE-30 load balancer dynamically determines the new location of the server, by polling the OTV information from the Nexus 7000 switch. The ACE-30 polls the Nexus 7000 switch every minute to determine the location of the server resources.

In this use case, two clients have established connections to the servers in the local data center. Once the connections are established, a migration of VM4 using VMotion will be done to move VM4 to the remote data center. For this particular use case, the aggregate maximum threshold will be kept below 50%.

Once the migrration of VM4 has been completed, a third client will be started. All of Client 3's connections will be established in the local data center, having no effect on VM4, which had just previously moved to the remote data center. VM4's connections will remain stable and active, while Client 3's connections will be load balanced to the remaining 3 real servers in the local data center.

Next VM4 will be migrated back to the local data center. During both server migration activities es (from and back to the local data center), the Cisco ACE-30 learns the new location of VM4 in less than one minute and all connections to VM4 remain established. VM4 retains its existing connections during both migration events. Again, the following combination of narrative, CLI output and diagrams will lead through this use case.

Figure 6 shows the topology for this use case. Again, there is an extended server farm, with four real servers in DC1 (local) and two in DC2 (remote). In this beginning state, the four local servers have roughly 24 connections each.

Figure 6 DWS With VMotion, #1

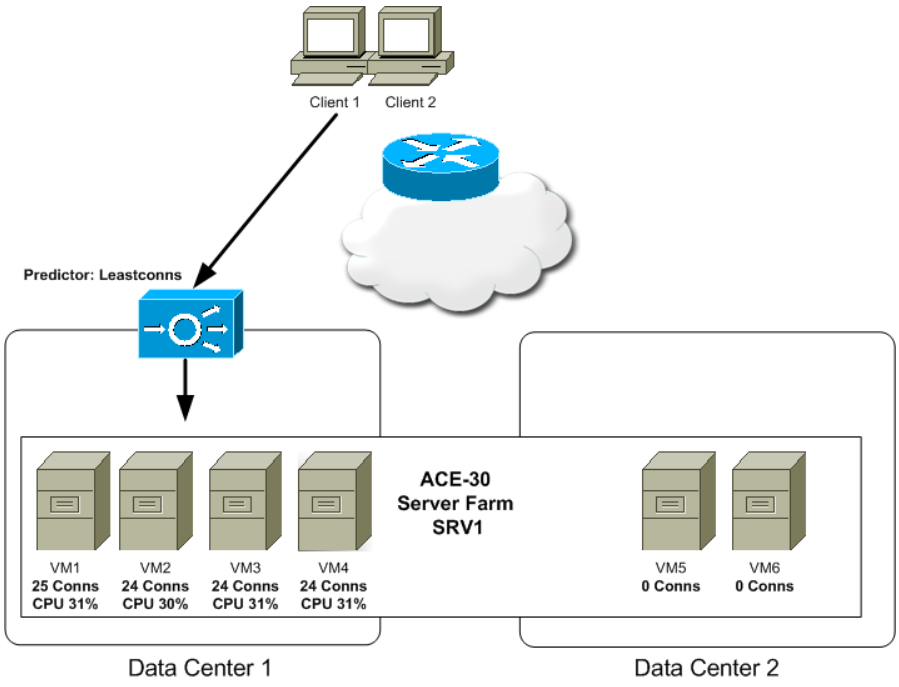

The server farm data from the ACE demonstrates the current connections that are alive on these four servers. Note that real server VM4 is considered local ("[L]") and has 24 current connections.

```
dc1c-ace-s9/cap-exp# show serverfarm SRV1

   Codes: L - local,   R - remote

 serverfarm     : SRV1, type: HOST
 total rservers : 7
 state          : ACTIVE
 DWS state      : ENABLED_REMOTE_LB(Bursting traffic to local VMs)

                                                ----------connections-----------
     real                  weight state         current    total      failures
  ---+--------------------+-----+------------+----------+----------+---------
  rserver: VM1
      10.25.1.11:0          8   OPERATIONAL [L] 25         25         0
  rserver: VM2
      10.25.1.12:0          8   OPERATIONAL [L] 24         24         0
  rserver: VM3
      10.25.2.11:0          8   OPERATIONAL [L] 24         24         0
  rserver: VM4
      10.25.2.12:0          8   OPERATIONAL [L] 24         24         0
```

```
   rserver: VM5
       10.25.1.21:0          8   OPERATIONAL [R] 0          0          0
   rserver: VM6
       10.25.1.22:0          8   OPERATIONAL [R] 0          0          0
```

