


# The Ceph RGW archive zone feature



# Who?

- Javier
- Computer Engineer and Software Developer
- Cloud and Virtualization team @ Igalia

 @javimunhoz

 [linkedin.com/in/javiermunhoz](https://www.linkedin.com/in/javiermunhoz)



# This talk

- The new archive zone feature
- The RGW/S3 archive zone in action

# The new archive zone feature


# The new archive zone feature

- The new archive zone federation feature enables full preservation of all objects (including history) in a separate zone
- Available since Ceph Nautilus (v14.2.0)



# The new archive zone feature

- Context
- Use Cases
- Requirements
- Open feature



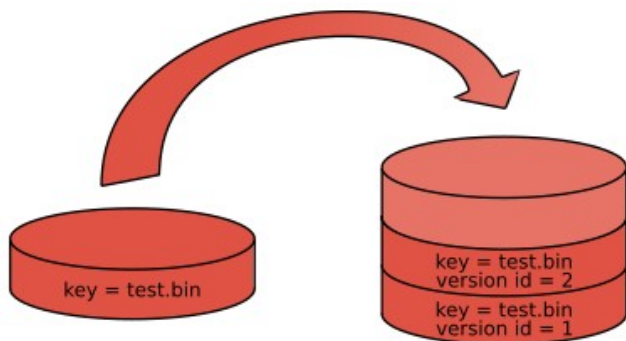
# The new archive zone feature

- Backup
  - “... copy of your current and active data which can be used for operational recoveries in the event that your data is lost or corrupted in some way ...”
- Archives
  - “... to be used as repositories for data which needs to be kept for a long period of time but is not necessarily important for operational activities ...”
- Some comments on ...
  - location, speed, searchability, duration ...



# The new archive zone feature

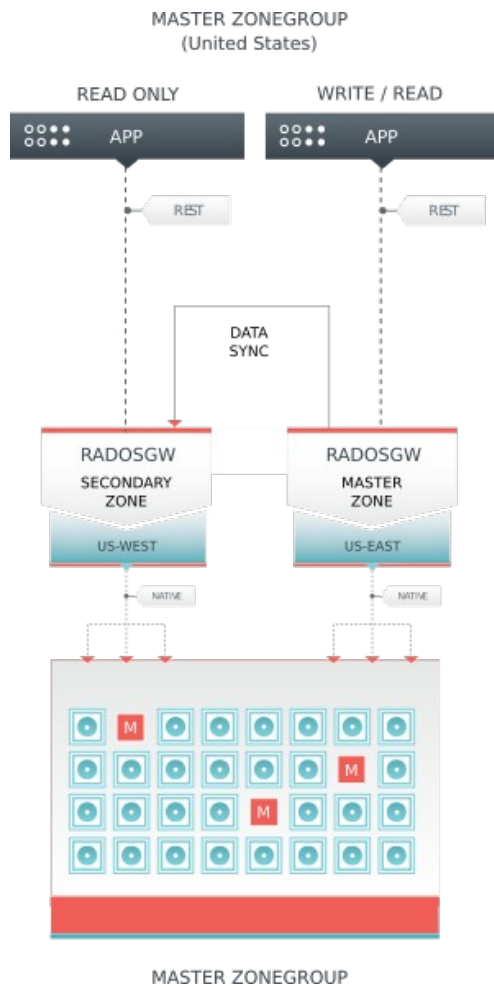
- RGW/S3 Object Versioning



```
$ aws --endpoint=http://rgw:8080 s3api list-object-versions \
  --bucket test-bucket-4
{
  "DeleteMarkers": [
    {
      "Owner": {
        "DisplayName": "Test User",
        "ID": "testuser1"
      },
      "IsLatest": false,
      "VersionId": "NSiqwXFnv7pA8ixFzIXZJw3xi0xls06",
      "Key": "test-key-1",
      "LastModified": "2018-11-14T12:14:05.736Z"
    }
  ],
  "Versions": [
    {
      "LastModified": "2018-11-14T12:27:32.613Z",
      "VersionId": "gydfTpqLM7ex7WJChvVN1TsrivDCZ1b",
      "ETag": "\"b9c85244be9733bc79eca588db7bf306\"",
      "StorageClass": "STANDARD",
      "Key": "test-key-1",
      "Owner": {
        "DisplayName": "Test User",
        "ID": "testuser1"
      },
      "IsLatest": true,
      "Size": 151024
    },
    {
      "LastModified": "2018-11-14T12:13:55.945Z",
      "VersionId": "T7i7USZG-DaqHB0truDhIzaov0tSnVZ",
      "ETag": "\"b9c85244be9733bc79eca588db7bf306\"",
      "StorageClass": "STANDARD",
      "Key": "test-key-1",
      "Owner": {
        "DisplayName": "Test User",
        "ID": "testuser1"
      },
      "IsLatest": false,
      "Size": 151024
    }
  ]
}
```


