Hybrid microservices disaster recovery reference architecture: Implementing cross-site replication for Cloudant

This document is created as part of the Hybrid Microservices Disaster Recovery Reference Architecture. The purpose of this reference architecture is to demonstrate how to construct an application written with microservices that can fail over from an active site to a passive site. As part of that reference architecture, this document will show setting up a stateless microservice in Bluemix® connected to the Cloudant® database (DB) in Bluemix with a multisite replication.

Environment

The following diagram explains the environment used for this use case. This document explains the steps required for achieving the item shown as (7) in the diagram.

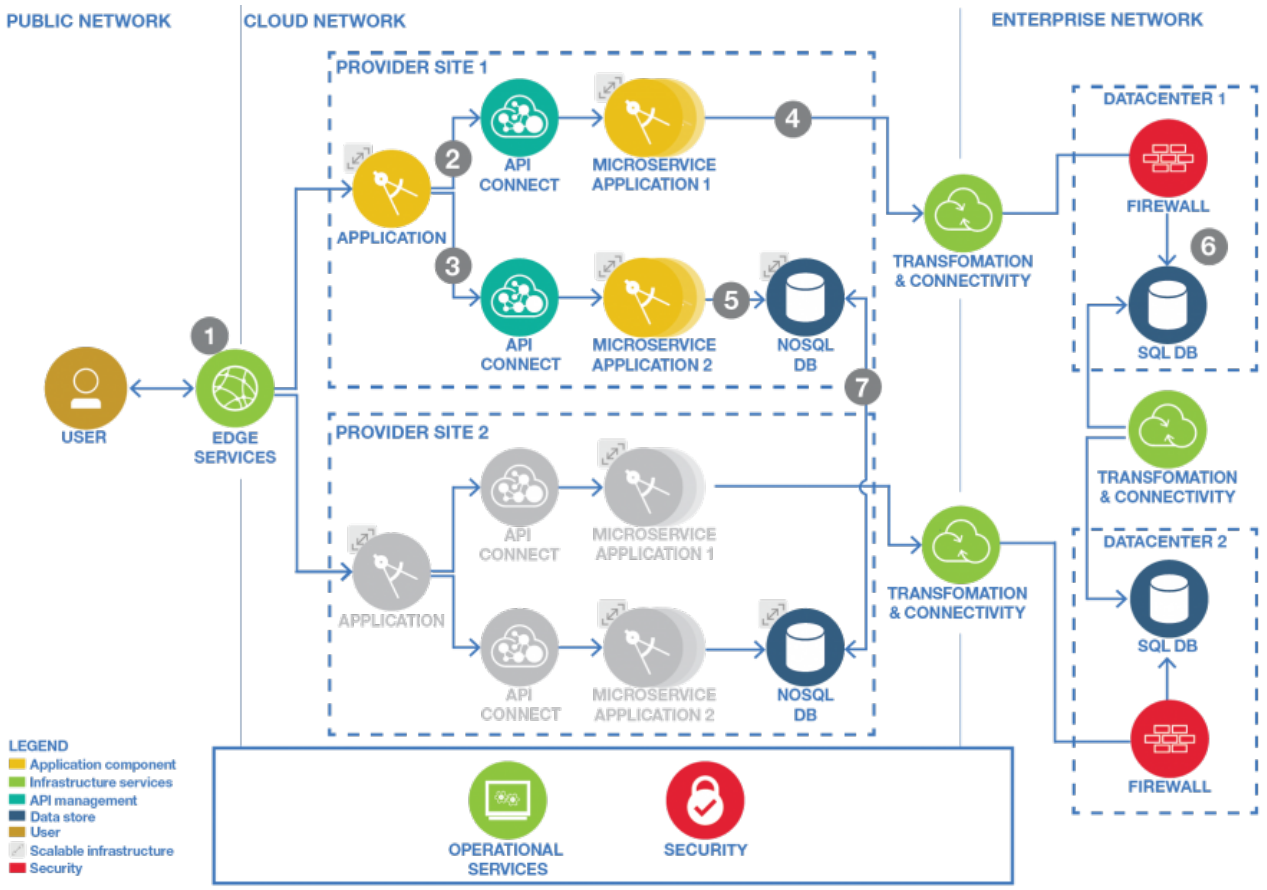

Setup procedure

Follow the steps in this section to recreate the configuration of the application used in our example.