Looking at the vCenter probe detail on the ACE, notice all four local real servers are being used to calculate the aggregate stats for this server farm.

```
dc1c-ace-s9/cap-exp# show probe VCENTER-DC1 detail

 probe       : VCENTER-DC1
 type        : VM
 state       : ACTIVE
 description :
----------------------------------------------
   interval          :  10      vm-controller :  VCENTER-DC1
   cpu-load:
      burst-threshold:
       max threshold :  50      min threshold :  25
   mem-load:
      burst-threshold:
       max threshold :  99      min threshold :  99
                 ------------------ probe results ------------------
   associations   ip-address     cpu-load  mem-load   health
   ------------ ---------------+---------+---------+-----------------------
   serverfarm  : SRV1
                aggregate-stats    30       39       BURST_LOCAL

   No. Passed probe    :          86         No. Failed probe  : 0
   No. Probes skipped  :           0         Last status code  : 0
   Last probe time     : Wed Jun 15 19:55:11 2011
   Last fail time      : Never

     real      : VM1[0]
                 10.25.1.11         31       41       SUCCESS

     real      : VM2[0]
                 10.25.1.12         30       40       SUCCESS

     real      : VM3[0]
                 10.25.2.11         31       39       SUCCESS

     real      : VM4[0]
                 10.25.2.12         31       36       SUCCESS
```

Figure 7 illustrates the server migration event that moves real server VM4 from DC1 to DC2. In the diagram, notice that VM4 maintains the current connections through the server migration event.