# The new archive zone feature

- Replication
  - Many ways
  - Unidirectional vs Bi-directional
  - Synchronous vs Asynchronous
  - Data collisions
  - Latency
  - Restore time (“down time”)



# The new archive zone feature

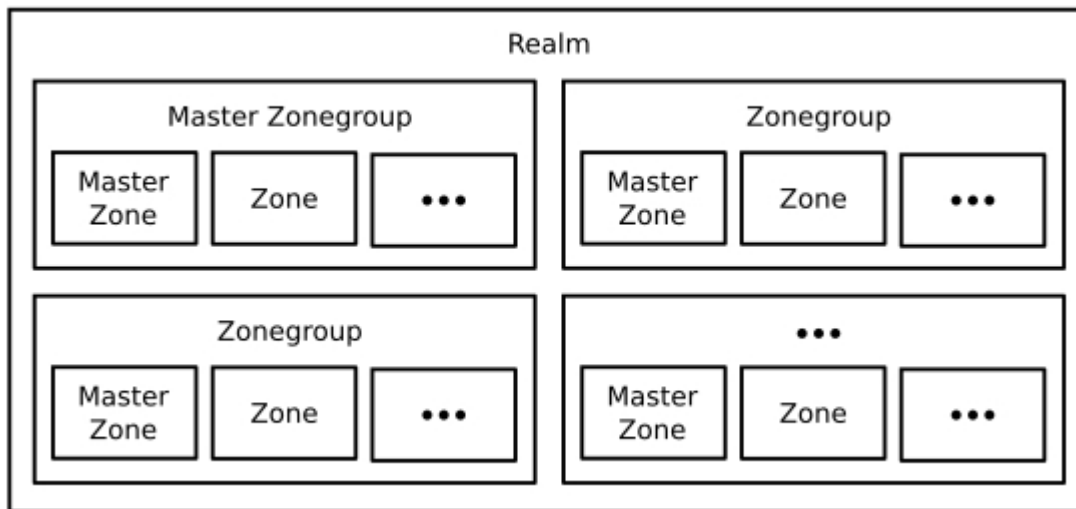
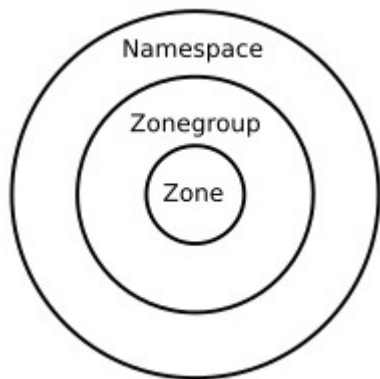
- Restore time (“down time”)
  - hardware repair time + system recovery time
- Recovery time (RTO)
- Data loss (RPO)



