

ADVANCED TECHNOLOGY GROUP (ATG)



Accelerate with ATG Webinar

IBM Storage Ceph S3 Object Storage Deep Dive

John Shubeck – ATG Storage Technical Specialist

Date: May 30, 2024



Accelerate with ATG Technical Webinar Series

Advanced Technology Group experts cover a variety of technical topics.

Audience: Clients who have or are considering acquiring IBM Storage solutions. Business Partners and IBMers are also welcome.

To automatically receive announcements of upcoming Accelerate with IBM Storage webinars, Clients, Business Partners and IBMers are welcome to send an email request to accelerate-join@hursley.ibm.com.



2024 Upcoming Webinars – Register Here!

[IBM TS7700 Tape Solution Overview 201](#) – June 18th, 2024

[Why IBM Cloud Object Storage System and, why now?](#) – June 20th, 2024

[Unleash the Power of the IBM FlashSystem 5300](#) – June 25th, 2024

[IBM C-Type SAN Analytics: Real-time and Always-on SAN Performance visibility at Scale](#) – July 23rd, 2024

Important Links to bookmark:



ATG Accelerate Site: <https://ibm.biz/BdSUFN>

ATG MediaCenter Channel: <https://ibm.biz/BdfEgQ>

Offerings

Client Technical Workshops

- IBM DS8900F Advanced Functions
- IBM Fusion & Ceph: A Deep Dive into Next Gen Storage
- IBM FlashSystem Deep Dive & Advanced Functions
- IBM Cyber Resiliency with IBM Storage Defender

TechZone Test Drive / Demo's

- IBM Storage Scale and Storage Scale System GUI
- IBM Storage Virtualize Test Drive
- IBM DS8900F Storage Management Test Drive
- Managing Copy Services on the DS8000 Using IBM Copy Services Manager Test Drive
- IBM DS8900F Safeguarded Copy (SGC) Test Drive
- IBM Cloud Object Storage Test Drive - (Appliance based)
- IBM Cloud Object Storage Test Drive - (VMware based)
- IBM Storage Protect Live Test Drive
- **IBM Storage Ceph Test Drive - (VMware based)**

Please reach out to your IBM Representative or Business Partner for more information.

***IMPORTANT* The ATG team serves clients and Business Partners in the Americas, concentrating on North America.**

Registration Open!

Storage @ IBM TechXchange Conference 2024

October 21-24, 2024

Mandalay Bay | Las Vegas
#IBMTechXchange

Key Learnings

- Practical how-to advice
- Patterns and best practices
- Success stories, IBM PoV, proven techniques

Featured Products

IBM Storage Defender

IBM Storage Fusion

IBM Storage Scale + IBM Storage Ceph

IBM Tape + IBM SAN

IBM Storage FlashSystem + IBM Storage DS8000

Collaborate. Learn. Play.

Community

IBM Champions

User Groups

Tech Peers

Business Partners





Sandbox

Network

Learn

Collaborate

Play



Accelerate your Career

Labs (Instructor-Led, Self-paced)

IBM Certification Testing

Earn up to 25 hours in CPE credits

Breakout Sessions

Trends and Directions

User Groups

Product Deep Dives


Meet the Expert

Professional Development

Show the Code

Birds of a Feather

Academic/Research



Roadmaps

Go deep with people in the know and set the stage for where IBM is going in the future



<https://www.ibm.com/community/ibm-techxchange-conference/>

Game On!



Accelerate with ATG Survey

Please take a moment to share your feedback with our team!

