

# SAP BusinessObjects BI Platform on AWS

## HA/DR Guide for Linux

November 2018

Amazon Web Services (AWS)

### Contents

- About This Guide .....2
- Prerequisite Knowledge .....3
  - AWS Services .....3
  - SAP BusinessObjects BI Platform on AWS .....3
- High Availability .....3
  - HA for SAP BusinessObjects BI Platform on AWS .....4
  - Planning the Deployment in a Primary Region .....5
  - Installing SAP BusinessObjects BI Platform for HA ..... 12
  - HA for SAP Data Services ..... 17
  - HA Testing ..... 18
- Disaster Recovery ..... 18
  - DR for SAP BusinessObjects BI Platform on AWS ..... 18
  - Planning the Deployment in the DR Region ..... 19
  - Installing the SAP BusinessObjects BI Platform in the DR Region .....24
  - DR for SAP Data Services .....26
- Summary .....27
- Appendix A: Resources for Database HA Solutions ..... 28



Appendix B: SAP Notes for HA/DR Configuration.....	30
Appendix C: Additional Tips.....	31
Tagging AWS Resources.....	31
Third-Party Software Components .....	31

## About This Guide

The purpose of this guide is to provide an overview of how to configure high availability (HA) and disaster recovery (DR) for SAP BusinessObjects Business Intelligence (BI) Platform on AWS. This guide will explore how features native to AWS in combination with SAP BusinessObjects BI Platform installation and configuration techniques can greatly improve the availability of an SAP deployment. This guide is not an exhaustive list of all possible configuration options, but covers solutions common to typical deployment scenarios.

This guide isn't intended to replace the SAP BusinessObjects BI Platform installation and administration guides, operating system documentation, or RDBMS documentation.

The procedures and examples in this guide are based on the following:

- A typical, large-scale deployment on AWS that includes two Availability Zones and three subnets in each Availability Zone. You can change this configuration to support your own requirements for SAP BusinessObjects BI Platform servers and tiers.
- An internal Application Load Balancer in front of the web servers, but you can use another internal or internet-facing load balancer.
- Amazon Relational Database Service (Amazon RDS) for MySQL as an example Central Management Server (CMS) and auditing database for SAP BusinessObjects BI Platform. However, you can use any of the [databases supported by SAP](#). HA configuration instructions for other databases aren't included in this guide; see the database-specific documentation on the SAP website. (See [Appendix A](#) for some helpful links.)
- Amazon Elastic File System (Amazon EFS) for input and output filestores.

**Note** The SAP notes referenced in this guide require an SAP ONE Support Launchpad user account. For more information, see the [SAP Support website](#).

## Prerequisite Knowledge

### AWS Services

Before you follow the configuration instructions in this guide, we recommend that you become familiar with the following AWS services. (If you are new to AWS, see [Getting Started with AWS](#).)

- [Amazon Elastic Compute Cloud \(Amazon EC2\)](#)
- [Amazon Elastic Block Store \(Amazon EBS\)](#)
- [Amazon Virtual Private Cloud \(Amazon VPC\)](#)
- [AWS CloudFormation](#)
- [Amazon EFS](#)
- [Amazon RDS](#)

### SAP BusinessObjects BI Platform on AWS

This guide assumes that you're already familiar with implementing and operating SAP solutions on AWS. Please read the SAP notes listed in the following table before continuing. SAP BusinessObjects BI Platform is supported on AWS as described in [SAP Note 2438592](#). All AWS guides for SAP BusinessObjects BI Platform can be found on the [SAP on AWS website](#).

SAP Note	Description
<a href="#">1588667</a>	SAP on AWS: Overview of related SAP notes and web links
<a href="#">1656099</a>	SAP on AWS: Supported products, platforms, and landscapes
<a href="#">2442979</a>	Amazon S3 recommendations for SAP BusinessObjects Business Intelligence Platform
<a href="#">2438592</a>	BI Platform 4.2 Cloud Support

### High Availability

HA design for a software application protects single points of failure (SPOFs). A SPOF is a critical component of an application whose failure can cause service outage for users. The server-side architecture of SAP BusinessObjects BI Platform consists of five tiers: web, management, storage, processing, and data. (For details, see the administrator's guide on the [SAP BusinessObjects Business Intelligence Platform](#) website.) SAP BusinessObjects BI Platform 4.x provides HA of platform services using CMS clustering, which could provide customers with the required level of redundancy. CMS servers are in the management tier.

SAP BusinessObjects BI Platform tiers are designed as follows to eliminate SPOFs and to provide redundant installation options for native, highly available components:

- **Management tier:** Includes the CMS servers, event servers, and associated services. SAP BusinessObjects BI Platform cannot function without a CMS server. Multiple CMS servers can run in a cluster on different machines. A cluster consists of two or more CMS servers working together on a common CMS system database. The CMS is a SPOF, so you must create a cluster of CMS servers running in more than one Availability Zone for a highly available design.
- **Storage tier:** Includes input and output file repository servers. Install these servers redundantly so that failure of any single server doesn't cause a service outage. The file system used by these servers to store files such as documents, reports, and universes must be on a shared file system. This file system is a SPOF and therefore must be highly available.
- **Web tier and processing tier:** These tiers perform functions like receiving and processing user requests. These tiers are not SPOFs. However, if a server that provides a specific service isn't available, users cannot use that service. To avoid such situations, install these servers redundantly so that the failure of any single server doesn't cause a service outage.
- **Data tier:** Consists of the CMS system database and the auditing data store. The CMS database is a SPOF. Install a highly available database using vendor-specific database HA technologies. The specific method used depends on the type of CMS database you're using.

## HA for SAP BusinessObjects BI Platform on AWS

In this guide, we'll provide an example HA architecture that closely resembles a typical on-premises installation, and we'll also show how AWS features in combination with SAP BusinessObjects BI Platform installation options support an HA solution that extends beyond a single data center. For the AWS Cloud, this means that the application is highly available within an AWS Region and can survive the failure of a single Availability Zone. The HA solution discussed in this guide includes these design features:

- SAP BusinessObjects BI Platform nodes are distributed across multiple EC2 instances within a virtual private cloud (VPC). Although Availability Zone failure is a rare occurrence, to protect your application against such scenarios, install each tier of SAP BusinessObjects BI Platform in multiple Availability Zones.
- Applications are deployed on a web server, following best practices from the [Web Application Deployment Guide for Unix](#) on the SAP Help Portal.

- The CMS database is deployed with a failover node that is in a different Availability Zone from the primary node. This guide uses Amazon RDS for MySQL as an example of the SAP BusinessObjects BI Platform CMS database. Amazon RDS is a web service that makes it easier to set up, operate, and scale a relational database in the cloud. It provides cost-efficient, resizable capacity for an industry-standard relational database and manages common database administration tasks. In this guide, we use an Amazon RDS MySQL 5.6 Multi-AZ deployment for the CMS database. See the [AWS documentation](#) for details on how you can configure HA for an Amazon RDS MySQL database in a few steps.
- Input and output filestores are deployed on a shared file system such as an Amazon EFS file system. Amazon EFS provides simple, scalable, elastic file storage for use with AWS Cloud services and on-premises resources.

**Note** Amazon EFS isn't available yet in all AWS Regions; see a [list of available regions](#). If you're planning to use a region that doesn't support Amazon EFS, you can choose a partner network file system (NFS) solution from AWS Marketplace (such as Cloud Volumes ONTAP or SoftNAS) or you can deploy your own solution.

## Planning the Deployment in a Primary Region

Good planning is a key step in ensuring a successful HA deployment for SAP BusinessObjects BI Platform on AWS. Consider these guidelines in your planning:

- Perform the sizing exercise using the [SAP BusinessObjects BI4 Sizing Guide](#). You can determine the resource requirements of each tier based on the sizing exercise.
- Create your architecture document for SAP BusinessObjects BI Platform. Based on sizing decisions, determine the distribution of BI components across EC2 instances and subnets. The level of redundancy in your distributed architecture depends on your HA requirements: your recovery time objective (RTO) and recovery point objective (RPO). For example, if you design your system to be available at full capacity during an Availability Zone failure with zero RTO, you can deploy your system so that servers in a single Availability Zone can process all user requests. If your business can tolerate losing partial capacity temporarily, you can provision lost instances by using Amazon Machine Images (AMIs) and AWS CloudFormation templates in the Availability Zone that's available at the time. There may be other options as well, depending on your cost and recovery time requirements.
- Choose a CMS database with an HA feature. SAP BusinessObjects BI Platform cannot function if the CMS database isn't available. The method of replication between the primary and standby databases depends on your RTO and RPO requirements, and must

be consistent with your application recovery times. If you use an Amazon RDS database as the CMS database, AWS manages HA setup and failover, as explained later in this guide.