Create the basic microservices application

The follow prerequisites are required to create the basic microservices application:

- You must have a Bluemix account
- The Bluemix account must have sufficient remaining quota to deploy the application

You can begin by following the steps in this blog posts to create the basic microservices application: Microservices store sample on Bluemix.

The steps need to be repeated at each of the disaster recovery sites. In our example, we will deploy the application in two regions:

- U.S. South (primary site)
- United Kingdom (DR site 1)

When you follow the steps above, note that in step 3 of the instructions you will see the following screen:

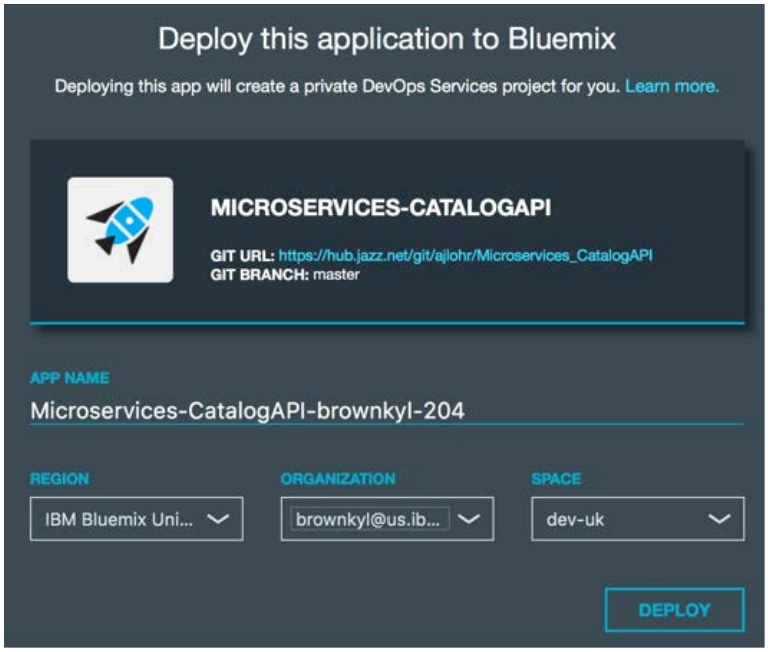

To deploy the application into each region, you need to select the region in the left-hand drop-down. Be sure to click Deploy once for each of the two regions. Once you have the applications up and running at each region, you should see all three applications running with the green status, as shown below. Note that the name of the applications may differ based on what you choose during the deployment.

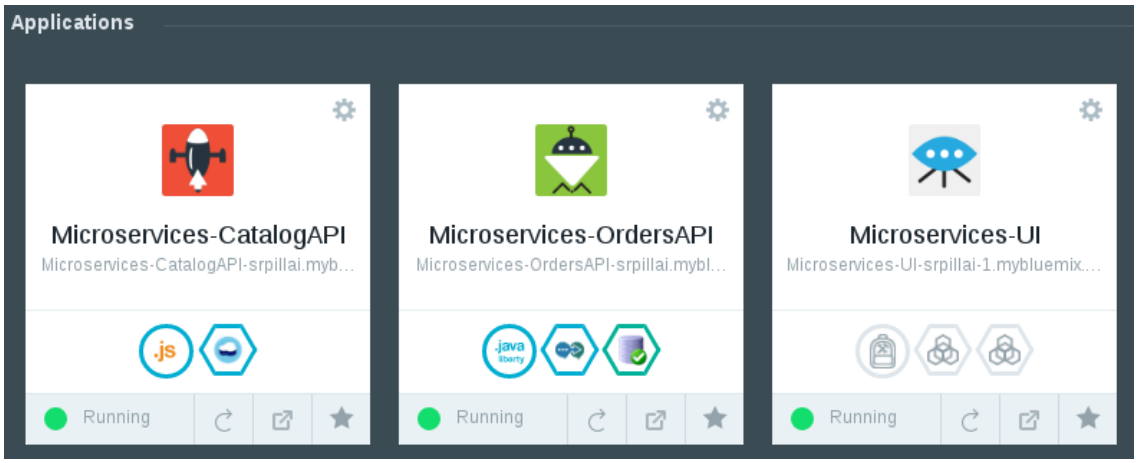

Cloudant database replication

Step 1. Identify the Cloudant credentials

Navigate to the primary (US-South) site in Bluemix and select the Microservices CatalogAPI Application from the dashboard:

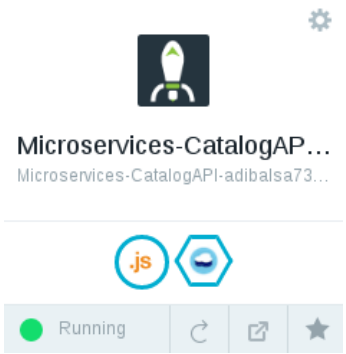

Select Environment Variables from the left pane and examine the Cloudant credentials from the VCAP_SERVICES tab:

Make sure to make note of the credential details. These details will be required in order to perform the database replication. The bold values below are used in that setup (your credential values will be different):


```
{
"cloudantNoSQLDB": [
  {
   "name": "myMicroservicesCloudant",
   "label": "cloudantNoSQLDB",
   "plan": "Shared",
   "credentials": {
    "username": "8fd85255-21d3-457c-9cc2-dcf1b4781a8e-bluemix",
    "password": "8b3901b441b2a0a7d4dcf6dc6188bc2b0b59d7e2765141ef20ff099a97c37dfd",
    "host": "8fd85255-21d3-457c-9cc2-dcf1b4781a8e-bluemix.cloudant.com","port": 443,
    "url": "https://8fd85255-21d3-457c-9cc2-dcf1b4781a8e-bluemix:
8b3901b441b2a0a7d4dcf6dc6188bc2b0b59d7e2765141ef20ff099a97c37dfd@8fd85255-21d3-457c-9cc2-dcf1b4781a8e-
bluemix.cloudant.com"
    }
   }
  ]
 }
```


Next, navigate to the Disaster Recovery (DR) site (UK) in the Bluemix console. Verify that you have the CatalogAPI application up and running at the DR site and perform the steps above, taking note of the Cloudant credentials. The following credentials are the ones you must take note of (will differ for each implementation)

	Primary site (US-South)	**DR site (UK)**
Username	8fd85255-21d3-457c-9cc2-dcf1b4781a8e-bluemix	18d94a42-e910-49f5-b74e-c5552a18e48d-bluemix
Password	8b3901b441b2a0a7d4dcf6dc6188bc2b0b59d7e276514 1ef20ff099a97c37dfd	aff1aa6642d3c9875c01890be14da32823e6cf4825cd18 0459550dfcc2cbf75d
URL	https://8fd85255-21d3-457c-9cc2-dcf1b4781a8e- bluemix: 8b3901b441b2a0a7d4dcf6dc6188bc2b0b59d7e276514 1ef20ff099a97c37dfd@8fd85255-21d3-457c-9cc2- dcf1b4781a8e-bluemix.cloudant.com	https://18d94a42-e910-49f5-b74e-c5552a18e48d-bluemix:aff1aa6642d3c9875c01890be14da32823e6cf 4825cd180459550dfcc2cbf75d@18d94a42- e910-49f5-b74e-c5552a18e48d-bluemix.cloudant.com

Step 2. Identify the database

From the dashboard, navigate to the catalog application and click on the Cloudant DB icon.

This will bring up the Cloudant page. Click the Launch button in the top-right corner.

