

IBM WebSphere Commerce V7 Quarterly Releases – 2Q2015

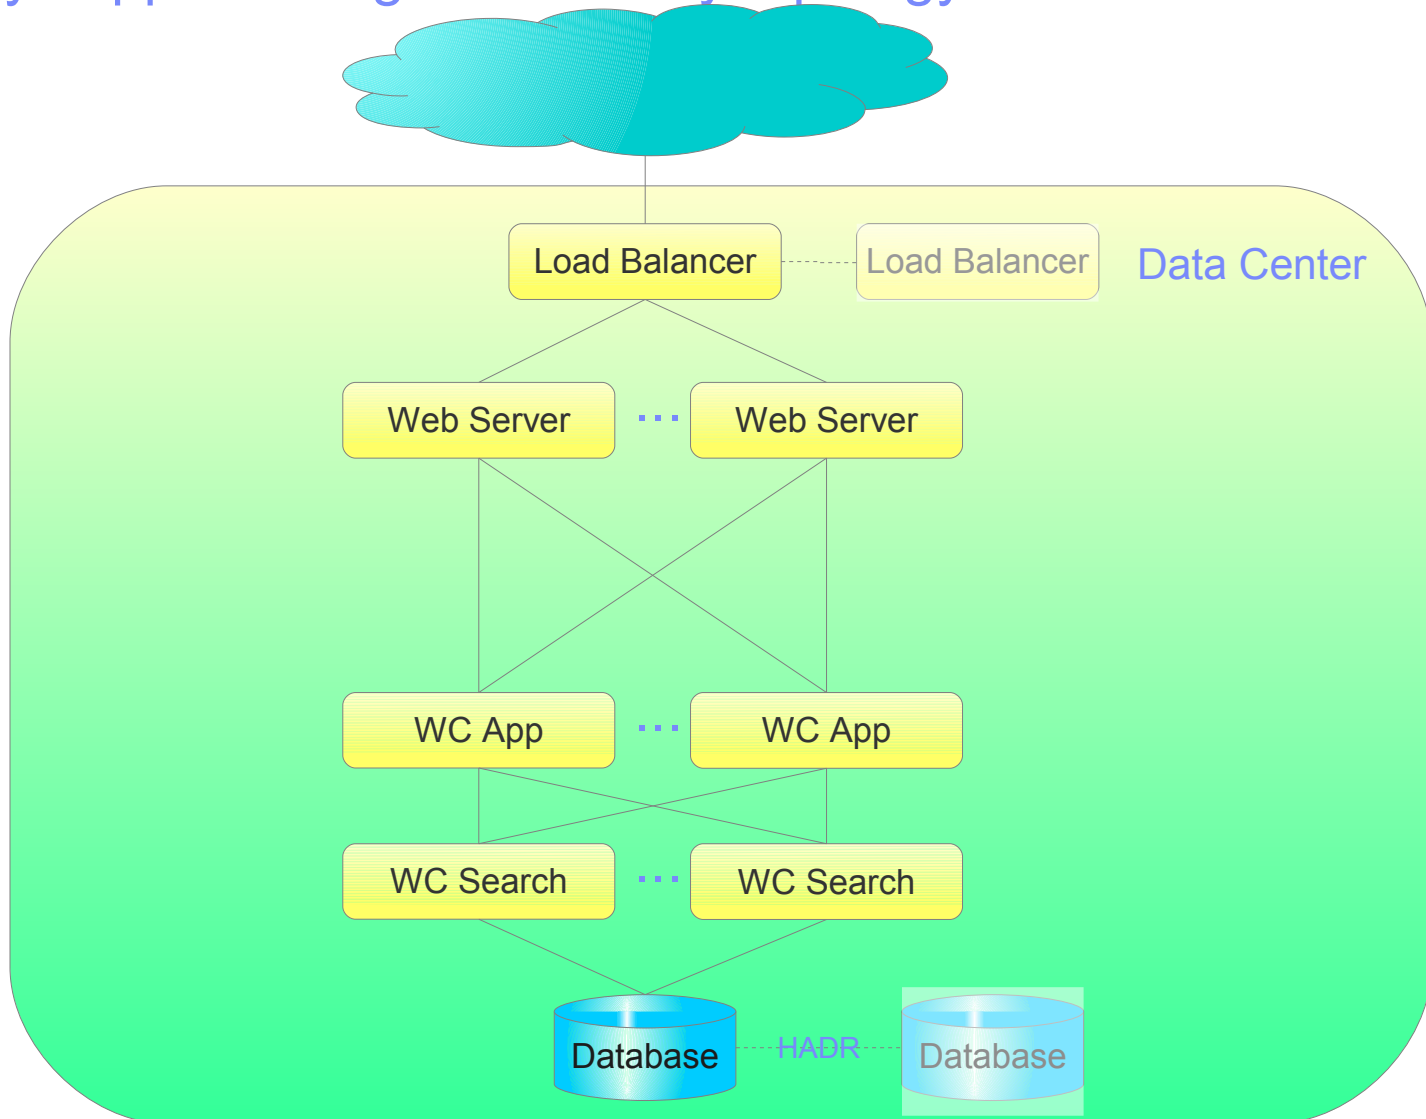
Dual-Cell Topology Support



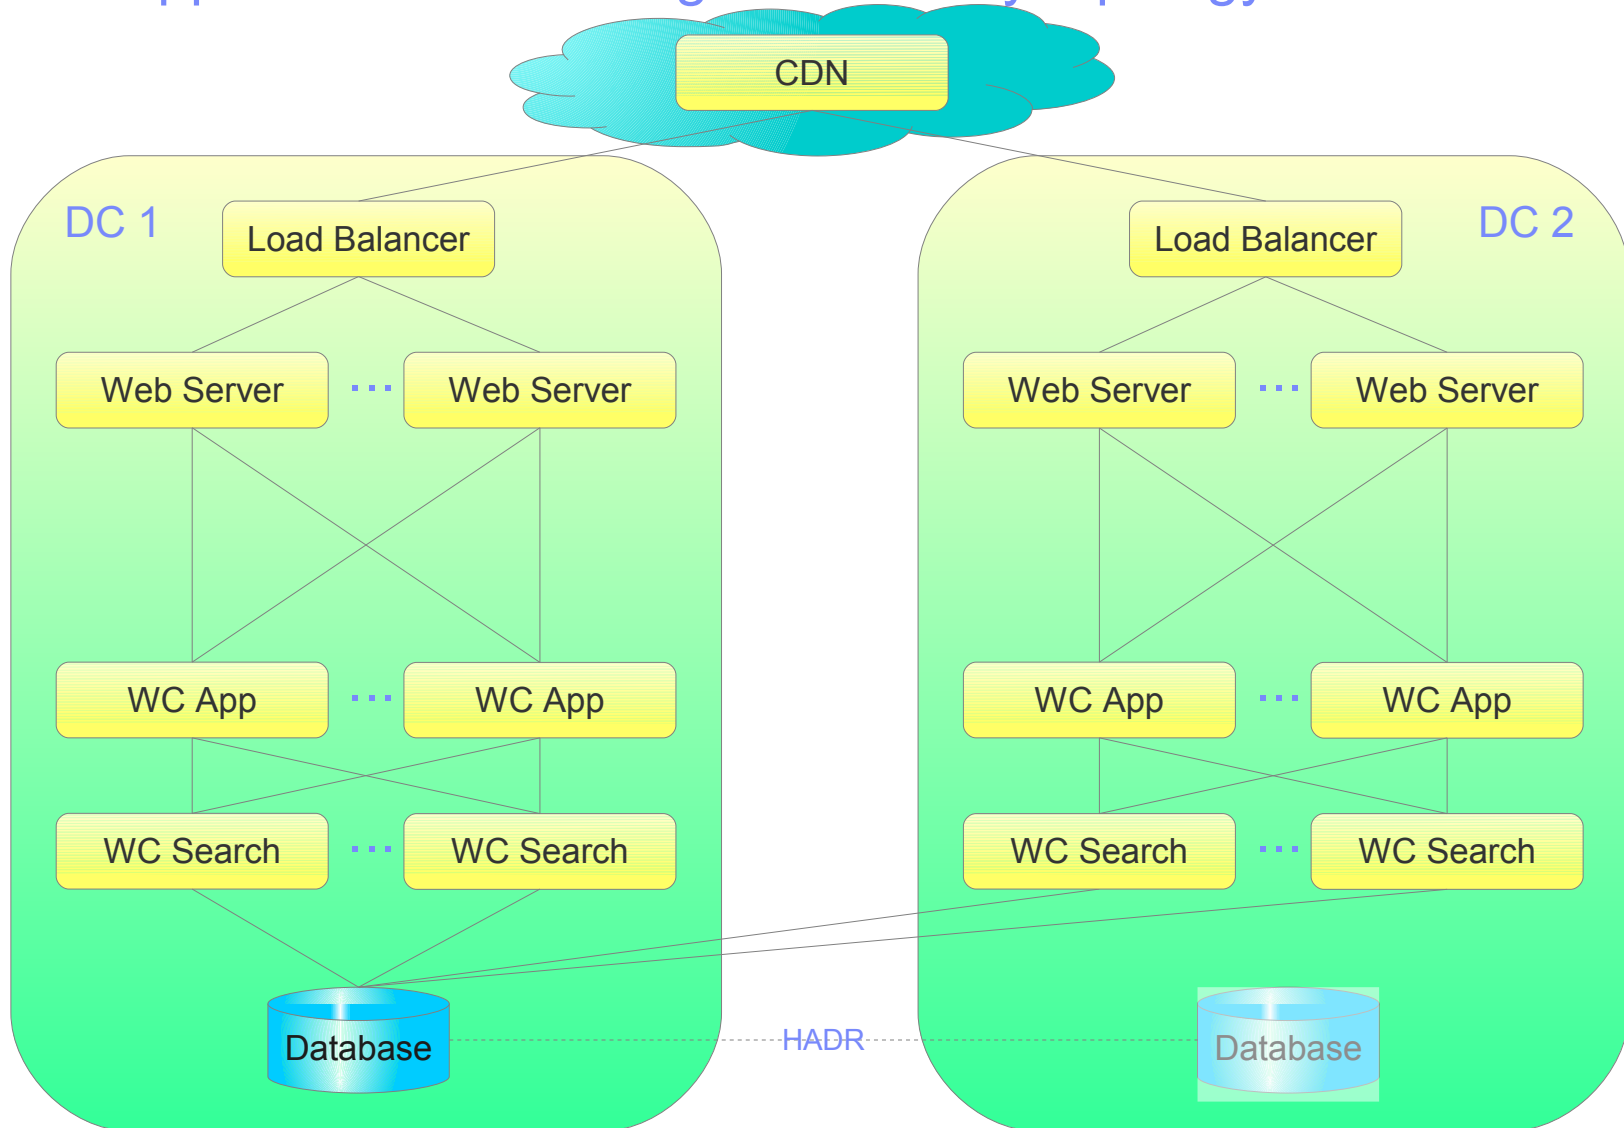
Introducing Dual-Cell Support

- Some scenarios call for deploying WebSphere Commerce in two separate data centers (DCs)
 - Disaster resilience
 - Extra capacity in spare data center
 - Allowing DC-wide maintenance at low traffic time without affecting shoppers
- Traditionally, this topology is not officially supported by IBM
- We now declare support for this topology
 - With some caveats noted in later slides

Currently supported high-availability topology



New supported Dual-Cell high-availability topology



Notes on Dual-Cell Support

- Improved high-availability
 - Resilience to data-center-wide issues, site will continue to work at half capacity even if an entire data center (DC) is lost
- Alternative scenario: DC2 is a stand-by DC
 - Only taking traffic at high-peak times
 - Rest of year, used for other activities, e.g. performance testing
- Support for staged roll-outs of application updates
 - Try new code in DC1 before applying to DC2
- Latency between data centers (DCs) must be low (a few milliseconds)
 - Geographically close
 - Fibre-linked
- Application Server Cell cannot be shared between DCs
 - Each DC has a separate cell
 - Sharing cells across DCs is not recommended