You can access this 6-question survey via [Menti.com](https://www.menti.com) with code 1708 6924 or

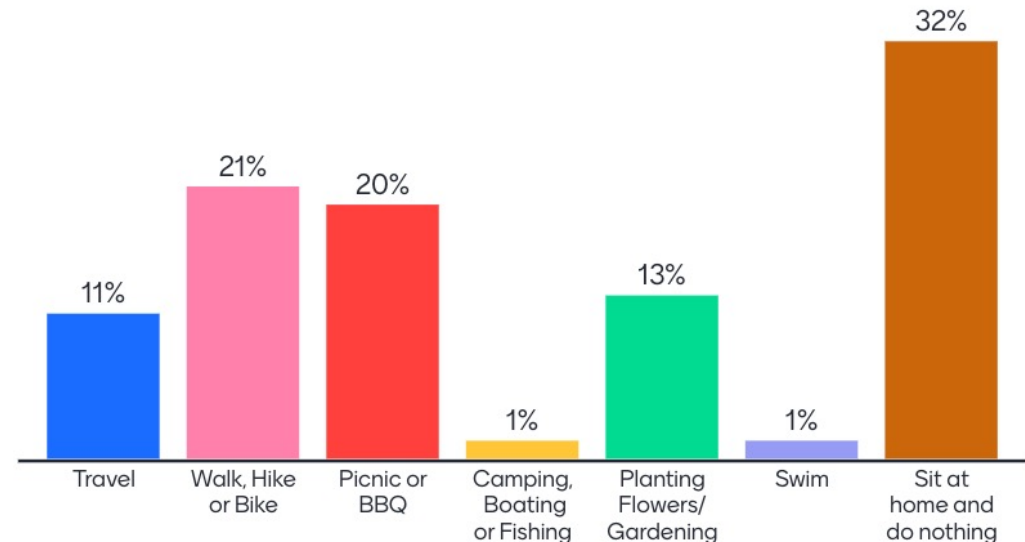
Direct link <https://www.menti.com/alwhyze7z1gz>

Or

QR Code



What are you planning to do for Memorial Day Weekend?



ADVANCED TECHNOLOGY GROUP (ATG)



Accelerate with ATG Webinar

IBM Storage Ceph S3 Object Storage Deep Dive

John Shubeck – ATG Storage Technical Specialist

Date: May 30, 2024



About the Presenter



John Shubeck is an information technology professional with over 42 years of industry experience spanning both the customer and technology provider experience. John is currently serving as a Senior Storage Technical Specialist on IBM Object Storage platforms across all market segments in the Americas.

Introducing our panelists




Shu Mookerjee is a Level 2 Certified Technical Specialist with over twenty years at IBM, working in a variety of roles including sales, management and technology. For the last decade, he has focused exclusively on storage and has been the co-author of four (4) Redbooks. Currently, Shu is part of the Advanced Technology Group where he provides education, technical guidance, Proofs of Concept and Proofs of Technology to IBMers, business partners and clients.

Introducing our panelists



Jerrod Carr is an IBM Principal Storage Technical Specialist in IBM Storage Solutions. Jerrod Carr has been in the Storage industry for over 21 years selling hardware and software for various large technology companies. With beginnings in the Cleversafe IBM team for 8 years providing expertise in Cloud Object Storage, the last 3 years working on the Americas SWAT team as a Senior Storage Specialist providing unstructured data experience to the various markets.



Todd Johnston is a proven applied technologist bringing decades in customer advocacy, technology implementations, service provider enablement, sales engineering, and solution architecture. Todd is especially focused on Ceph innovation and early adoption, technical agility, and disruptive technologies. Todd's expertise in lies storage software, virtualization, and container and hybrid cloud infrastructures. Todd joins us from the IBM SWAT practice in our cross functional team.