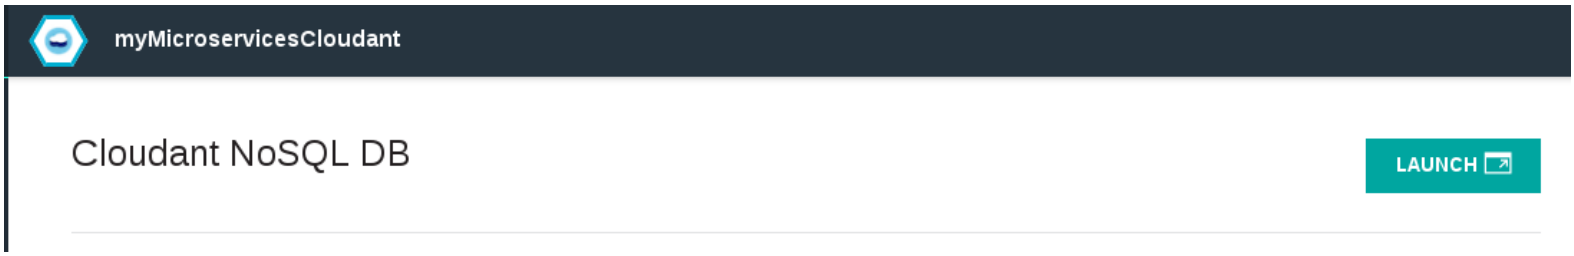

This will bring up the Cloudant DB configuration page for the specific database of that site. In the example described here, it should list a database named “items” as shown below:

The database name must be appended to the end of the URL to obtain the direct link for each specific database. For example, in the section above (Identifying Cloudant credentials), we have identified the URL for the overall database server. To get the URL for the items database, use the following syntax:

https://$USERNAME:$PASSWORD@$REMOTE_USERNAME.cloudant.com/$DATABASE_NAME

In this case, append the name /items at the end. This will result in the speciRic database URL, as shown below:

Primary site items database direct URL:

https://8fd85255-21d3-457c-9cc2-dcf1b4781a8e-bluemix: 8b3901b441b2a0a7d4dcf6dc6188bc2b0b59d7e2765141ef20ff099a97c37dfd@8fd85255-21d3-457c-9cc2-dcf1b4781a8e- bluemix.cloudant.com/items

Disaster site items database direct URL:

https://18d94a42-e910-49f5-b74e-c5552a18e48d- bluemix:aff1aa6642d3c9875c01890be14da32823e6cf4825cd180459550dfcc2cbf75d@18d94a42-e910-49f5-b74e-c5552a18e48d-bluemix.cloudant.com/items

Step 3. Replicate the Cloudant database

Click on the Replication tab in the left pane. Provide the replication details as shown in the picture below:

- Select an ID (optional). In this example, we use US_to_UK_Replication.
- Select the source database as remote and provide the URL of the other site's database link.
- Select the target database as My Databases and type in “items” (the local database to that site). Make sure to select the checkbox to make this replication continuous.