- Design the Amazon VPC IP address range, CIDR block, and subnet ranges before you start the installation.

### *Designing Network and Security Groups for the Primary Region*

Define the security groups that will be used to control access to instances for administrative functions, application and DB-level communications, and isolation of failed resources.

**Note** Security groups are firewall rules that you define at the instance or network interface level to open or close specific ports for network communication. You'll need to come up with your own set of rules and configure these based on your application connectivity, setup, and integration requirements.

Here are some of the key considerations for configuring security groups for SAP BusinessObjects BI Platform:

- Users will connect to the web server with web browsers or CMS servers and by using desktop SAP BusinessObjects BI Platform client tools.
- Web servers will communicate with CMS servers and other BusinessObjects BI Platform servers.
- CMS servers will communicate with the CMS and auditing databases.
- The BusinessObjects BI Platform processing tier servers will communicate with the data sources. Data sources could be SAP or non-SAP systems in your landscape where SAP BusinessObjects BI Platform runs the reports.

See [SAP Note 2276646](#) to find out the ports used by different SAP BusinessObjects BI Platform components for communication. The SAP deployment and networking teams should work closely to understand what network traffic to allow in each tier and to configure the tiers accordingly. The following ideas should help provide some structure and guidance:

- Set up a virtual private gateway and one customer gateway. These provide VPN connectivity between the corporate data center and the VPC.
- Set up route table configurations for all the traffic to and from the corporate data center over the VPN tunnel.

- Define all communications on required protocols and ports by using network access control lists (ACLs).
- Set up security groups on management servers with restricted access from certain on-premises networks or IP addresses.
- Set up security groups with limited inbound and outbound protocols and ports for each instance.

**Note** Servers within a particular VPC subnet might need to access resources on the internet for actions such as software updates. You can provide this access by adding an internet gateway to the VPC and using a network address translation (NAT) gateway or NAT instance placed within a public subnet to protect internal resources. Another method is to create network routes to direct the traffic to traverse the VPN tunnel, into the corporate data center, and out through corporate proxy servers. See the blog posts *VPC Subnet Zoning Patterns for SAP on AWS (Part I, II and III)* in the [AWS for SAP](#) blog for guidance on designing VPCs for SAP applications.

This guide uses an example of a typical, large-scale deployment with two Availability Zones to maximize availability and durability, and three subnets per Availability Zone to distribute different SAP BusinessObjects BI Platform tiers. However, based on your sizing requirements, you can also use more than two Availability Zones to install and distribute the SAP BusinessObjects BI Platform nodes. There are many ways to distribute the servers among Availability Zones, EC2 instances, and subnets. In this example, we've designed the architecture for HA/DR, with the specifications listed in the following table. See the [Disaster Recovery](#) section for details on DR design and setup.

- The web tier is installed in subnet 1 in both Availability Zones.
- The management, storage, and processing tiers are installed in subnet 2 in both Availability Zones.
- The data tier (CMS database) is installed in the DB subnet group in subnet 3 in both Availability Zones.

AWS component	Key consideration	Specifications for primary region	Information used for DR region
Region	Consider latency requirements, distance from end users.	us-west-2 US West (Oregon)	us-east-1 US East (N. Virginia)
Availability Zone 1	Choose at least two Availability Zones.	us-west-2a	us-east-1a



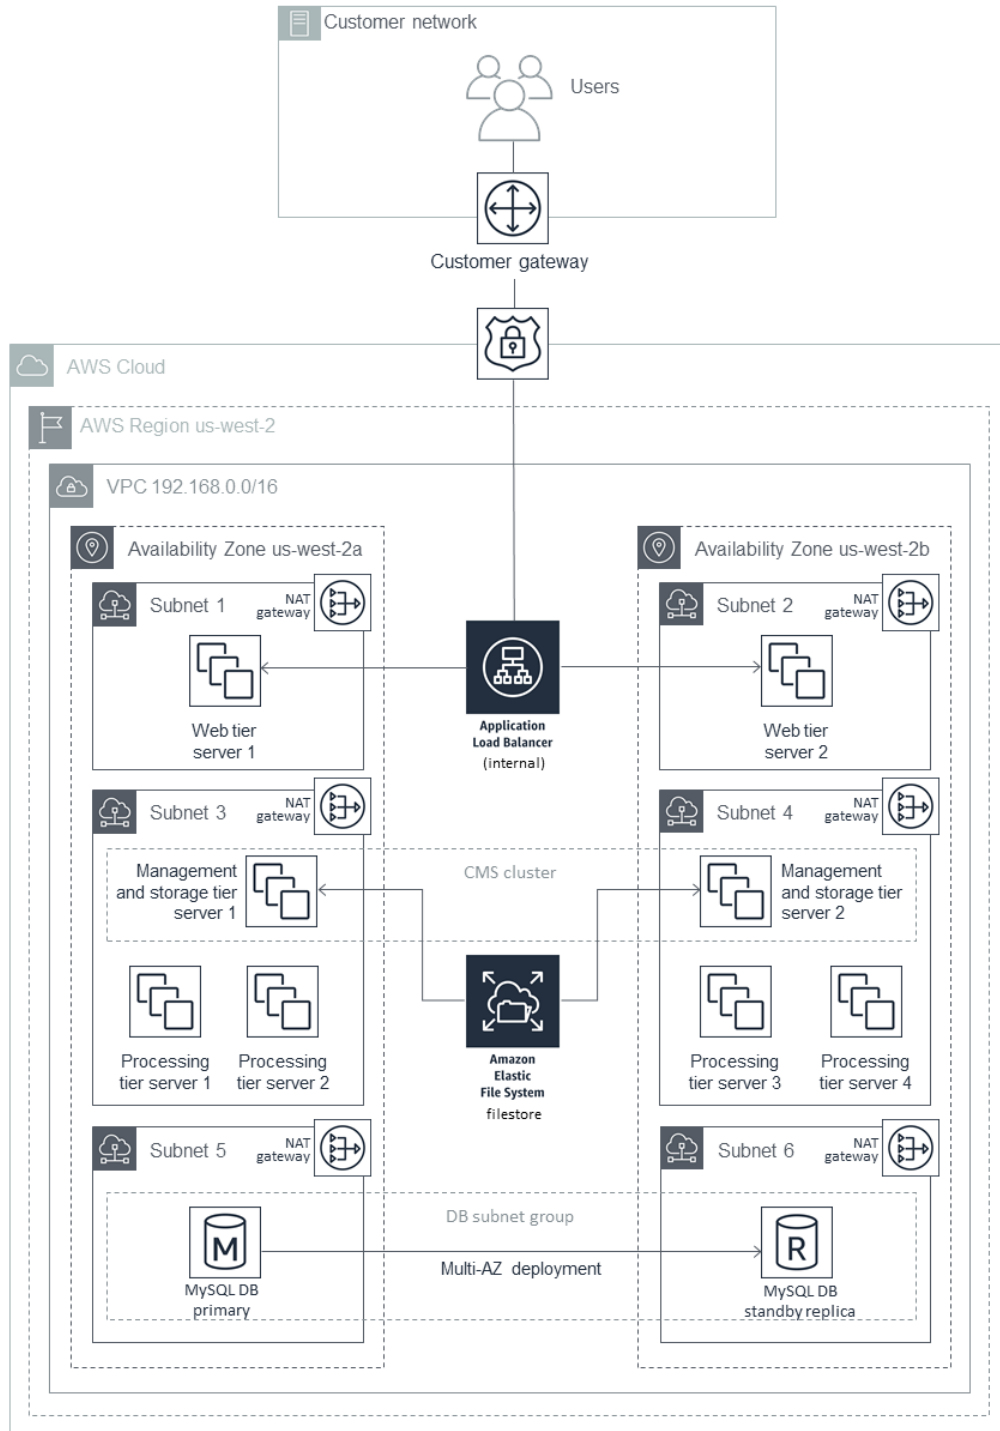
AWS component	Key consideration	Specifications for primary region	Information used for DR region
Availability Zone 2	Choose at least two Availability Zones.	us-west-2b	Availability Zone 2 isn't available for DR in this example.
VPC IP range/CIDR block	Range shouldn't overlap with existing internal IP range. The IP range should be sized appropriately to support the number of hosts and planned growth.	192.168.0.0/16	10.0.0.0/16
Web tier subnet	CIDR IP range for Availability Zone 1 subnets should accommodate future growth.	192.168.5.0/24	10.0.5.0/24
Management, storage, and processing tier subnet		192.168.6.0/24	10.0.6.0/24
DB group subnet		192.168.7.0/24	10.0.7.0/24
Web tier subnet	CIDR IP range for Availability Zone 2 subnets should accommodate future growth.	192.168.8.0/24	Availability Zone 2 isn't available for DR in this example.
Management, storage and processing tier subnet		192.168.9.0/24	
DB group subnet		192.168.10.0/24	

**Note** The information in this table is used as an example to describe the HA/DR architecture. The AWS Regions, Availability Zones, and IP addresses listed in the table will be used throughout this guide as needed. You can change these details for your own implementation, based on your design requirements.