source: <http://tiny.cc/qjiu4y>

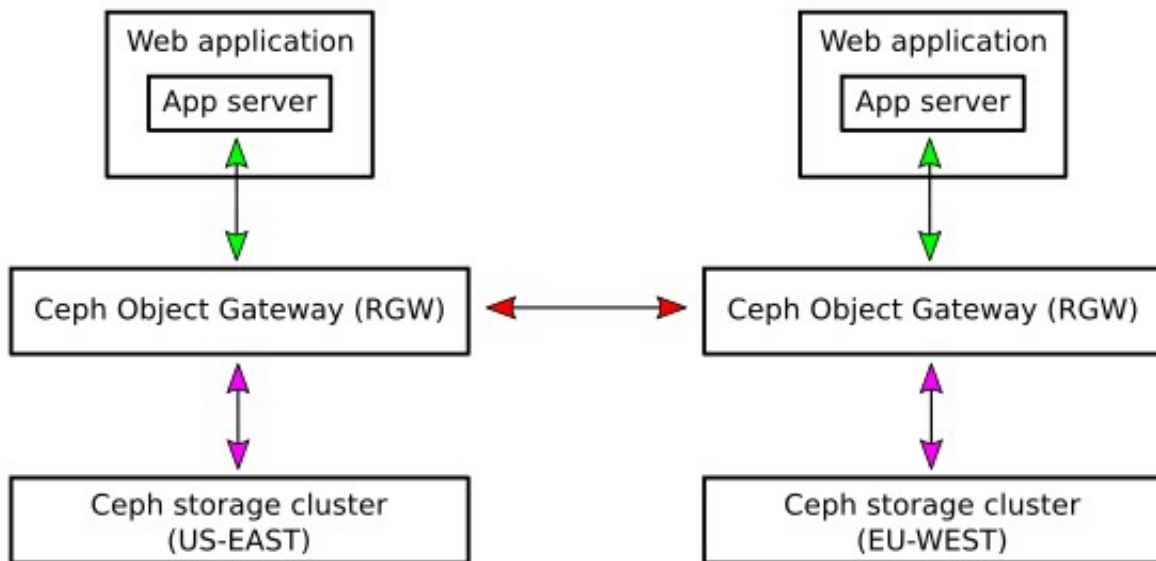
# The new archive zone feature

- The RGW federation model



# The new archive zone feature

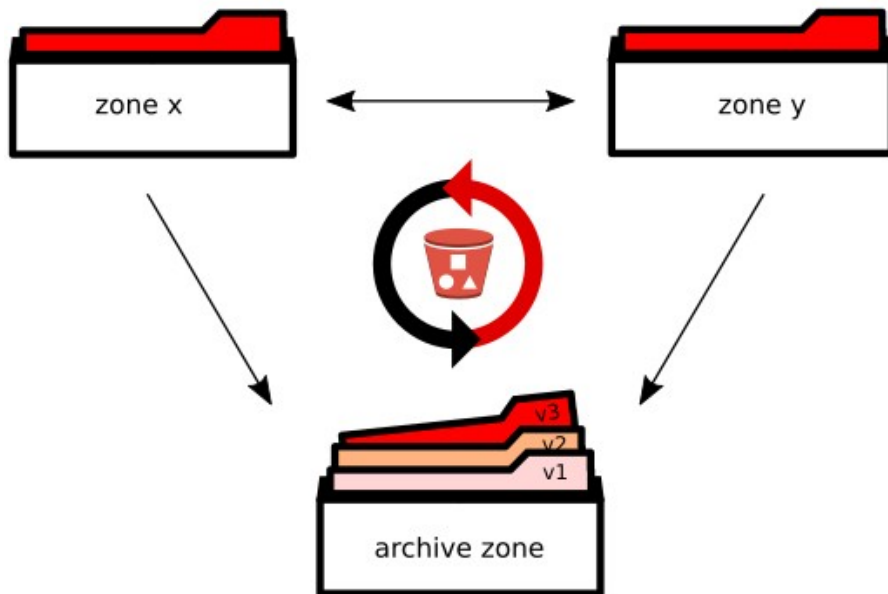
- The RGW multi-site replication



# The new archive zone feature

- The RGW multi-site replication
  - Global object storage
  - Namespaces support
  - Active-Active support
  - Metadata sync is synchronous
  - Data sync is asynchronous (co-routines)
  - Locally immediately consistent, remote eventually consistent
  - Failover / Failback

# The new archive zone feature



# The new archive zone feature

- Extending the RGW federation model
  - Mechanism vs Policy
    - replication framework + plugins
    - RGW Sync modules
  - Implementation
    - archive module + framework modifications
    - special care with the handling of empty buckets and the replication of metadata that required atomic bucket renaming
    - the configuration and handling of the archive zone takes place through the tool 'radosgw-admin' and the endpoint/s associated with the archive zone
    - ...




# The RGW/S3 archive zone in action

# The RGW/S3 archive zone in action



# The RGW/S3 archive zone in action



# The RGW/S3 archive zone in action

1

<http://docs.ceph.com/docs/nautilus/radosgw/multisite/>

2

**# set up archive zone**

```
$ radosgw-admin zone modify --rgw-zone=mad --tier-type=archive
```

```
$ radosgw-admin zone modify bcn --sync-from-all=false
```

```
$ radosgw-admin period update --commit
```

# The RGW/S3 archive zone in action

```
$ aws --endpoint=$BCN s3api create-bucket \  
  --bucket test-bucket-1  
  
$ aws --endpoint=$BCN s3api put-object \  
  --bucket test-bucket-1 \  
  --key test-file-1-in \  
  --body /tmp/test-1.file  
{  
  "ETag": "\"f7d53f25b67715fd4959eb7787de7902\""  
}  
  
$ aws --endpoint=$BCN s3api put-object \  
  --bucket test-bucket-1 \  
  --key test-file-1-in \  
  --body /tmp/test-2.file  
{  
  "ETag": "\"931606baaa7a2b4ef61198406f8fc3f4\""  
}
```