Installation and Maintenance Notes

- As mentioned, we support the configuration, but we are not providing standard installation scripts
- Setting up the second DC does involve a number of tweaks, which are documented in the Knowledge Center
 - Example: A dummy database must be used in DC2 during setup and then reconfigured to point to DC1's database
- Since the application server cell cannot span DCs, maintenance operations must be duplicated between the DCs, for example:
 - Installing maintenance (fix packs and interim fixes) must be coordinated between DCs
 - Deploying custom code updates needs to be done separately in the DCs
 - As such, you must ensure customizations work in a “mixed mode” between the code logic and database

Trademarks, disclaimer, and copyright information

IBM, the IBM logo, ibm.com, Coremetrics, DB2, PowerVM, Rational, WebSphere, and z/VM are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of other IBM trademarks is available on the web at "[Copyright and trademark information](http://www.ibm.com/legal/copytrade.shtml)" at <http://www.ibm.com/legal/copytrade.shtml>

Other company, product, or service names may be trademarks or service marks of others.

THE INFORMATION CONTAINED IN THIS PRESENTATION IS PROVIDED FOR INFORMATIONAL PURPOSES ONLY. WHILE EFFORTS WERE MADE TO VERIFY THE COMPLETENESS AND ACCURACY OF THE INFORMATION CONTAINED IN THIS PRESENTATION, IT IS PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. IN ADDITION, THIS INFORMATION IS BASED ON IBM'S CURRENT PRODUCT PLANS AND STRATEGY, WHICH ARE SUBJECT TO CHANGE BY IBM WITHOUT NOTICE. IBM SHALL NOT BE RESPONSIBLE FOR ANY DAMAGES ARISING OUT OF THE USE OF, OR OTHERWISE RELATED TO, THIS PRESENTATION OR ANY OTHER DOCUMENTATION. NOTHING CONTAINED IN THIS PRESENTATION IS INTENDED TO, NOR SHALL HAVE THE EFFECT OF, CREATING ANY WARRANTIES OR REPRESENTATIONS FROM IBM (OR ITS SUPPLIERS OR LICENSORS), OR ALTERING THE TERMS AND CONDITIONS OF ANY AGREEMENT OR LICENSE GOVERNING THE USE OF IBM PRODUCTS OR SOFTWARE.

© Copyright International Business Machines Corporation 2015. All rights reserved.

IBM WebSphere Commerce V7 Quarterly Releases – 2Q2015

Dual-Cell Topology Support



Welcome to the dual-cell topology support module of the technical enablement for the second quarter 2015 release of IBM WebSphere Commerce version 7.

Introducing Dual-Cell Support

- Some scenarios call for deploying WebSphere Commerce in two separate data centers (DCs)
 - Disaster resilience
 - Extra capacity in spare data center
 - Allowing DC-wide maintenance at low traffic time without affecting shoppers
- Traditionally, this topology is not officially supported by IBM
- We now declare support for this topology
 - With some caveats noted in later slides

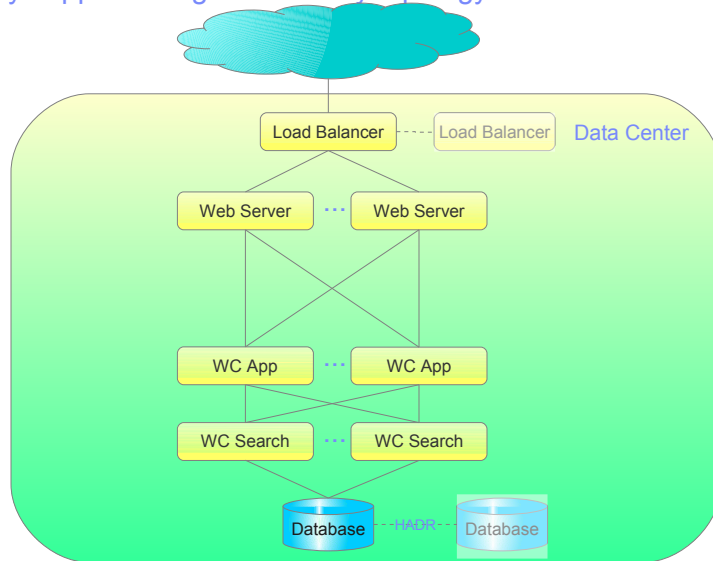
Some clients have requirements that call for a production site split over two separate data centres.