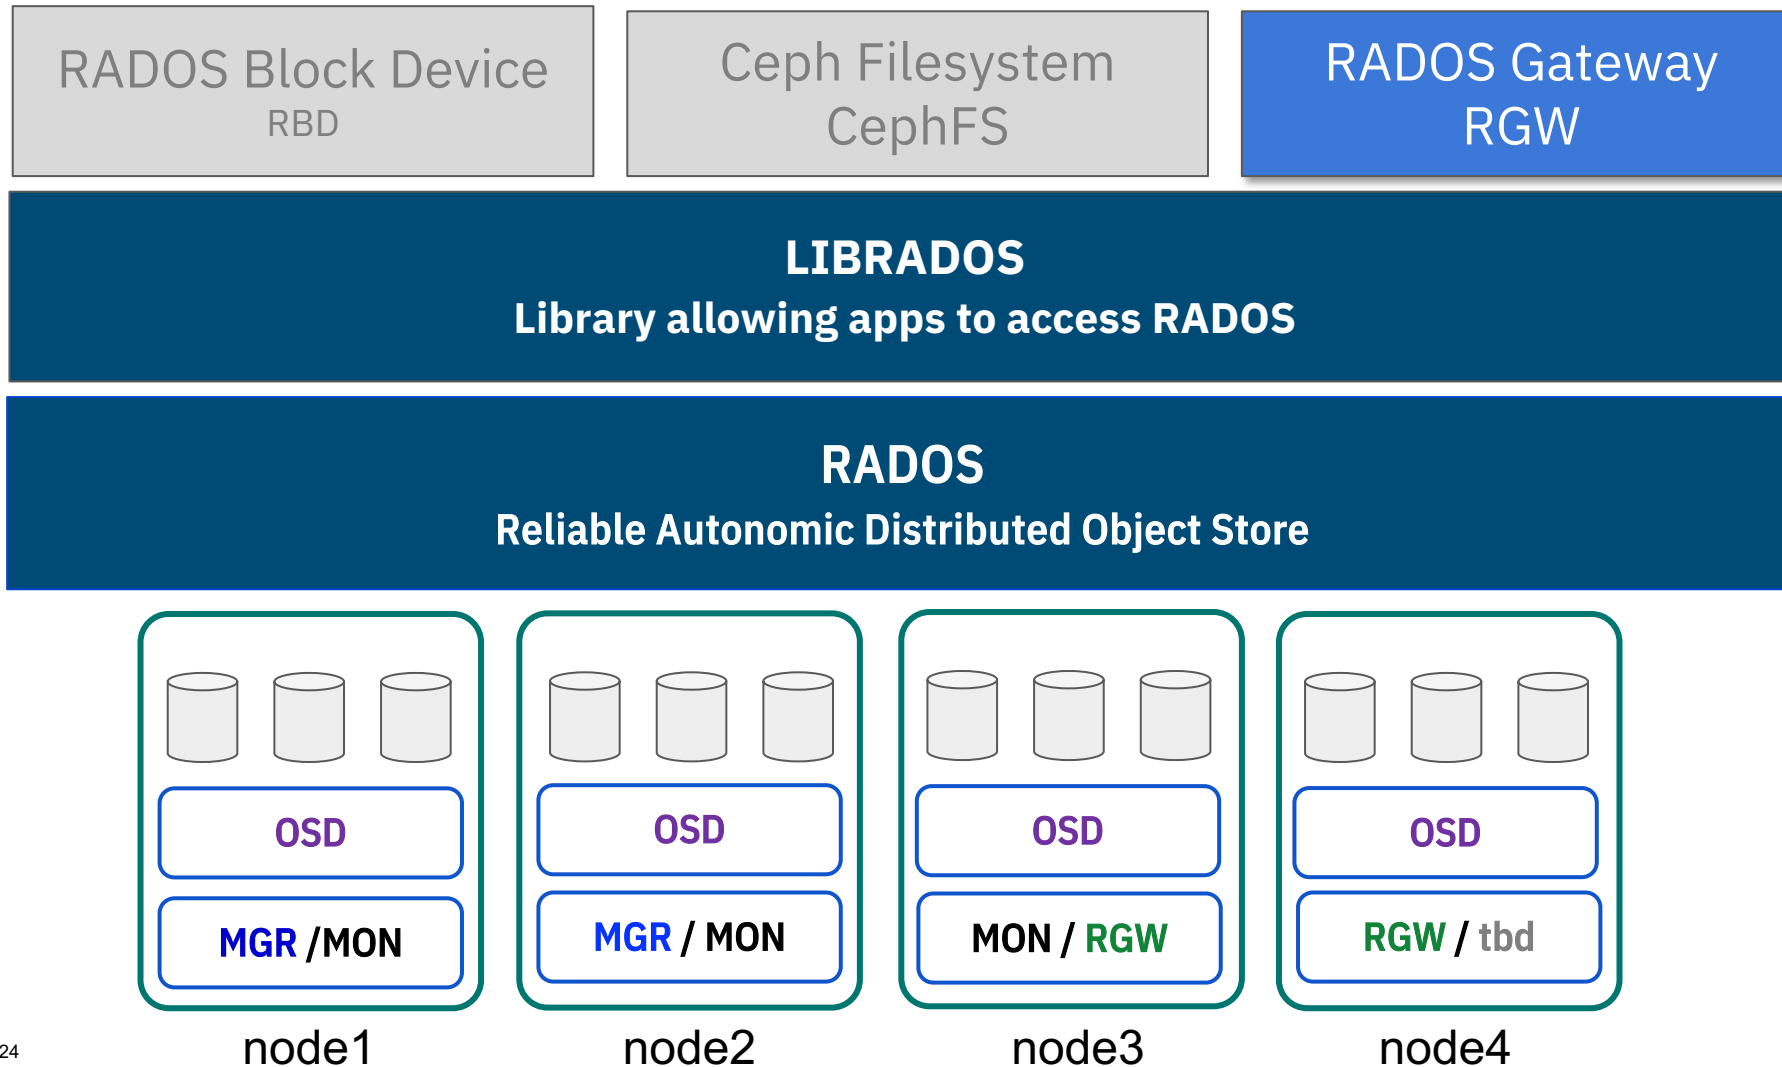
Summary of topics



- Overview of the RADOS Gateway (RGW)
- Configuring the RGW Daemons, Users, Buckets
- Configuring SSL for secure S3 access
- Two site replication (Realm, Zone, Zonegroup)
- Data placement groups
- Multi tenancy introduction
- Configuring multiprotocol (i.e. S3 and NFS)
- Design considerations for POC and Production
- Exploring the S3 API compatibility and feature support