Figure 1 shows the high-level architecture of all application and network components used in this example. This architecture includes an internal Application Load Balancer between the two web servers, but you can use other load balancers in your design. In this case, the load balancer can route requests only from clients that have access to the VPC. Typically, this will allow access from your corporate network and from the VPC itself. If your users access SAP BusinessObjects BI Platform from the internet, you can choose to deploy an internet-facing load balancer. An internet-facing load balancer has a publicly resolvable DNS name, so it can route requests from clients over the internet to the EC2 instances registered with the load balancer. For configuration steps, see [How do I connect a public-facing load balancer to EC2 instances that have private IP addresses?](#) in the AWS Knowledge Center.



**Note** The AWS Regions and Availability Zones shown in the following diagrams are just examples. This architecture can be used in any AWS Region.



**Figure 1: HA architecture for SAP BusinessObjects BI Platform**

This example puts the web tier in a dedicated subnet. End users will access the web tier servers where the web applications are deployed. You can use security groups to restrict connectivity between the subnets only to the required ports. For example, you can restrict connectivity from the application subnet to the database subnet to allow access only to the database listener port. The Multi-AZ architecture provides both load balancing and HA.

The Application Load Balancer will distribute the user load across the two web servers in two Availability Zones. The web servers forward the user requests to the CMS servers. Two EC2 instances (one in each Availability Zone) host the management and storage tier. Similarly, there are four instances (two in each Availability Zone) that host the processing tier. You can scale servers in or out further in any tier, according to your requirements. CMS servers installed in both Availability Zones query the CMS database in Availability Zone 1 during normal operations.

In case of an Availability Zone failure, as shown in Figure 2, the health check for the web server in Availability Zone 1 fails, and the Application Load Balancer forwards requests only to the web server in Availability Zone 2. The management and processing tiers in Availability Zone 2 process the user requests. Amazon RDS MySQL switches the standby replica in Availability Zone 2 to primary, and the CMS server connects to the new primary CMS database in Availability Zone 2. If the servers in Availability Zone 2 alone cannot process the user load, you can choose to redeploy additional servers using AMIs and AWS CloudFormation templates.

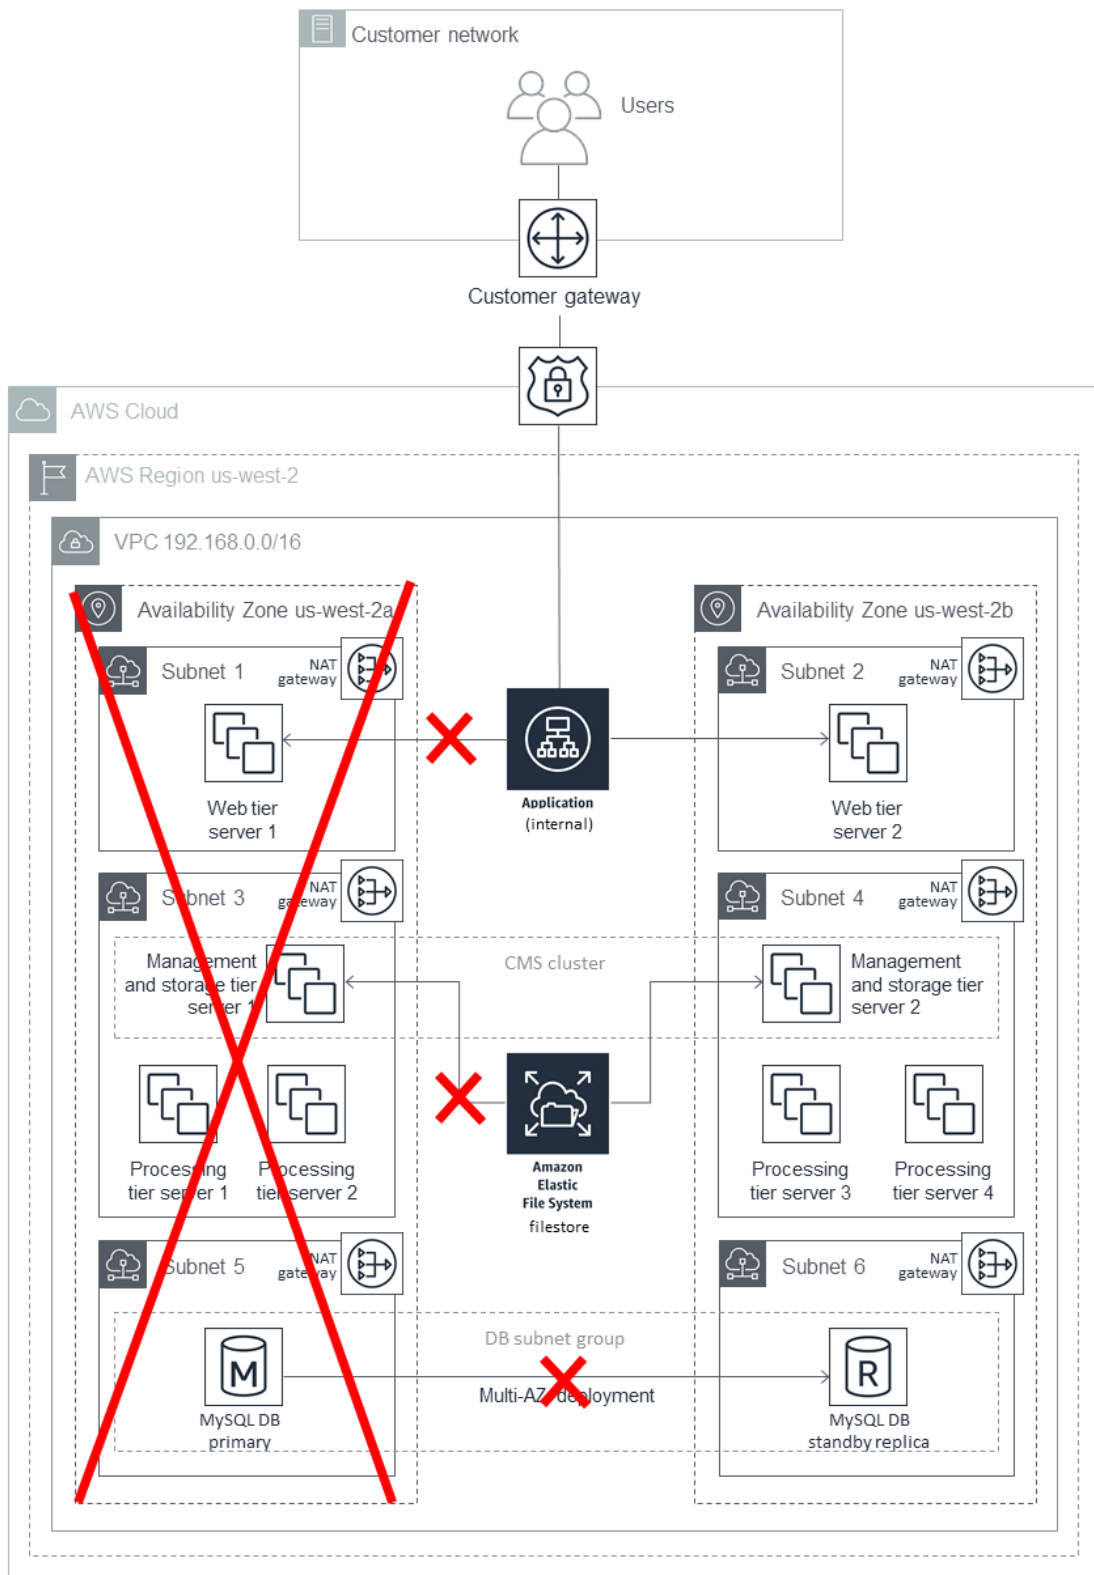


Figure 2: What happens during a failure in Availability Zone 1

## Installing SAP BusinessObjects BI Platform for HA

### Creating a VPC

1. Create a VPC in one of the available AWS Regions by following the instructions in the [Amazon VPC documentation](#).

**Note** You can also use the [Amazon VPC Quick Start](#) to provision your VPC.

2. Specify a contiguous IP address range in CIDR block format (e.g., 192.168.0.0/16). You must choose a range that doesn't overlap with an already existing range used internally on your corporate network.
3. Create six new subnets, associate each with the new VPC, and split them across two different Availability Zones.
4. To continue with the installation process, you'll need access to the instances deployed into the VPC. You can use a [VPN connection](#) or [AWS Direct Connect](#) to establish a dedicated network connection from your premises to AWS. In both cases, be sure to create appropriate route tables, associate them with the new subnets, and adjust the network ACLs with rules to meet internal security requirements.