# The RGW/S3 archive zone in action

```
$ aws --endpoint=$BCN s3api list-object-versions \  
  --bucket test-bucket-1  
{  
  "Versions": [  
    {  
      "LastModified": "2019-04-09T20:12:40.354Z",  
      "VersionId": "null",  
      "ETag": "\"931606baaa7a2b4ef61198406f8fc3f4\"",  
      "StorageClass": "STANDARD",  
      "Key": "test-file-1-in",  
      "Owner": {  
        "DisplayName": "Test User",  
        "ID": "testuser1"  
      },  
      "IsLatest": true,  
      "Size": 133792  
    }  
  ]  
}
```

```
$ aws --endpoint=$MAD s3api list-object-versions \  
  --bucket test-bucket-1  
{  
  "Versions": [  
    {  
      "LastModified": "2019-04-09T20:12:40.354Z",  
      "VersionId": "fTIW43biJSe6kOZsB8SIUYp7Aesghbz",  
      "ETag": "\"931606baaa7a2b4ef61198406f8fc3f4\"",  
      "StorageClass": "STANDARD",  
      "Key": "test-file-1-in",  
      "Owner": {  
        "DisplayName": "Test User",  
        "ID": "testuser1"  
      },  
      "IsLatest": true,  
      "Size": 133792  
    },  
    {  
      "LastModified": "2019-04-09T20:12:30.008Z",  
      "VersionId": "WJdcKZqWvWC2mPDqu4AiWZG2ivYx.QQ",  
      "ETag": "\"f7d53f25b67715fd4959eb7787de7902\"",  
      "StorageClass": "STANDARD",  
      "Key": "test-file-1-in",  
      "Owner": {  
        "DisplayName": "Test User",  
        "ID": "testuser1"  
      },  
      "IsLatest": false,  
      "Size": 141528  
    }  
  ]  
}
```

# The RGW/S3 archive zone in action

```
$ aws --endpoint=$BCN s3api delete-object \  
  --bucket test-bucket-1 \  
  --key test-file-1-in  
  
$ aws --endpoint=$BCN s3api delete-bucket \  
  --bucket test-bucket-1  
  
$ aws --endpoint=$BCN s3api list-buckets  
{  
  "Owner": {  
    "DisplayName": "Test User",  
    "ID": "testuser1"  
  },  
  "Buckets": []  
}  
  
$ aws --endpoint=$MAD s3api list-buckets  
{  
  "Owner": {  
    "DisplayName": "Test User",  
    "ID": "testuser1"  
  },  
  "Buckets": [  
    {  
      "CreationDate": "2019-04-09T20:11:59.137Z",  
      "Name": "test-bucket-1-deleted-d5d4071b41e69cd7bbf6207037b416dc"  
    }  
  ]  
}
```

# The RGW/S3 archive zone in action

```
$ aws --endpoint=$MAD s3api list-object-versions --bucket test-bucket-1-deleted-d5d4071b41e69cd7bbf6207037b416dc
{
  "Versions": [
    {
      "LastModified": "2019-04-09T20:12:40.354Z",
      "VersionId": "FTIW43biJSe6kOZsB8SIUYp7Aesghbz",
      "ETag": "\"931606baaa7a2b4ef61198406f8fc3f4\"",
      "StorageClass": "STANDARD",
      "Key": "test-file-1-in",
      "Owner": {
        "DisplayName": "Test User",
        "ID": "testuser1"
      },
      "IsLatest": true,
      "Size": 133792
    },
    {
      "LastModified": "2019-04-09T20:12:30.008Z",
      "VersionId": "WJdcKZqWvWC2mPDqu4AiwZG2ivYx.QQ",
      "ETag": "\"f7d53f25b67715fd4959eb7787de7902\"",
      "StorageClass": "STANDARD",
      "Key": "test-file-1-in",
      "Owner": {
        "DisplayName": "Test User",
        "ID": "testuser1"
      },
      "IsLatest": false,
      "Size": 141528
    }
  ]
}
```