Running in a dual-cell architecture has a number of advantages, including:

- disaster resilience, that is if an entire data centre is taken out of service by, say a power outage, the second data center can take over traffic.
- extra capacity in a spare data center, for example keeping the second data center for load testing outside peak periods, but as additional capacity for live traffic during high-traffic, peak, seasons.
- allowing scheduled maintenance across an entire data center, without impacting end-customers. For example, the main load balancers can be taken offline for firmware upgrades in the primary data center, while traffic is still served from the secondary data center.

We have not traditionally supported this type of configuration, even though we have customers using this topology. However, after extensive testing, we can now declare support for this topology, with some caveats as we will cover later.

Currently supported high-availability topology



3

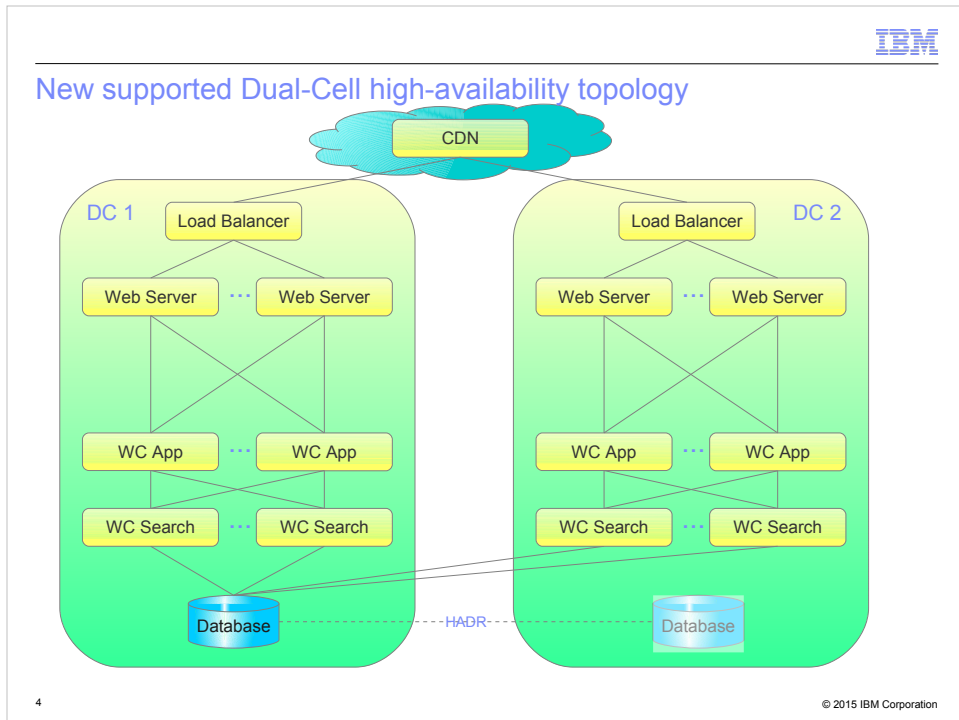
© 2015 IBM Corporation

This slide shows the traditionally supported high-availability topology.

We see a single datacenter with a load balancer, sending traffic to a group of two or more Web servers, passing traffic to two or more application servers, backed by two or more search servers, all reading and updating a single active database.

To provide high availability, most clients have standby load balancers and databases.

New supported Dual-Cell high-availability topology



4

© 2015 IBM Corporation

On this slide we see the new supported dual-cell topology.

We see two data centers, each very similar to the previous slide. However, there are some important points:

- A global load balancer, here shown to be a content delivery network, must provide the traffic balancing between the two data centers
- Each data center is shown here with only one in-center load balancer, for simplicity. However, in most situations, you would still have standby load balancers in both data centers.
- Instead of in-datacenter standby databases, data center 1 has the primary, active, database, and data center 2 has the standby, fail-over, database. In case of an HADR fail-over, data center 2 will have the active database and both data centers will switch over to this database.