**Note** The procedure in this section is based on the large-scale distributed deployment example. You can revise the procedure to support your architecture. For example, you can use additional subnets combined with network ACLs to further isolate or restrict access to different tiers of the SAP BusinessObjects BI Platform environment, and you can create multiple public and private subnets for isolation. Any subnet that is associated with an internet gateway will become a public subnet.

### Building the Infrastructure

After you create the VPC, deploy supporting infrastructure instances that will provide key services used by the SAP environment from within the VPC. These services might include:

- Active Directory services (see [Quick Start](#))
- DNS
- Network address translation (NAT) services
- Remote Desktop gateways (see [Quick Start](#))
- Bastion hosts (see [Quick Start](#))

When supporting infrastructure services are in place, you can continue with the next section to start deploying SAP components into the VPC.

### *Preparing and Installing the CMS Database*

This guide uses Amazon RDS MySQL for the CMS database. You can create a separate database for the auditing database if it's required. HA isn't a requirement for the auditing database, so this guide doesn't cover the steps for provisioning it.

1. Create a DB subnet group for an RDS instance by following the instructions in the [Amazon RDS documentation](#). Use two of the subnets you created earlier to create a DB subnet group.
2. In the [Amazon RDS console](#), launch an Amazon RDS MySQL DB instance by following the instructions in the [Amazon RDS documentation](#).
  - a. Choose a supported DB version based on the [Product Availability Matrix](#) on the SAP website, and select the instance type and storage based on your sizing output.
  - b. On the **Specify DB details** page, in the **Instance specifications** section, choose **Create replica in different zone**.
  - c. The **Choose use case** page asks if you are planning to use the DB instance you are creating for production. If you choose **Production - MySQL**, the Multi-AZ failover option is pre-selected.
  - d. On the **Configure advanced settings** page, provide information about the infrastructure you already provisioned, such as settings for the VPC, DB subnet group, and security group. In addition, you can provide custom options for encryption, backup retention period, maintenance window, and so on. You will also create a user to administer this database.
  - e. For the database name, you can provide the name you want to use for the CMS database. You can also change the database port from the default value of 3306 to your choice of port.
  - f. Choose **Create database**, and then wait for the DB instance status to change to **available** in the Amazon RDS console.
  - g. Choose the **Instances** view and note the **Endpoint** name. In case of failover to another Availability Zone, this endpoint enables an application to reconnect to a new primary database instance without having to change anything.
  - h. (Optional) Create a CNAME in Amazon Route 53 for the database cluster endpoint. Use this CNAME during the installation of SAP BusinessObjects BI Platform nodes.

If you're planning a different database than Amazon RDS MySQL, use the database failover solution appropriate for your use case. Some solutions provide automatic failover on AWS

in the event of primary database failure, including the failover of the virtual IP address used for database access. This method provides seamless failover for your applications.

### *Preparing the OS*

Follow these steps to build and configure EC2 instances before installing SAP BusinessObjects BI Platform:

1. Launch a new Amazon EC2 Linux instance in Availability Zone 1. Prepare the server for the installation of SAP BusinessObjects BI Platform by following the instructions in the [SAP BusinessObjects BI Platform installation guide](#).
2. Create an EBS volume and attach it to the EC2 Linux instance. For instructions, see the [Amazon EBS user guide](#).
3. Create and mount a file system by following the instructions in the [Amazon EBS user guide](#). Instructions include enabling automatic mounting each time you reboot the Linux instance.
4. Optionally, you can create an AMI of the instance, as described in the [AWS Toolkit for Visual Studio user guide](#), and launch all other SAP BusinessObjects BI Platform nodes by using this AMI.
5. Launch other EC2 instances according to your design, using one of these methods:
  - Manually perform steps 1-3 for all EC2 instances required by your architecture.
  - Create an AWS CloudFormation template, using the AMI you created in step 4 as a resource. Use this template to launch all additional instances in their respective Availability Zones. This option saves time and helps prevent manual errors during instance provisioning.

### *Installing CMS Nodes*

1. Log in to each CMS server in the SAP BusinessObjects BI Platform node.
2. Launch the installation as described in the [SAP BusinessObjects BI Platform installation guide](#) and choose:  
**Custom / Expand > Expand an existing SAP BusinessObjects BI platform deployment > Instances > Servers > Platform Services**
3. For the first CMS server installation, choose **Start a new SAP BusinessObjects BI platform deployment**.
4. For all additional CMS server installations, choose **Expand an existing SAP BusinessObjects BI platform deployment**.

5. Enter the details of the first CMS server when prompted. This will cluster the additional CMS servers with the first one.

**Note** You can also perform a silent installation by using a response file and the procedure described in the SAP installation guide. This method enables you to quickly rerun the installation again on the same server or on another server, as required. You can further automate the installation by using an AWS CloudFormation template. The template will enable you to provision EC2 instances and install the BusinessObjects BI Platform application in just a few clicks.

### *Installing the Processing Tier*

1. Log in to the processing tier server in the SAP BusinessObjects BI Platform node.
2. Launch the installation as described in the [SAP BusinessObjects BI Platform installation guide](#) and choose:  
**Custom / Expand > Expand an existing SAP BusinessObjects BI platform deployment > Instances > Servers**
3. Choose the servers and other features that you want to install on the processing tier node, and run the installation.
4. Repeat these steps for all your processing tier servers.

### *Installing the Web Tier*

1. Log in to each web tier server in the SAP BusinessObjects BI Platform node.
2. Launch the installation as described in the [SAP BusinessObjects BI Platform installation guide](#) and choose:  
**Custom / Expand > Expand an existing SAP BusinessObjects BI platform deployment > Instances > WebTier**
3. Create a response file for silent installation by choosing **Custom / Expand > Web Tier**. You can install an integrated Tomcat web server by using the SAP installer, or you can use another supported web server and install applications by using the SAP installer.

### *Post-Installation Configuration for HA*

When you have installed the SAP BusinessObjects BI Platform on all servers, you must configure the cluster name, Amazon EFS, and the load balancer to complete the HA configuration.



### Configuring the SAP BusinessObjects BI Platform Cluster Name

Set up a cluster name by following the instructions in the [SAP BusinessObjects BI Platform Administrator Guide](#). Use this cluster name in the properties files of all web applications and client tools.

### Configuring Amazon EFS for Input and Output Filestores

**Note** Amazon EFS isn't available yet in all AWS Regions; see a [list of available regions](#). If you're planning to use a region that doesn't support Amazon EFS, you can choose a partner NFS solution from AWS Marketplace (such as Cloud Volumes ONTAP or SoftNAS) or you can deploy your own solution.

By default, the File Repository Server (FRS) filestore for SAP BusinessObjects BI Platform is located in the local installation directory. You can set up the filestore on a shared storage system like Amazon EFS, and then mount the Amazon EFS file system on all servers that are configured to run storage tier servers:

1. Create an Amazon EFS file system in the VPC where you installed SAP BusinessObjects BI Platform, and create mount targets in the subnets that run storage tier servers.
2. Configure Amazon EFS for automatic mounting on FRS servers by following the instructions in the [Amazon EFS documentation](#), and run the command **mount -a** to mount the file system.
3. Follow the instructions in the [SAP Community Wiki](#) to change the FRS path to Amazon EFS.

### Configuring the Load Balancer for Web Server Access

To distribute the user load evenly across the web tier servers, you can use a load balancer between the web users and the web servers. In this guide, we'll discuss the use of [Elastic Load Balancing \(ELB\)](#) for this purpose. An Application Load Balancer automatically scales its request handling capacity in response to incoming application traffic. Follow these steps to configure an Application Load Balancer for SAP BusinessObjects BI Platform:

1. Enable Secure Sockets Layer (SSL) communications for SAP BusinessObjects BI Platform by following the instructions in the [administrator guide](#). See also: [Enabling SSL in BI Platform 4.2 SP05](#) on the SAP blog.
2. In the [Amazon EC2 console](#), [create an Application Load Balancer](#) in the VPC where SAP BusinessObjects BI Platform is running. Specify the Availability Zones and subnets of all the web tier servers.

3. [Configure a security group](#) that allows users to connect to the Application Load Balancer on the SSL port.
4. [Create a target group](#) to register web servers as the targets to the load balancer. For **Target type**, choose **ip** and specify the IP address and SSL port of the web servers to register as targets.
5. [Enable sticky sessions](#).
6. (Optional) Create a CNAME in Route 53 for the Application Load Balancer DNS name. Use this CNAME to access SAP BusinessObjects BI Platform.