Figure 7 DWS With VMotion, #2

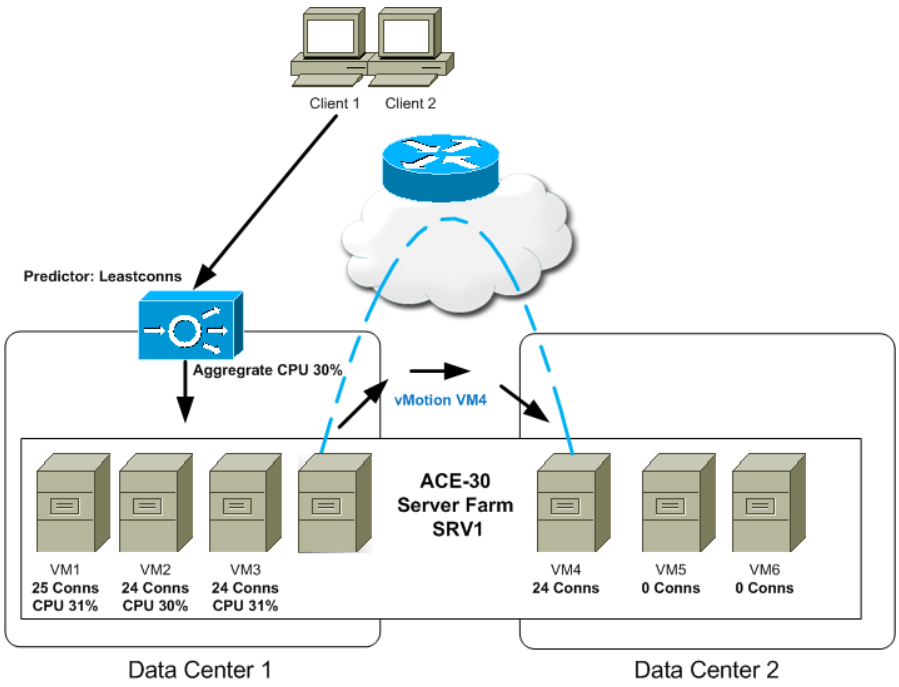

The **show serverfarm** output verifies that not only is VM4 seen as a remote server now ("[R]"), but it has retained its 4 connections.

```
dc1c-ace-s9/cap-exp# show serverfarm SRV1

   Codes: L - local,   R - remote

 serverfarm     : SRV1, type: HOST
 total rservers : 7
 state          : ACTIVE
 DWS state      : ENABLED_REMOTE_LB(Bursting traffic to local VMs)

                                                ----------connections-----------
     real                  weight state         current    total      failures
  ---+--------------------+-----+------------+----------+----------+---------
  rserver: VM1
      10.25.1.11:0          8   OPERATIONAL [L] 25         26         0
  rserver: VM2
      10.25.1.12:0          8   OPERATIONAL [L] 25         25         0
  rserver: VM3
      10.25.2.11:0          8   OPERATIONAL [L] 24         24         0
  rserver: VM4
      10.25.2.12:0          8   OPERATIONAL [R] 24         24         0
  rserver: VM5
      10.25.1.21:0          8   OPERATIONAL [R] 0          0          0
  rserver: VM6
      10.25.1.22:0          8   OPERATIONAL [R] 0          0          0
```

Now that VM4 has been relocated to DC2, the probe detail indicates that only three local real servers are being now used in the aggregate load calculation in this DC1 server farm. The information was updated in less than a minute after the migration was completed.

```
dc1c-ace-s9/cap-exp# show probe VCENTER-DC1 detail

 probe       : VCENTER-DC1
 type        : VM
 state       : ACTIVE
 description :
----------------------------------------------
   interval          :  10      vm-controller :  VCENTER-DC1
   cpu-load:
      burst-threshold:
       max threshold :  50      min threshold :  25
   mem-load:
      burst-threshold:
       max threshold :  99      min threshold :  99
              ------------------ probe results ------------------
   associations   ip-address     cpu-load  mem-load   health
   ------------ ---------------+---------+---------+----------------------
   serverfarm  : SRV1
                 aggregate-stats    30        42        BURST_LOCAL

   No. Passed probe    :          1          No. Failed probe  : 0
   No. Probes skipped  :          0          Last status code  : 0
   Last probe time     : Wed Jun 15 19:56:55 2011
   Last fail time      : Never

     real      : VM1[0]
                 10.25.1.11         31        43        SUCCESS

     real      : VM2[0]
                 10.25.1.12         30        39        SUCCESS

     real      : VM3[0]
                 10.25.2.11         31        44        SUCCESS
```

Now a third client will be started. Client 3's connections are load balanced to the servers in the local data center, due to the DWS threshold not being met or exceeded. This is illustrated in Figure 8 and shown in more detail in the device CLI output below.