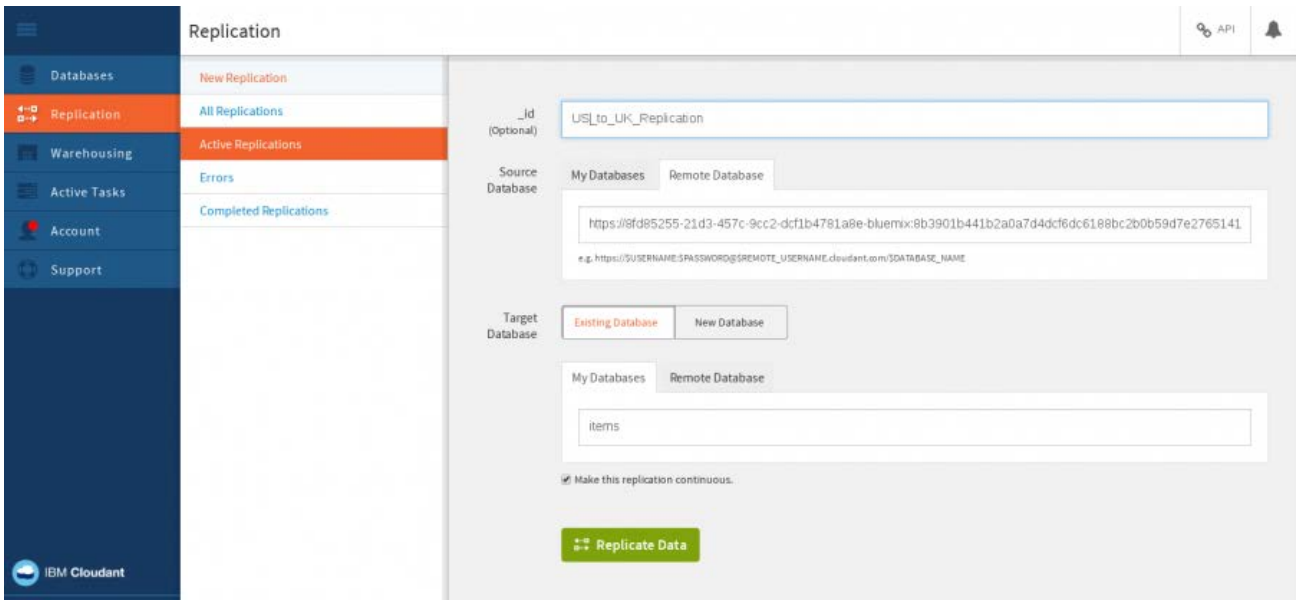

Click **Replicate Data**. This brings up a pop-up to enter the password. Enter the password for the local database to that site (note that you noted this password down in an earlier step).

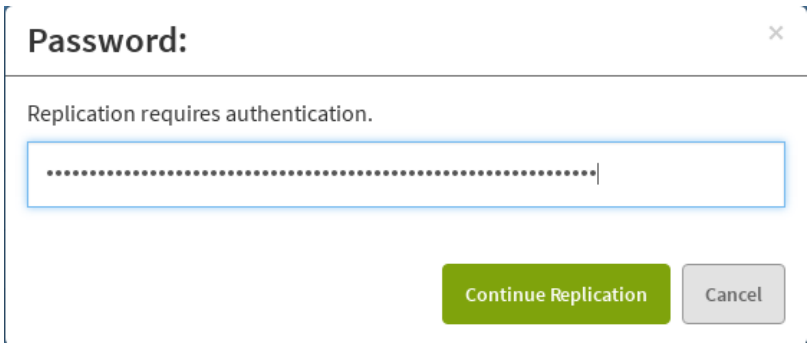

This will bring up the active replications section and show the current replication status.

If you go to the Databases tab on the left pane, you may also notice that there is a replicator database that also appears.

You will need to repeat the step at the other site as well in order to achieve two-way replication.

These steps ensure that the databases at the primary and disaster recovery site remain in sync. Whenever there is a failure at the primary site, the application can fail over to the disaster recovery site and ensure that it has the same data (within the limits of eventual consistency).

Cloudant database API access

The Cloudant database operations can be performed via application programming interfaces (APIs). The APIs can be called via HTTP in a browser or other program or through the command line with CURL. This comes handy while performing activities programmatically. Here are a few steps to make API calls to the Cloudant database.

Configure CURL APIs for Cloudant database access

To perform operations on Cloudant DB through APIs, follow the API reference documentation for Cloudant.

Follow these steps to configure Client URL Request Library (CURL) APIs from a Linux/UNIX system. The generic set of parameters required for any API call include:

- User name
- Password
- Cloudant URL
- Database name
- Content type

Defining these parameters as variable helps to easily execute the APIs. For the example described here, the following variables are set in a shell script named cloudant_environment.sh.

```
vi cloudant_environment.sh
```

```
####### Variable definition for Cloudant connectivity ######
alias prettyjson='python -m json.tool'    # helps to format the JSON output
USER_USA="8fd85255-21d3-457c-9cc2-dcf1b4781a8e-bluemix"    # Username at Primary Site (USA)
PASS_USA="8b3901b441b2a0a7d4dcf6dc6188bc2b0b59d7e2765141ef20ff099a97c37dfd"    # Password $USER_USA
```

```
URL_USA="https://$USER_USA:$PASS_USA@$USER_USA.cloudant.com"     # URL for Primary site
USER_UK="18d94a42-e910-49f5-b74e-c5552a18e48d-bluemix"     # User at DR site (UK)
PASS_UK="aff1aa6642d3c9875c01890be14da32823e6cf4825cd180459550dfcc2cbf75d"     # Password $USER_UK
URL_UK="https://$USER_UK:$PASS_UK@$USER_UK.cloudant.com"     # URL for DR site
DATABASE="items"     # Database name
DBURL_USA="$URL_USA/$DATABASE"     # Database URL at Primary site
DBURL_UK="$URL_UK/$DATABASE"     # Database URL at DR Site
C_TYPE="Content-Type: application/json"     # Content type (Json) header for API
```