## HA for SAP Data Services

The HA architecture of SAP Data Services is very similar to the HA architecture of SAP BusinessObjects BI Platform. Before you can install SAP Data Services, you must have a working CMS server for SAP BusinessObjects BI Platform. You can install SAP Data Services on existing SAP BusinessObjects BI Platform nodes. If you do not have a SAP BusinessObjects BI Platform installation, SAP BusinessObjects Information Platform Services (IPS) can provide the basic CMS functions required by SAP Data Services. For either option, you can follow the steps in the previous sections to install the CMS server (by clustering two IPS CMS servers together), the CMS database, and a filestore with HA. You must use two Availability Zones for the CMS servers, a highly available shared file system like Amazon EFS, and a CMS database with a Read Replica in a separate Availability Zone. Here are the high-level steps:

1. Create databases for the SAP Data Services repository in Amazon RDS MySQL. You can use your existing SAP BusinessObjects BI Platform or SAP BusinessObjects IPS CMS database engine to create a separate database for the repository. Alternatively, you can create a separate Amazon RDS MySQL instance for the repository.
2. Install SAP BusinessObjects BI Platform or SAP BusinessObjects IPS with a minimum of two CMS servers in two separate Availability Zones. See the [SAP Data Services installation guide](#) for IPS installation instructions.
3. Configure the filestore shared file system as described in the [previous section](#).
4. Install SAP Data Services on all SAP BusinessObjects BI Platform servers or IPS servers.
5. Create an SAP Data Services central repository in the database created in step 1 by following the steps in the section *Repository management* in the [SAP Data Services administrator guide](#).
6. Configure HA for SAP Data Services batch jobs by following the instructions in [SAP Note 1938068](#).

## HA Testing

[SAP Note 1229417](#) provides four tests you must perform to ensure that the CMSs are clustered correctly and that cluster functionality is working. Additionally, when performing the tests, you must also ensure that your deployment can tolerate Availability Zone failure. You can confirm this by stopping all the servers in one Availability Zone and making sure that all the services are still available. To test the CMS database, you can perform the failover manually. In the [Amazon RDS console](#), choose **Instances**, **Instance actions**, **Failover**. After all servers in one Availability Zone are shut down and database failover is complete, test the availability of the system and all the services. Repeat the same test for the second Availability Zone.

## Disaster Recovery

The DR approach you take for SAP BusinessObjects BI Platform, as for any other enterprise application, depends on your RTO and RPO requirements. As discussed in [SAP Note 2056228](#), there are two options for building a DR site for SAP BusinessObjects BI Platform:

- Fully or selectively using SAP Lifecycle Manager (LCM) or Data Federation Services to promote or distribute the content from the primary system.
- Periodically copying over the CMS database and FRS contents, and using that to start a secondary system when required.

## DR for SAP BusinessObjects BI Platform on AWS

DR for SAP BusinessObjects BI Platform on the AWS Cloud refers to a scenario in which the primary AWS Region where SAP BusinessObjects BI Platform application is running is unavailable. The goal of setting up a DR site is to be able to recover the application within your acceptable RTO and RPO.

There are no restrictions in AWS for using LCM or Data Federation Services for your DR environment. Note that using LCM requires either consuming resources on the source system or provisioning an additional system to run promotion management jobs. This option might also result in a higher RPO depending on the frequency of promotion management jobs. See [Promotion Management Architecture](#) on the SAP Community Wiki for the high-level architecture for this option.

In this guide, we'll discuss the second option for handling DR, which is to copy the CMS database and FRS contents. Using variants of this option, you can build the complete primary system within your recovery time limits. This option doesn't require resources from the primary system except for the backup copy of the database and file system.

## Planning the Deployment in the DR Region

Planning the deployment of a DR system is critical. Consider these items in your planning:

1. Choose a region for your DR site. Make sure that this DR region supports the services and resources that you are using for your SAP BusinessObjects BI Platform. Each AWS Region is completely isolated from other AWS Regions to provide the best possible fault tolerance and stability. You may choose a DR region that is geographically closer to your users. For the examples in this guide, we'll assume US West (Oregon) as the primary region and US East (N. Virginia) as the DR region.
2. Determine the RTO and RPO that are acceptable to your business. This decision determines the methodologies you choose for copying the CMS database and FRS content from the primary site to the DR site. See [Using Amazon Web Services for Disaster Recovery](#) on the AWS website for options based on your RTO and RPO requirements. In this guide, because we're using Amazon RDS MySQL for the CMS database, we have two options for copying the database:
  - Perform periodic snapshots of the Amazon RDS MySQL database and copy these snapshots to the DR region (results in lower RTO and RPO, and lower cost).
  - Create an Amazon RDS MySQL Read Replica in the DR region, and promote this Read Replica to master in the event of a disaster or DR drill (results in higher RTO and RPO, and higher cost).

For details, see the section [Installing and Configuring the CMS Database for the DR Region](#) later in this guide.

3. Determine the capacity of the DR site. You might want the DR site to have the same capacity and resources as the primary site, or you might decide to run it with fewer resources.
4. For the examples in this guide, we're using a Route 53 CNAME with an Application Load Balancer for user connectivity. You can change the value of the CNAME to point to the load balancer of the DR region in the event of a disaster or DR drill. This will direct users to the DR environment. You can also use other methods to redirect users if you aren't using Route 53.
5. Determine the failover process. Define the disaster situations beforehand and put a mechanism in place to direct users to the DR site if a disaster occurs. Define test cases to simulate these scenarios, and perform a DR drill.

### *Designing Network and Security Groups for the DR Region*

Network and security groups in the DR region will follow the same design as in the primary region, as discussed in the section [Designing Network and Security Groups for the Primary Region](#) earlier in this guide. The [table in that section](#) lists the VPC and subnet CIDRs that the DR region uses for the examples in this guide. Figure 1 shows the DR architecture used in this guide. This guide does not discuss the use of a secondary Availability Zone for HA in the DR region, but based on your requirements, you can use one to install additional servers.

### *Multiple Region Multi-VPC Connectivity*

See [Multiple Region Multi-VPC Connectivity](#) on the AWS website to find options for connecting multiple VPCs in different AWS Regions. If you aren't using an Amazon RDS database, you can use the database replication method supported by your database vendor to build a DR database and continuously copy the data. SAP HANA System Replication is one such example. AWS provides options to establish a connection between the primary and DR VPCs in two separate AWS Regions securely and reliably. [Multiple Region Multi-VPC Connectivity](#) discusses two ways to connect the VPCs:

- **Routing over AWS networks:** This option enables you to use an AWS backbone network for connectivity. AWS provides a high bandwidth, global network infrastructure to support your regional networking needs. AWS Regions are connected to multiple Internet Service Providers (ISPs) and to a private global network backbone, which provides lower cost and more consistent cross-region network latency when compared with the public internet.
- **Routing over non-AWS networks:** This option enables you to route DR network traffic through your corporate data center. Customers can use an internet-based VPN or their own corporate network backbone.

Each option provides different benefits for bandwidth reliability, security, HA, and use of existing infrastructure, and might have limitations as well. You can choose the option that meets your requirements. You can also use AWS Direct Connect gateways to connect your AWS Direct Connect connection over a private virtual interface to one or more VPCs located in the same or different regions in your account. See [AWS Direct Connect Gateways](#) in the AWS documentation for the rules and limitations of using this approach.

**Note** The AWS Regions and Availability Zones shown in the following diagrams are just examples. This architecture can be used in any AWS Region.