# The RGW/S3 archive zone in action

```
$ aws --endpoint=$BCN s3api create-bucket --bucket test-bucket-1

$ aws --endpoint=$BCN s3api put-object --bucket test-bucket-1 --key test-file-1-in --body /tmp/test-3.file
{
  "ETag": "\"b89143b082f931507c92d5be5dfae035\""
}

$ aws --endpoint=$BCN s3api list-object-versions --bucket test-bucket-1
{
  "Versions": [
    {
      "LastModified": "2019-04-09T20:16:36.381Z",
      "VersionId": "null",
      "ETag": "\"b89143b082f931507c92d5be5dfae035\"",
      "StorageClass": "STANDARD",
      "Key": "test-file-1-in",
      "Owner": {
        "DisplayName": "Test User",
        "ID": "testuser1"
      },
      "IsLatest": true,
      "Size": 63704
    }
  ]
}
```

# The RGW/S3 archive zone in action

```
$ aws --endpoint=$MAD s3api list-buckets
{
  "Owner": {
    "DisplayName": "Test User",
    "ID": "testuser1"
  },
  "Buckets": [
    {
      "CreationDate": "2019-04-09T20:15:51.668Z",
      "Name": "test-bucket-1"
    },
    {
      "CreationDate": "2019-04-09T20:11:59.137Z",
      "Name": "test-bucket-1-deleted-d5d4071b41e69cd7bbf6207037b416dc"
    }
  ]
}
```

```
$ aws --endpoint=$MAD s3api list-object-versions --bucket test-bucket-1
{
  "Versions": [
    {
      "LastModified": "2019-04-09T20:16:36.381Z",
      "VersionId": "3dP4-A0JSVXoQ8ICrlnhEksKz8DDpTT",
      "ETag": "\"b89143b082f931507c92d5be5dfae035\"",
      "StorageClass": "STANDARD",
      "Key": "test-file-1-in",
      "Owner": {
        "DisplayName": "Test User",
        "ID": "testuser1"
      },
      "IsLatest": true,
      "Size": 63704
    }
  ]
}
```

# The RGW/S3 archive zone in action

```
$ aws --endpoint=$BCN s3api delete-object --bucket test-bucket-1 --key test-file-1-in
```

```
$ aws --endpoint=$BCN s3api list-object-versions --bucket test-bucket-1
```

```
$ aws --endpoint=$MAD s3api list-object-versions --bucket test-bucket-1
```

```
{  
  "Versions": [  
    {  
      "LastModified": "2019-04-09T20:16:36.381Z",  
      "VersionId": "3dP4-A0JSVXoQ8ICrlnhEksKz8DDpTT",  
      "ETag": "\"b89143b082f931507c92d5be5dfae035\"",  
      "StorageClass": "STANDARD",  
      "Key": "test-file-1-in",  
      "Owner": {  
        "DisplayName": "Test User",  
        "ID": "testuser1"  
      },  
      "IsLatest": true,  
      "Size": 63704  
    }  
  ]  
}
```

# The RGW/S3 archive zone in action

```
$ aws --endpoint=$MAD s3api list-object-versions --bucket test-bucket-1-deleted-d5d4071b41e69cd7bbf6207037b416dc
{
  "Versions": [
    {
      "LastModified": "2019-04-09T20:12:40.354Z",
      "VersionId": "fTIW43biJSe6kOZsB8SIUYp7Aesghbz",
      "ETag": "\"931606baaa7a2b4ef61198406f8fc3f4\"",
      "StorageClass": "STANDARD",
      "Key": "test-file-1-in",
      "Owner": {
        "DisplayName": "Test User",
        "ID": "testuser1"
      },
      "IsLatest": true,
      "Size": 133792
    },
    {
      "LastModified": "2019-04-09T20:12:30.008Z",
      "VersionId": "WJdcKZqWvWC2mPDqu4AiwZG2ivYx.QQ",
      "ETag": "\"f7d53f25b67715fd4959eb7787de7902\"",
      "StorageClass": "STANDARD",
      "Key": "test-file-1-in",
      "Owner": {
        "DisplayName": "Test User",
        "ID": "testuser1"
      },
      "IsLatest": false,
      "Size": 141528
    }
  ]
}
```

# Community



Cephalocon 2018 APAC



Florian Haas @xahteivi · 8h

Folks! In May, the #Ceph community gathers in Barcelona for the annual #Cephalocon. Yes, in Barcelona. Have I mentioned Barcelona? It's in Barcelona. Yes I'm speaking but it's also all about Ceph and it's IN BARCELONA.

You want to be there.



Cephalocon Barcelona 2019 - Ceph  
[ceph.com](https://ceph.com)



@ceph

- [ceph.com](https://ceph.com)
- [ceph.com/community](https://ceph.com/community)
- [ceph.com/get-involved](https://ceph.com/get-involved)

# Thanks!

## Q+A



 @javimunhoz