Once the above variables are created in a shell script, source that shell script so the variables are present in the current shell session as shown below:

```
chmod u+x ./cloudant_environment.sh
. ./cloudant_environment.sh
echo $USER_USA
8fd85255-21d3-457c-9cc2-dcf1b4781a8e-bluemix
```

Once the variables are sourced, the Cloudant APIs can be called using CURL. (First make sure that CURL and Python are installed on the system where these commands are being executed). The following example shows retrieving the DB information for a specific database from Cloudant:

```
curl -X 'GET' -H "$C_TYPE" $DBURL_USA 2>/dev/null |prettyjson
{"compact_running": false,
"db_name": "items",
"disk_format_version": 6,
"disk_size": 1132176,
"doc_count": 121,
"doc_del_count": 2,
"instance_start_time": "0",
"other": {
    "data_size": 37761
},
"purge_seq": 0,
"sizes": {
    "active": 55375,
    "external": 37761,
    "file": 1132176
},
"update_seq": "123- g1AAAAILeJzLYWBg4MhgTmHQSElKzi9KdUhJMtJLytVNSczRLS5JzEtJLEox0kvOyS9NScwr0ctLLckB6mBKZEiS___
_f1aiCKpeYyL0JikAySR7sHYu0q1OcgBpjwdrlyLD9gSQ9nqwdh7Sbc9jAZIMDUAKaML8rEQO0l0AMWIBxIj9WYkC5L
riAMSI-1mJfOS64gHECGBYCGYBAJeQrtU"
}
```

The above command uses the variable $C_TYPE and $DBURL_USA and prettyjson from the sourced shell script.

Handling disaster

The data between the primary and disaster recovery sites is continuously replicated. During a disaster, switching the application edge services to the disaster recovery site's application URL will be sufficient. There is no special consideration to be done at the Cloudant DB. The steps to switch the edge services is described as part of the DNS switching documentation.

To fail back (re-establish) the link between two sites once the primary site is up and running, you can follow the same replication procedure described in the Cloudant database replication section. However, there is a set of special scenarios other than disaster that you also need to consider. In particular, consider the following situations:

- Synchronization between the sites are lost
- Primary / disaster recovery site database is down/out of sync

Synchronization between the sites are lost

The synchronization between the sites can be lost due to multiple factors, including:

- Network connectivity issues between the sites
- Configuration errors at any of the sites