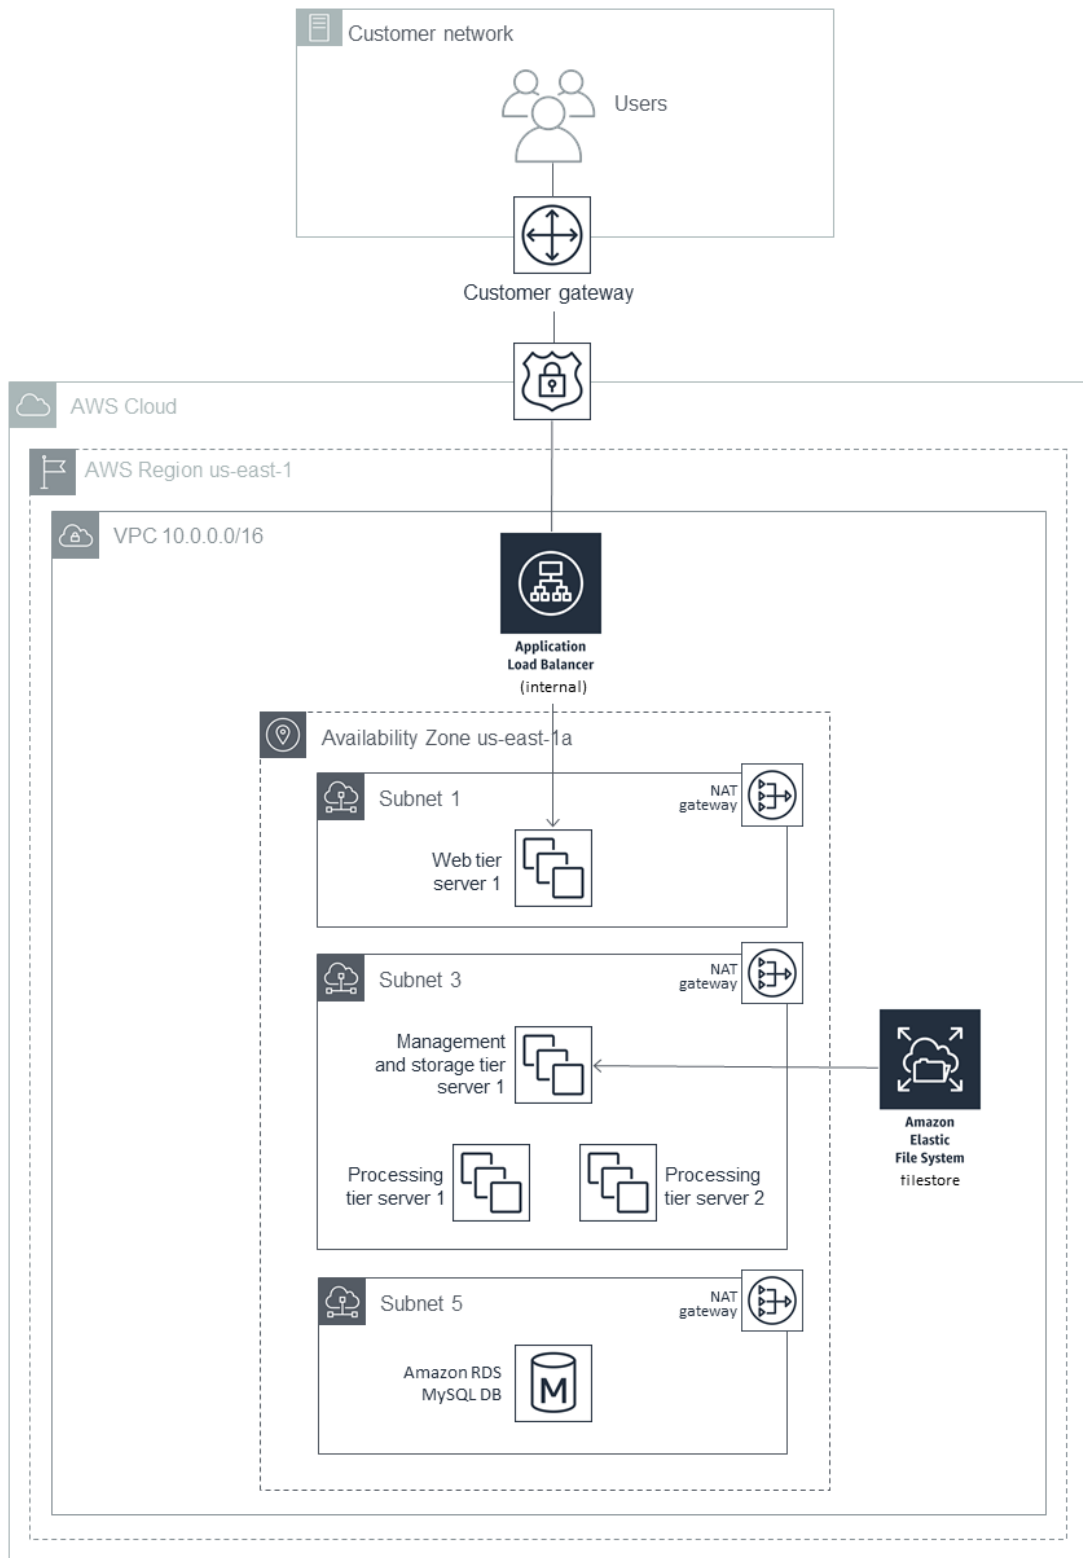


Figure 1: DR architecture for SAP BusinessObjects BI Platform

Figure 4 shows the user connectivity using Route 53 as an example. You can also use your own DNS. During normal operations, the value of the CNAME record for SAP BusinessObjects BI Platform points to the Application Load Balancer of the primary region. The load balancer distributes the load between the available web servers that deploy the SAP BusinessObjects BI Platform web applications.