Figure 8 DWS With VMotion, #3

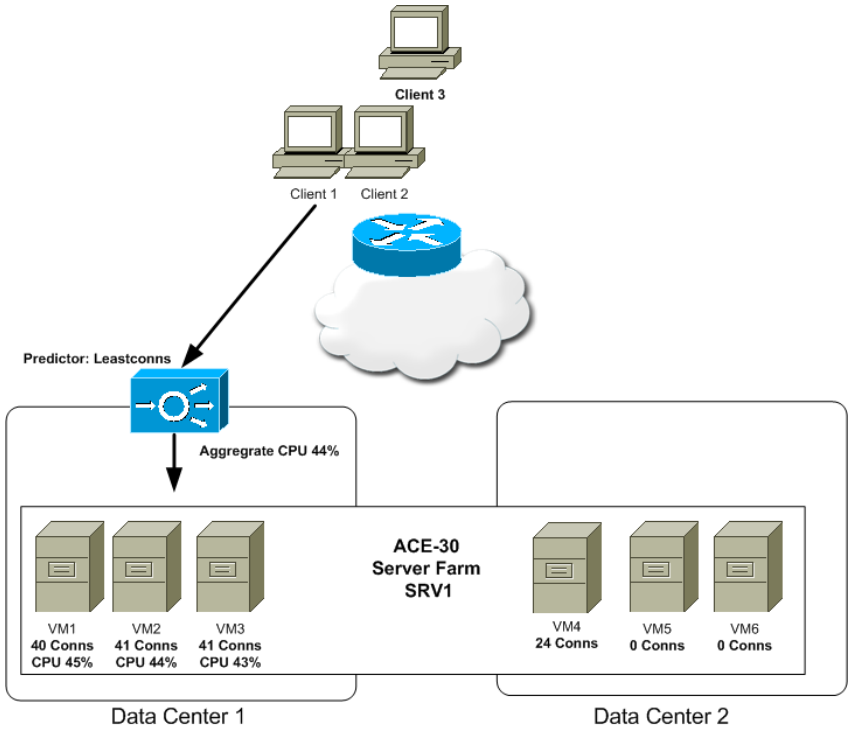


```
dc1c-ace-s9/cap-exp# show serverfarm SRV1

   Codes: L - local,   R - remote

 serverfarm     : SRV1, type: HOST
 total rservers : 7
 state          : ACTIVE
 DWS state      : ENABLED_REMOTE_LB(Bursting traffic to local VMs)
 ---------------------------------
                                                ----------connections-----------
      real                  weight state        current    total      failures
   ---+--------------------+-----+------------+----------+----------+---------
   rserver: VM1
       10.25.1.11:0          8   OPERATIONAL [L] 40         42         0
   rserver: VM2
       10.25.1.12:0          8   OPERATIONAL [L] 41         42         0
   rserver: VM3
       10.25.2.11:0          8   OPERATIONAL [L] 41         41         0
   rserver: VM4
       10.25.2.12:0          8   OPERATIONAL [R] 24         24         0
   rserver: VM5
       10.25.1.21:0          8   OPERATIONAL [R] 0          0          0
   rserver: VM6
       10.25.1.22:0          8   OPERATIONAL [R] 0          0          0
   rserver: VM7
       10.25.4.22:0          8   OUTOFSERVICE[L] 0          0          0

 dc1c-ace-s9/cap-exp# show probe VCENTER-DC1 detail

 probe       : VCENTER-DC1
 type        : VM
 state       : ACTIVE
 description :
```

```
-----------------------------------------------
   interval          :  10      vm-controller :  VCENTER-DC1
   cpu-load:
      burst-threshold:
       max threshold :  50      min threshold :  25
   mem-load:
      burst-threshold:
       max threshold :  99      min threshold :  99
                 ------------------ probe results ------------------
   associations   ip-address     cpu-load  mem-load   health
   ------------ ---------------+---------+---------+----------------------
   serverfarm  : SRV1
                aggregate-stats    44        50       BURST_LOCAL

   No. Passed probe    :          31          No. Failed probe  : 0
   No. Probes skipped  :           0          Last status code  : 0
   Last probe time     : Wed Jun 15 20:01:56 2011
   Last fail time      : Never

     real      : VM1[0]
                 10.25.1.11         45        48       SUCCESS

     real      : VM2[0]
                 10.25.1.12         44        52       SUCCESS

     real      : VM3[0]
                 10.25.2.11         43        50       SUCCESS
```

After Client 3 established it's connections, we see that that the new aggregate threshold for three clients is at 44%, and that the number of active connections being load balanced increased for only those servers in the local data center.

Now, in Figure 9, VM4 is migrated back to the local data center. All 24 connections to VM4 remain stable and connected, and the new aggregate threshold calculation will once again be based upon the four servers in the local data center.