Check the status of the replication either through the Cloudant web or using an API. The following command helps to get the active replication status. Make sure to configure the API prerequisites as described in the Configure CURL APIs for cloudantDB access section.

```
curl -X 'GET' -H "$C_TYPE" $URL_USA/_active_tasks 2>/dev/null |prettyjson
```


```
[
{
    "changes_pending": 0,
    "checkpoint_interval": 30000,
"checkpointed_source_seq":
"123-g1AAAAPTeJy90rsKwjAUBuCggiIiCk4uOrpY2ly0nfRN9KSJSGkraJ31TfRN9E30TWouXdyMoMsJBHI-
fv6kCKHWti7QRPB4t5dLwbHHs6mAdHooIBewF9iL091RQF54uSxS9aIGiI_Ksky2dRhk6qLJaBDG4cxhzzvKPkD5WE2-
qNy2cYFK8Dehwyr3sHyp3VXlDq3LfDmP8Nd5ySfuWrunyu0YF0cRBuk7rHLPmzfURGd1KPqi7ZaxBcGUUfLTjq19tfZ
N272qZxISiX_as7Xv1n5ou2vs2CeMAP1D7qe1Td99Y8_mAQ-wy_9OXkf5Peo",
    "continuous": true,
    "doc_id": "us_to_uk_replication",
    "doc_write_failures": 0,
    "docs_read": 0,
    "docs_written": 0,
    "missing_revisions_found": 0,
    "node": "dbcore@db2.bm-dal-standard2.cloudant.net",
    "pid": "<0.23952.1932>",
    "replication_id": "d511ac1ca5bfa6b3c7c0b129fc919a03+continuous",
    "revisions_checked": 125,
    "source": "https://8fd85255-21d3-457c-9cc2-dcf1b4781a8e-bluemix.cloudant.com/items/",
"source_seq": "123- g1AAAAPTeJy90rsKwjAUBuCggiIiCk4uOrpY2ly0nfRN9KSJSGkraJ31TfRN9E30TWouXdyMoMsJBHI-
fv6kCKHWti7QRPB4t5dLwbHHs6mAdHooIBewF9iL091RQF54uSxS9aIGiI_Ksky2dRhk6qLJaBDG4cxhzzvKPkD5WE2-
qNy2cYFK8Dehwyr3sHyp3VXlDq3LfDmP8Nd5ySfuWrunyu0YF0cRBuk7rHLPmzfURGd1KPqi7ZaxBcGUUfLTjq19tfZ
N272qZxISiX_as7Xv1n5ou2vs2CeMAP1D7qe1Td99Y8_mAQ-wy_9OXkf5Peo",
    "started_on": 1473251573,
    “target": "https://18d94a42-e910-49f5-b74e-c5552a18e48d-bluemix:*****@18d94a42-e910-49f5-b74e-c5552a18e48d- bluemix.cloudant.com/items/",
    "type": "replication",
    "updated_on": 1473254446,
    "user": "8fd85255-21d3-457c-9cc2-dcf1b4781a8e-bluemix"
    }
  ]
```


In the command output above, the blue text shows that the continuous replication is running without any errors for the us_to_uk_replication ID. In case of a replication failure, you should re-establish the replication as described in the Cloudant database replication section.

Primary/disaster recovery site is down or out of sync

When one of the sites is out of sync or corrupted, the data can be replicated from the site that has the most recent data. Perform the following steps to re-establish the synchronization:

- Delete the out-of-sync database as shown below. For example, assume that the US Cloudant DB is out of sync.

```
curl -X 'DELETE' -H "$C_TYPE" $DBURL_USA
```

```
{"ok":true}
```


- Recreate an empty database as shown below:

```
curl -X 'PUT' -H "$C_TYPE" $DBURL_USA
```

```
{"ok":true}
```


- Create an empty database. Check that the current doc_count is zero to ensure that the DB is empty (this step is optional, as there is a provision to create new database while initiating a DB replication).

```
curl -X 'GET' -H "$C_TYPE" $DBURL_USA 2>/dev/null |prettyjson |grep doc_count
```

```
"doc_count": 0,
```


Create a new replication as described in the Cloudant database replication section.

Replication performance

Replications can severely affect the performance of a Cloudant cluster. Performance testing is recommended to understand the impact on your environment under an increasing number of concurrent replications. Continuous replication can result in a large number of internal calls. This might affect costs for multi-tenant users of Cloudant systems. Continuous replication is disabled by default. Read more about Cloudant replication at https://docs.cloudant.com/replication.html.

Triggering a manual sync using API between two locations helps identify the synchronization speed. Following are some statistics of synchronization speed between two sites (US and UK) explained in this documentation.

Number of docs	Total DB size	Start time	End time	Replication duration
100	10MB	08 Sep 2016 13:08:57 GMT	08 Sep 2016 13:08:50 GMT	7 seconds
1000	100MB	08 Sep 2016 08:24:40 GMT	08 Sep 2016 08:24:43 GMT	57 seconds
4000	400MB	08 Sep 2016 12:20:44 GMT	08 Sep 2016 12:28:01 GMT	7 minutes, 17 seconds

References

- IBM Bluemix Public
- Online store application using microservices and Bluemix
- IBM Cloudant documentation