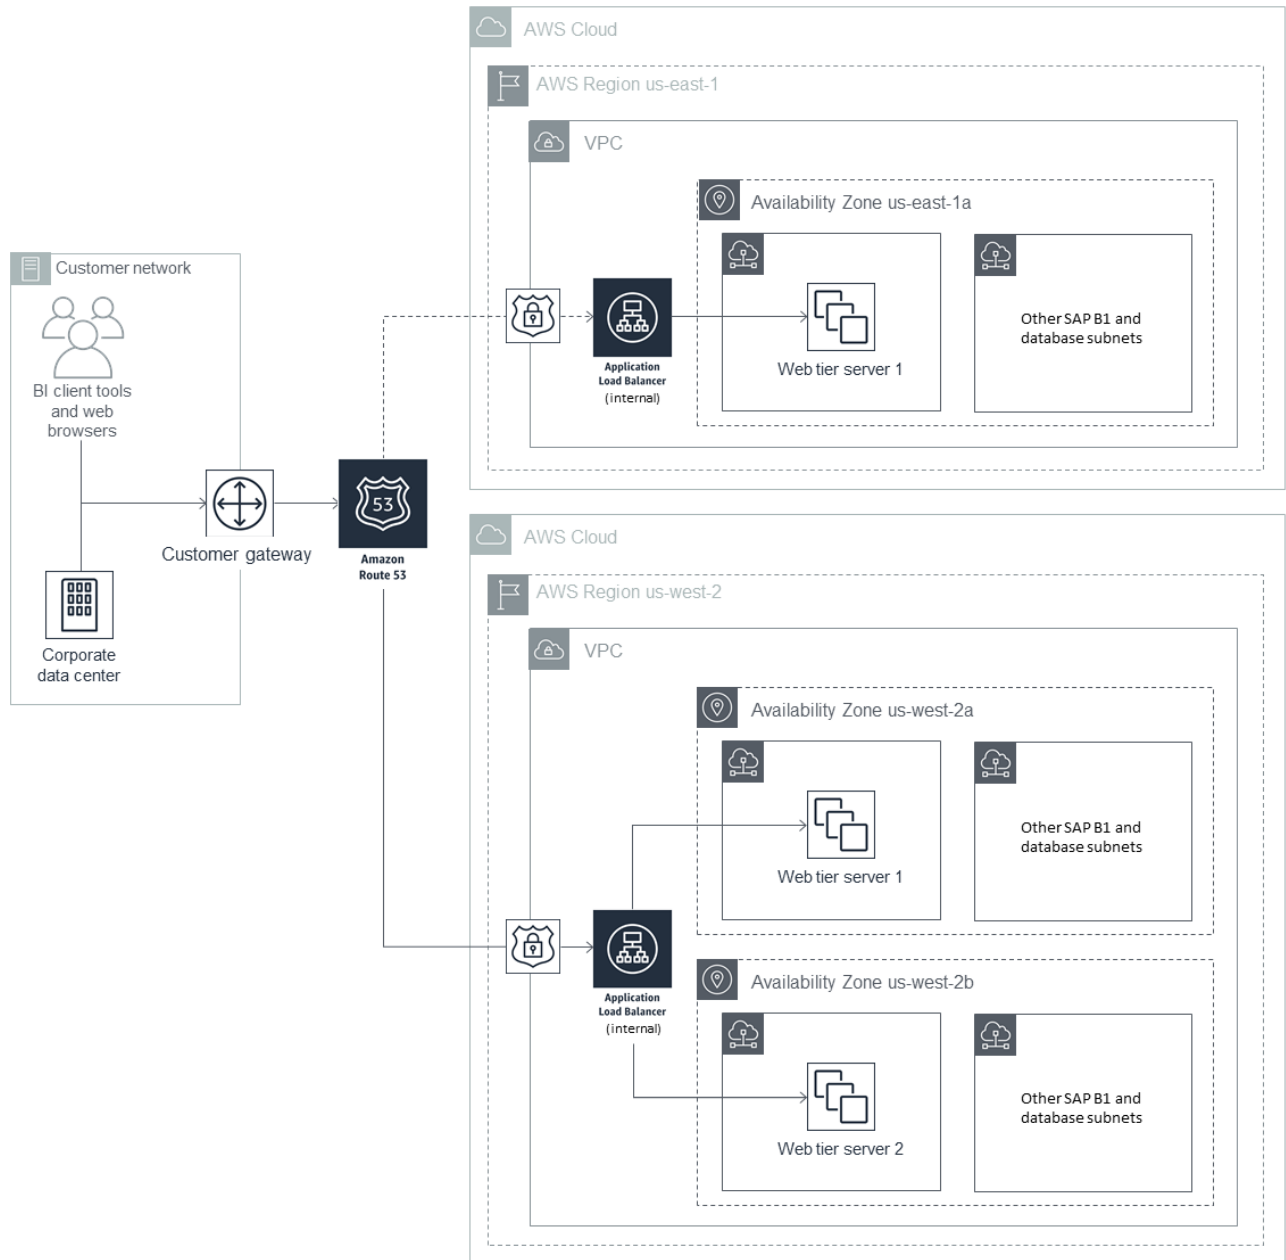
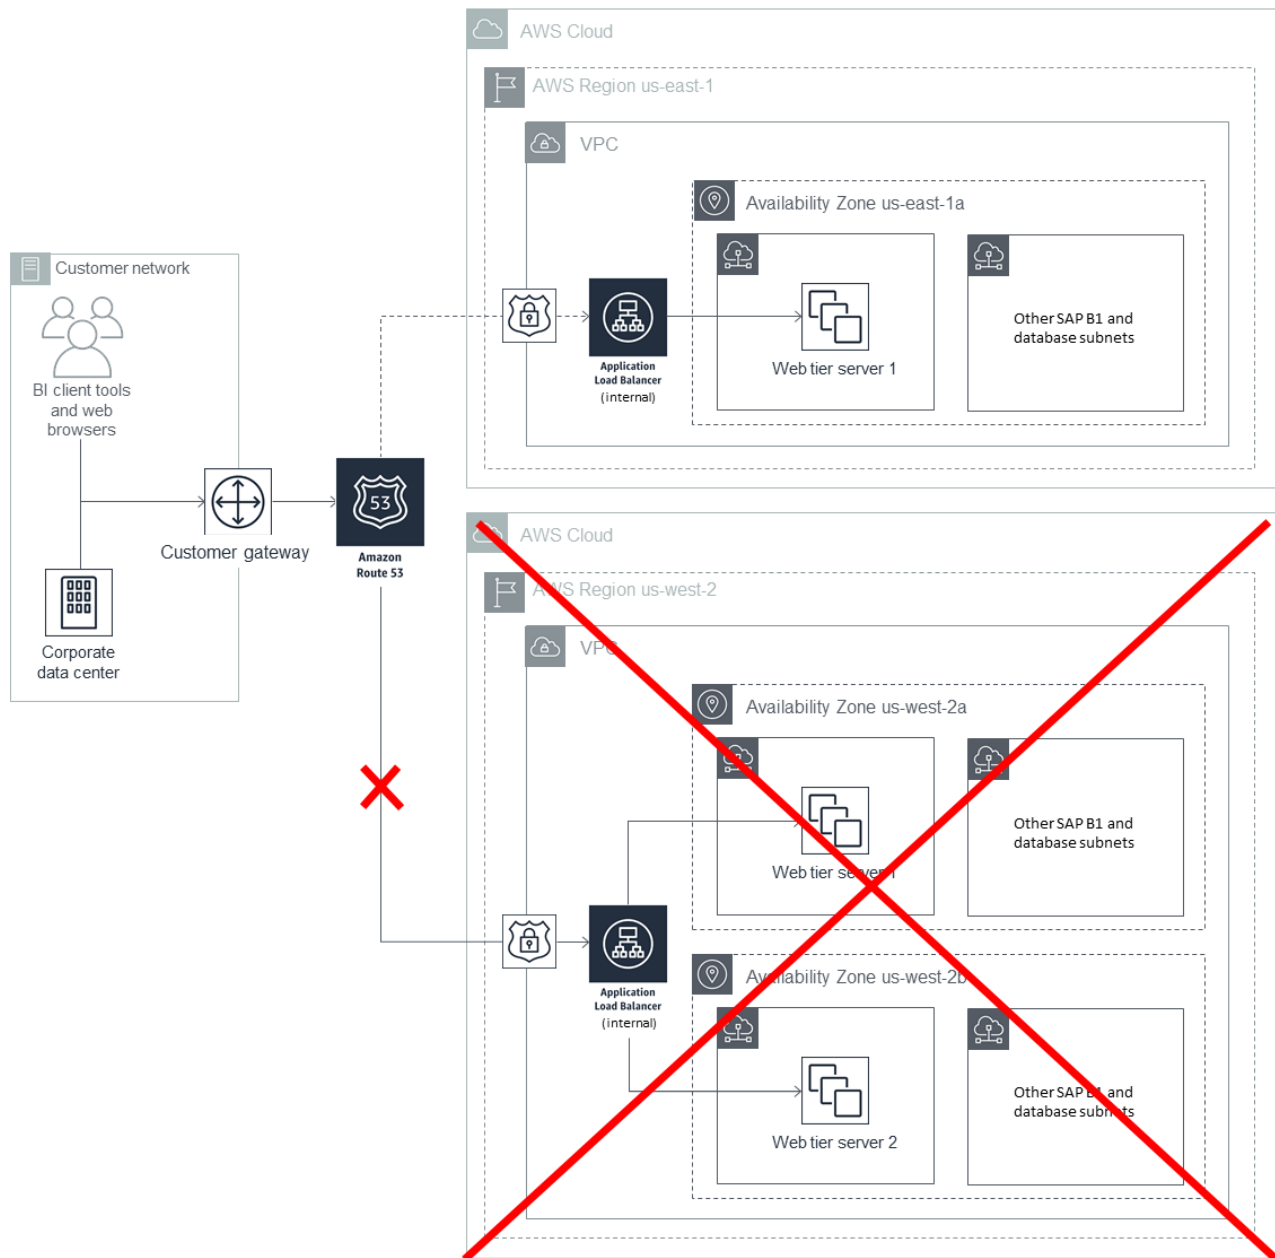


Figure 2: User connectivity with Route 53



In the event of DR or when performing a DR drill, you can either manually switch the DNS to point to the DR region or you can use automatic failover by using Route 53 health checks. For details, see [Configuring DNS Failover](#) in the Route 53 documentation. Redirect all the users to the DR region as shown in Figure 3.



**Figure 3: Redirecting users to a DR region**

## Installing the SAP BusinessObjects BI Platform in the DR Region

Once you have selected the DR region, follow the steps in the [Preparing the OS](#) section earlier in this guide to build SAP BusinessObjects BI Platform EC2 instances in that region.

### *Installing and Configuring the CMS Database for the DR Region*

The CMS database in the DR region must be a copy of the CMS database in the primary region. The method for copying this database depends on the database type and your RTO and RPO requirements. (See [Appendix A](#) for links to resources for commonly used databases.) This guide describes two options for an Amazon RDS MySQL database:

- Option 1: Create a cross-region Read Replica of the primary Amazon RDS MySQL database.

Amazon RDS uses the MySQL DB engine's built-in replication functionality to create a special type of DB instance called a Read Replica from a source DB instance. Updates made to the source DB instance are asynchronously copied to the Read Replica. Create a Read Replica of the primary CMS database by following the instructions in the [Amazon RDS documentation](#). The documentation also provides instructions for promoting a Read Replica to a primary DB instance. This option will provide a lower RPO and RTO than Option 2.

- Option 2: Copy primary database snapshots to the DR region and use them to build a DR CMS database.

Create a snapshot of the Amazon RDS MySQL DB and copy the snapshot to the DR region. Use this snapshot to create a new instance of the Amazon RDS MySQL DB, and use this instance to configure the SAP BusinessObjects BI Platform DR database. See the [Amazon RDS documentation](#) for instructions, features, and limitations of this option. For a continuous copy of the database content, you can create and schedule scripts to perform the snapshot and copy them to the DR region. You can also use open-source tools to deploy such solutions, but make sure to test any tools thoroughly before using them for your DR deployment. When your testing is complete and the primary region is active, you can delete the RDS instance. Create a new instance by using the latest snapshot when the DR system is required again.

### *Installing a CMS Node*

The steps to install the SAP BusinessObjects BI Platform application in the DR region are described in [SAP Note 2603573](#). This procedure uses SQL Anywhere to temporarily create Server Intelligence Agents (SIAs) on all instances with a local CMS database. In the next section, we'll reconfigure all nodes to form a single SAP BusinessObjects BI Platform cluster.

1. Log in to all CMS servers for your SAP BusinessObjects BI Platform DR nodes, based on your design.
2. Install the CMS servers by following the instructions in the SAP Note and using the same SIA names as in the production cluster.