Notes on Dual-Cell Support

- Improved high-availability
 - Resilience to data-center-wide issues, site will continue to work at half capacity even if an entire data center (DC) is lost
- Alternative scenario: DC2 is a stand-by DC
 - Only taking traffic at high-peak times
 - Rest of year, used for other activities, e.g. performance testing
- Support for staged roll-outs of application updates
 - Try new code in DC1 before applying to DC2
- Latency between data centers (DCs) must be low (a few milliseconds)
 - Geographically close
 - Fibre-linked
- Application Server Cell cannot be shared between DCs
 - Each DC has a separate cell
 - Sharing cells across DCs is not recommended

Some important points on the dual-cell support:

The support can be used for different scenarios:

- For example, improved high-availability, allowing an entire data center to go offline, while still maintaining functionality. Of course, if capacity is below what a single data center can handle, you must still implement some form of traffic throttling, for example if the outage occurs during peak.

An alternative use case is for only using two data center during high-traffic seasons, while using the stand-by data center for other purposes, for example load testing of new releases, during low-traffic times.

With the two data centers being separated, it is also possible to use the setup for a staged roll-out of new code. Either in an A/B testing scenario, or by taking the secondary data center offline, while testing a deployment.

You may have noticed that the HADR in the database layer is split across the data centers. This means that the bandwidth must be high, and network latency low, between the two data centers. This typically translates into a requirement that the data centers must be geographically close and fibre-linked. As such, this setup will not provide true disaster resilience, at least not against a area-wide incident.

WebSphere Application Server does not support pan-data center application cells. As

Installation and Maintenance Notes

- As mentioned, we support the configuration, but we are not providing standard installation scripts
- Setting up the second DC does involve a number of tweaks, which are documented in the Knowledge Center
 - Example: A dummy database must be used in DC2 during setup and then reconfigured to point to DC1's database
- Since the application server cell cannot span DCs, maintenance operations must be duplicated between the DCs, for example:
 - Installing maintenance (fix packs and interim fixes) must be coordinated between DCs
 - Deploying custom code updates needs to be done separately in the DCs
 - As such, you must ensure customizations work in a “mixed mode” between the code logic and database

Finally, some notes on installation and maintenance.

First, we stress that while we will support a dual-cell topology as described here, we do not provide installation or configuration scripts for setting up such an environment.

The configuration of the second data center does require some tweaks, which are documented in the Knowledge Center. The example here is that while configuring the secondary data center, you must use a dummy database, and once configured, you must switch over to point to the database in the primary data center.

All maintenance operations, such as fixpack installation, interim fix installation, or custom code deployments, but be performed separately in the two data center. A corollary to this is that at some point during a fix install or code deployment, code will be running against a different version than the database. Custom code must be developed to be support such a setup.



Trademarks, disclaimer, and copyright information

IBM, the IBM logo, ibm.com, Coremetrics, DB2, PowerVM, Rational, WebSphere, and z/VM are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of other IBM trademarks is available on the web at "[Copyright and trademark information](http://www.ibm.com/legal/copytrade.shtml)" at <http://www.ibm.com/legal/copytrade.shtml>

Other company, product, or service names may be trademarks or service marks of others.

THE INFORMATION CONTAINED IN THIS PRESENTATION IS PROVIDED FOR INFORMATIONAL PURPOSES ONLY. WHILE EFFORTS WERE MADE TO VERIFY THE COMPLETENESS AND ACCURACY OF THE INFORMATION CONTAINED IN THIS PRESENTATION, IT IS PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. IN ADDITION, THIS INFORMATION IS BASED ON IBM'S CURRENT PRODUCT PLANS AND STRATEGY, WHICH ARE SUBJECT TO CHANGE BY IBM WITHOUT NOTICE. IBM SHALL NOT BE RESPONSIBLE FOR ANY DAMAGES ARISING OUT OF THE USE OF, OR OTHERWISE RELATED TO, THIS PRESENTATION OR ANY OTHER DOCUMENTATION. NOTHING CONTAINED IN THIS PRESENTATION IS INTENDED TO, NOR SHALL HAVE THE EFFECT OF, CREATING ANY WARRANTIES OR REPRESENTATIONS FROM IBM (OR ITS SUPPLIERS OR LICENSORS), OR ALTERING THE TERMS AND CONDITIONS OF ANY AGREEMENT OR LICENSE GOVERNING THE USE OF IBM PRODUCTS OR SOFTWARE.

© Copyright International Business Machines Corporation 2015. All rights reserved.