Figure 9 DWS With VMotion, #4

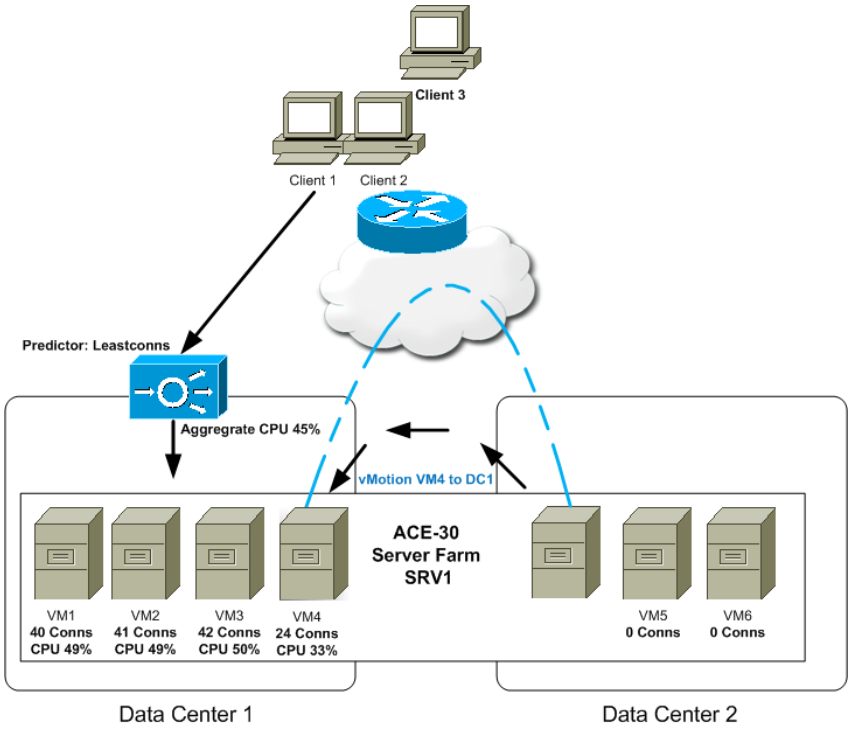


```
dc1c-ace-s9/cap-exp# show serverfarm SRV1

   Codes: L - local,   R - remote

 serverfarm     : SRV1, type: HOST
 total rservers : 7
 state          : ACTIVE
 DWS state      : ENABLED_REMOTE_LB(Bursting traffic to local VMs)
 ---------------------------------
                                                  ----------connections-----------
      real                   weight state         current    total      failures
   ---+--------------------+-----+------------+----------+----------+---------
   rserver: VM1
       10.25.1.11:0          8   OPERATIONAL [L] 40         42         0
   rserver: VM2
       10.25.1.12:0          8   OPERATIONAL [L] 41         42         0
   rserver: VM3
       10.25.2.11:0          8   OPERATIONAL [L] 42         42         0
   rserver: VM4
       10.25.2.12:0          8   OPERATIONAL [L] 24         24         0
   rserver: VM5
       10.25.1.21:0          8   OPERATIONAL [R] 0          0          0
   rserver: VM6
       10.25.1.22:0          8   OPERATIONAL [R] 0          0          0


dc1c-ace-s9/cap-exp# show probe VCENTER-DC1 detail

 probe       : VCENTER-DC1
 type        : VM
 state       : ACTIVE
 description :
```

```
-----------------------------------------------
   interval          :  10      vm-controller :  VCENTER-DC1
   cpu-load:
      burst-threshold:
       max threshold :  50      min threshold :  25
   mem-load:
      burst-threshold:
       max threshold :  99      min threshold :  99
                ------------------ probe results ------------------
   associations   ip-address     cpu-load  mem-load   health
   ------------ ---------------+---------+---------+----------------------
   serverfarm  : SRV1
                aggregate-stats    45        55        BURST_LOCAL

   No. Passed probe    :          35          No. Failed probe  : 0
   No. Probes skipped  :           0          Last status code  : 0
   Last probe time     : Wed Jun 15 20:09:38 2011
   Last fail time      : Never

     real      : VM1[0]
                 10.25.1.11         49        60        SUCCESS

     real      : VM2[0]
                 10.25.1.12         49        63        SUCCESS

     real      : VM3[0]
                 10.25.2.11         50        55        SUCCESS

     real      : VM4[0]
                 10.25.2.12         33        45        SUCCESS
```

DWS in Review

The Cisco Dynamic Workload Scaling solution integrates Cisco load balancing technology with VMware virtualization and Cisco OTV technology. Virtualization technology is gaining momentum in enterprise data centers with enterprises adopting it to optimize the use of computing resources, save costs and gain operational benefits. OTV is critical to the effective deployment of distributed data centers to support application availability and flexible workload mobility with virtualization technology. The ACE-30 technology is in the center of the overall solution and ties all the pieces together.

With the DWS solution, Cisco is again bringing innovation and leading the industry with the introduction of next-generation technology that shapes the data center. OTV technology is the results of years of experience in interconnecting data centers and providing Layer 2 and 3 technologies. DWS is an end-to-end solution to meet the data center challenges and is aligned with the broader set of data center innovations that will be changing data center networking in coming years.

More Information

- Cisco ACE and Cisco ACE Global Site Selector (GSS products)

  http://www.cisco.com/go/ace

- Layer 2 extension between remote data centers:

  http://www.cisco.com/en/US/prod/collateral/switches/ps5718/ps708/white_paper_c11_493718.html

  http://www.cisco.com/en/US/netsol/ns975/index.html

- Cisco OTV Technology:
  http://www.cisco.com/en/US/prod/switches/ps9441/nexus7000_promo.html.
- Cisco Nexus 7000 Series Switches:
  http://www.cisco.com/go/nexus7000
- Cisco Catalyst 6500 Series Switches:
  http://www.cisco.com/go/6500