### *Starting the SAP BusinessObjects BI Platform in the DR Region*

Create the SIA nodes for the CMS servers by following the procedure described in [SAP Note 2603573](#). When prompted for the CMS database details, use the connection information and credentials of the database created in the previous section. The CMS servers should now start correctly. You can now install the remaining tiers, such as the processing tier or web tier, as required by your design. The installation instructions for the primary region also apply to the DR region.

### *Creating an Amazon EFS Filestore for the DR Region*

The filestore for the DR region in the Amazon EFS file system must be a copy of the filestore in the primary region.

1. Stop the input and output file servers in the DR region and follow the instructions [earlier in this guide](#) to create an Amazon EFS file system.
2. Prepare S3 buckets in both the primary and DR regions for copying Amazon EFS content.
3. Copy the filestore in Amazon EFS in the primary region to the S3 bucket in that region by using simple scripts and setting up a frequency that matches your database snapshots. The frequency depends on your RPO.
4. Enable cross-region replication from the bucket in the primary region to the DR region.
5. Before starting the SAP BI servers in the DR region, copy the contents of the bucket to the storage tier servers' Amazon EFS mount point for the filestore.
6. Set up correct permissions for the Amazon EFS file system for the filestore.
7. Restart the input and output file servers in the DR region. You should now be able to verify that the SAP BusinessObjects BI Platform system in the DR region has all the data of the primary region, as of the time the database and filestore were last copied.

### *Performing the Failover*

When you complete all the steps described in the previous sections, the SAP BusinessObjects BI Platform system will be up and running in the DR region with primary region data. You can shut down the EC2 instances in the DR region when you don't need

them. Monitor the copy of the database and Amazon EFS file system to make sure that they're active according to your configuration. In the event of a disaster or DR drill, follow these steps:

1. Start the CMS database in the DR region using one of these options:
  - If you're using a Read Replica, open the [Amazon RDS console](#) in the DR region, and then choose **Instances**. Select the DB instance that is the Read Replica of the production database, and choose **Instance Actions, Promote Read Replica** to start the CMS database in the DR region.
  - If you're using snapshots, open the [Amazon RDS console](#) in the DR region, and then choose **Snapshots**. Select the latest snapshot and restore that snapshot to create a CMS database in the DR region.
2. If you're using snapshots, update the data source settings for SAP BusinessObjects BI Platform in the DR region by following the instructions in the section *To select a new or existing CMS database on UNIX* in the [SAP BusinessObjects BI Platform Administration Guide](#). Use the RDS instance created with the latest CMS database snapshot in the primary region.
3. Copy the contents of the S3 bucket that has the latest Amazon EFS content (from the primary region) to the Amazon EFS file system for the input and output filestore, and set up correct permissions.
4. Start SAP BusinessObjects BI Platform in the DR region and perform validation.
5. Redirect your users by changing the value of the CNAME used for SAP BusinessObjects BI Platform from the Application Load Balancer for the primary region to the load balancer for the DR region.

## DR for SAP Data Services

As described earlier in the [HA for SAP Data Services](#) section, you can install SAP Data Services on either an existing SAP BusinessObjects BI Platform node or on SAP BusinessObjects Information Platform Services (IPS). In both cases, the DR approach described previously in this section applies to SAP Data Services as well.

## Summary

SAP customers can now deploy SAP BusinessObjects BI Platform solution and landscapes on the scalable, on-demand Amazon EC2 platform in a highly available manner without having to invest in costly capital expenditures for the underlying infrastructure. By combining the flexibility of the AWS platform with SAP installation techniques, SAP customers can greatly improve the availability of their deployments. For further information, including case studies of customers who have deployed SAP systems on AWS, see [SAP Case Studies](#) on the AWS website.

## Appendix A: Resources for Database HA Solutions

**Note** The SAP notes listed in this appendix require an SAP ONE Support Launchpad user account. For more information, see the [SAP Support](#) website.

### SAP HANA

- [SAP HANA on AWS Implementation and Operations Guide](#) (AWS website)
- [Migrating SAP HANA Systems to X1 Instances on AWS](#) (AWS website)

SAP Note	Description
<a href="#">2309342</a>	SUSE Linux Enterprise High Availability Extension on AWS
<a href="#">2302728</a>	Supported scenarios with NEC Expresscluster on Amazon Web Services
<a href="#">1662610</a>	Support details for SIOS Protection Suite for Linux

### Microsoft SQL Server

- [High Availability Solutions \(SQL Server\)](#) (Microsoft website)
- [Business continuity and database recovery - SQL Server](#) (Microsoft website)

SAP Note	Description
<a href="#">965908</a>	SQL Server Database Mirroring and SAP Applications
<a href="#">1550337</a>	Database Mirroring causes blocking application instance

### IBM DB2

- [SAP on IBM DB2 for Linux, UNIX, and Windows](#) (SAP website)
- [High Availability Disaster Recovery for DB2 for LUW](#) (SAP website)

SAP Note	Description
<a href="#">1612105</a>	FAQ for DB2 High Availability Disaster Recovery (HADR)
<a href="#">1568539</a>	HADR – Virtual IP or Automatic Client Reroute

### MySQL

- [MySQL Enterprise High Availability](#) (MySQL website)

### MaxDB

- [SAP MaxDB Community](#) (SAP website)

- [SAP MaxDB Standby Database](#) (SAP website)
- [MaxDB Hot Standby](#) (SAP website)

SAP Note	Description
<a href="#">952783</a>	High Availability for the SAP MaxDB Database

## SAP ASE

- [SAP Applications on SAP Adaptive Server Enterprise](#) (SAP website)
- [SAP ASE HADR Users Guide](#) (SAP website)

## SQL Anywhere

- [Implementing High Availability for SAP Business Objects CMS and Audit Databases](#) (SAP website)

## Oracle

- [SAP on Oracle](#) (SAP website)
- [Data Guard Concepts and Administration](#) (Oracle website)



## Appendix B: SAP Notes for HA/DR Configuration

SAP Note	Title
<a href="#">1339944</a>	How to delete an unused node that is listed in the Central Management Console
<a href="#">1395589</a>	How to delete server references in the CMS
<a href="#">1256185</a>	How to delete an SIA from the CCM
<a href="#">1802343</a>	How to backup/restore the SubVersion Version Management System
<a href="#">1805615</a>	How to backup and restore the Promotion Management Overrides Database in BI 4.x
<a href="#">2253871</a>	How to rebuild index in Platform Search of BI 4.x
<a href="#">1275068</a>	Backup and restore BusinessObjects Enterprise XI 3.x / BI 4.x system from one environment to another environment (video)
<a href="#">2066464</a>	How to delete old watches that can still be seen in Monitoring after restoring a system from a backup
<a href="#">2306146</a>	Disaster recovery backup and restore best practices for restoring an entire BI 4.x system from one environment to another environment
<a href="#">2034684</a>	BI4.x - How to clone a BusinessObjects Server
<a href="#">2603573</a>	Creating a BI disaster recovery system
<a href="#">1927029</a>	How to create a backup of SAP BusinessObjects Business Intelligence 4.0/4.1 system using Promotion Management / LifeCycle Manager
<a href="#">1261514</a>	How do you change IP address in XI 3.x and BI4.x systems
<a href="#">1229417</a>	What are the requirements for a CMS cluster and how do you test to confirm cluster communication?

## Appendix C: Additional Tips

### Tagging AWS Resources

Adding tags to AWS objects will make it much easier to manage the SAP HA environment, and can also help you search for resources quickly. You can use Amazon EC2 API calls in conjunction with a special tag filter. For more information about tagging resources, see the [AWS documentation](#).

### Third-Party Software Components

Additional third-party components might be integral to running business processes within an SAP environment. After you determine your requirements, consider leveraging some of the concepts discussed in this guide, such as:

- Installing third-party software components on multiple instances
- Creating Amazon EBS-backed AMI images of key third-party systems, so you can launch them on demand in case of failures
- Using multiple interfaces to control access to specific software components
- Using multiple Availability Zones for critical third-party software components

© 2018, Amazon Web Services, Inc. or its affiliates. All rights reserved.

### **Notices**

This document is provided for informational purposes only. It represents AWS's current product offerings and practices as of the date of issue of this document, which are subject to change without notice. Customers are responsible for making their own independent assessment of the information in this document and any use of AWS's products or services, each of which is provided "as is" without warranty of any kind, whether express or implied. This document does not create any warranties, representations, contractual commitments, conditions or assurances from AWS, its affiliates, suppliers or licensors. The responsibilities and liabilities of AWS to its customers are controlled by AWS agreements, and this document is not part of, nor does it modify, any agreement between AWS and its customers.

The software included with this paper is licensed under the Apache License, Version 2.0 (the "License"). You may not use this file except in compliance with the License. A copy of the License is located at <http://aws.amazon.com/apache2.0/> or in the "license" file accompanying this file. This code is distributed on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License.