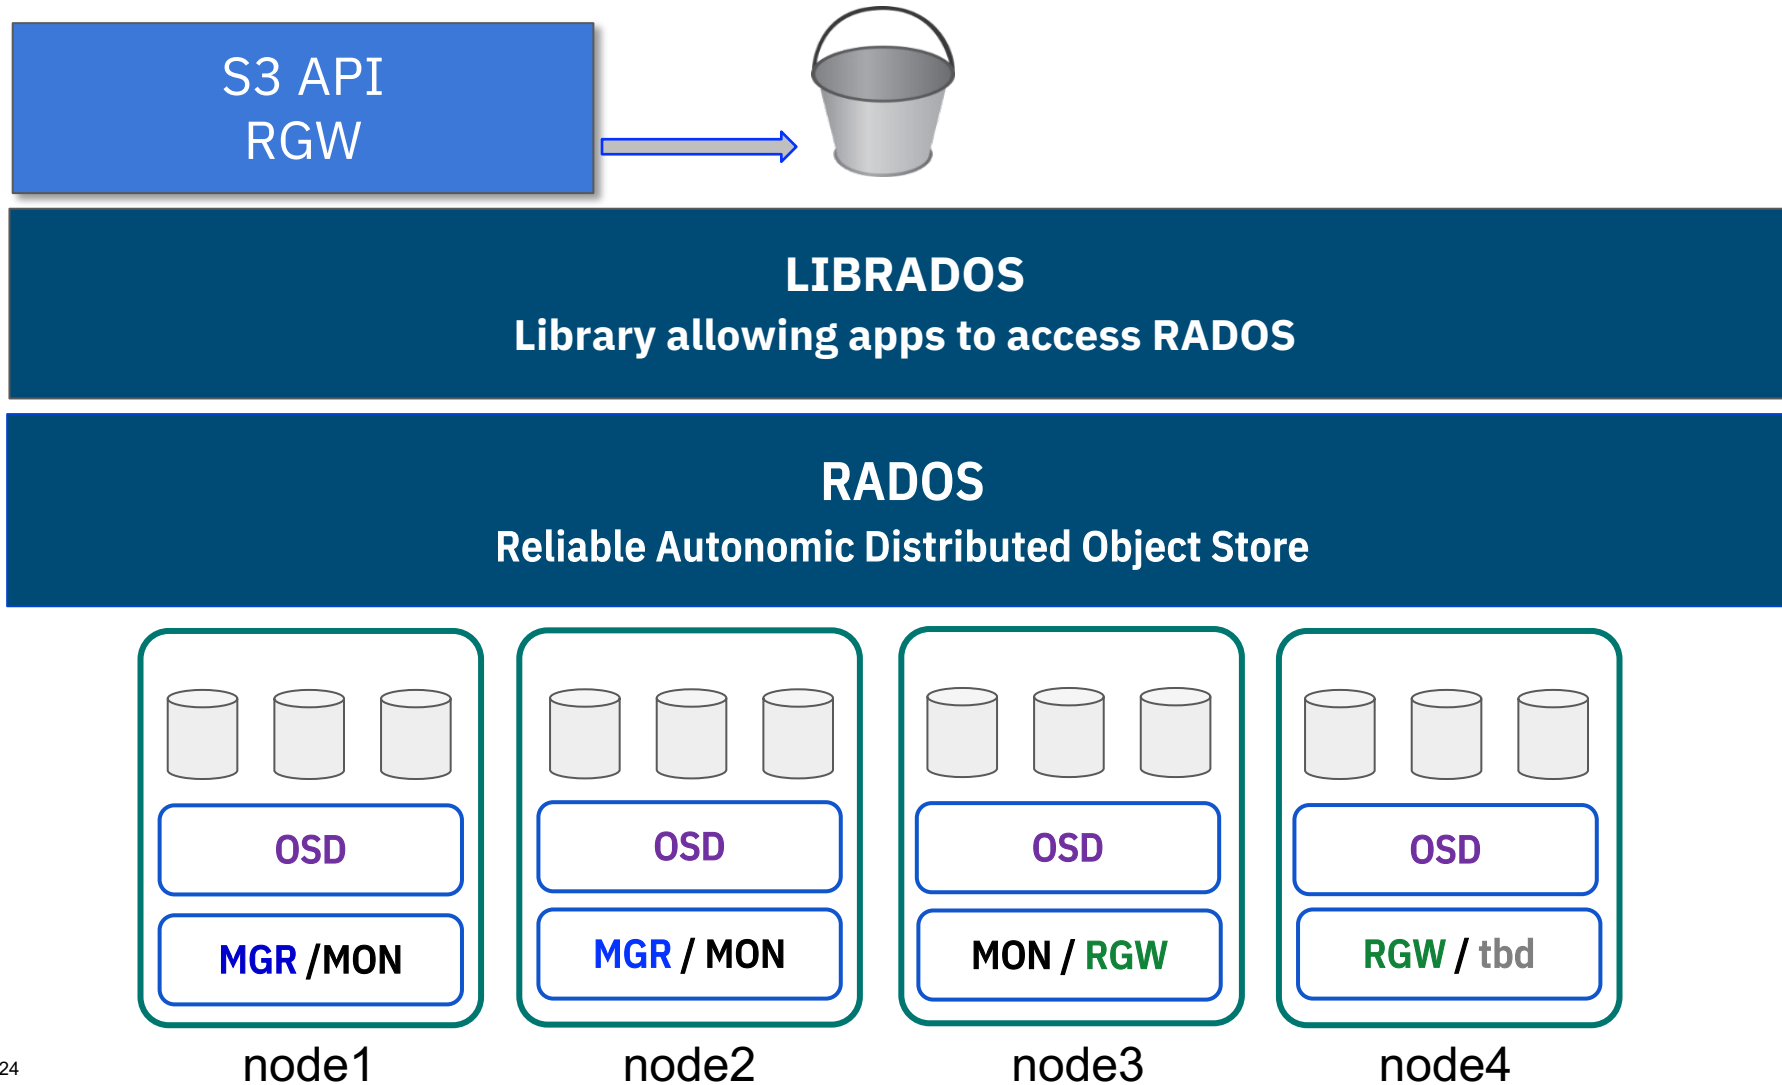
IBM Storage Ceph data access

Ceph Data Services (RBD, CephFS, RGW) for mainstream application use cases



IBM Storage Ceph – Overview of the RADOS Gateway

RADOS Gateway (RGW) layout

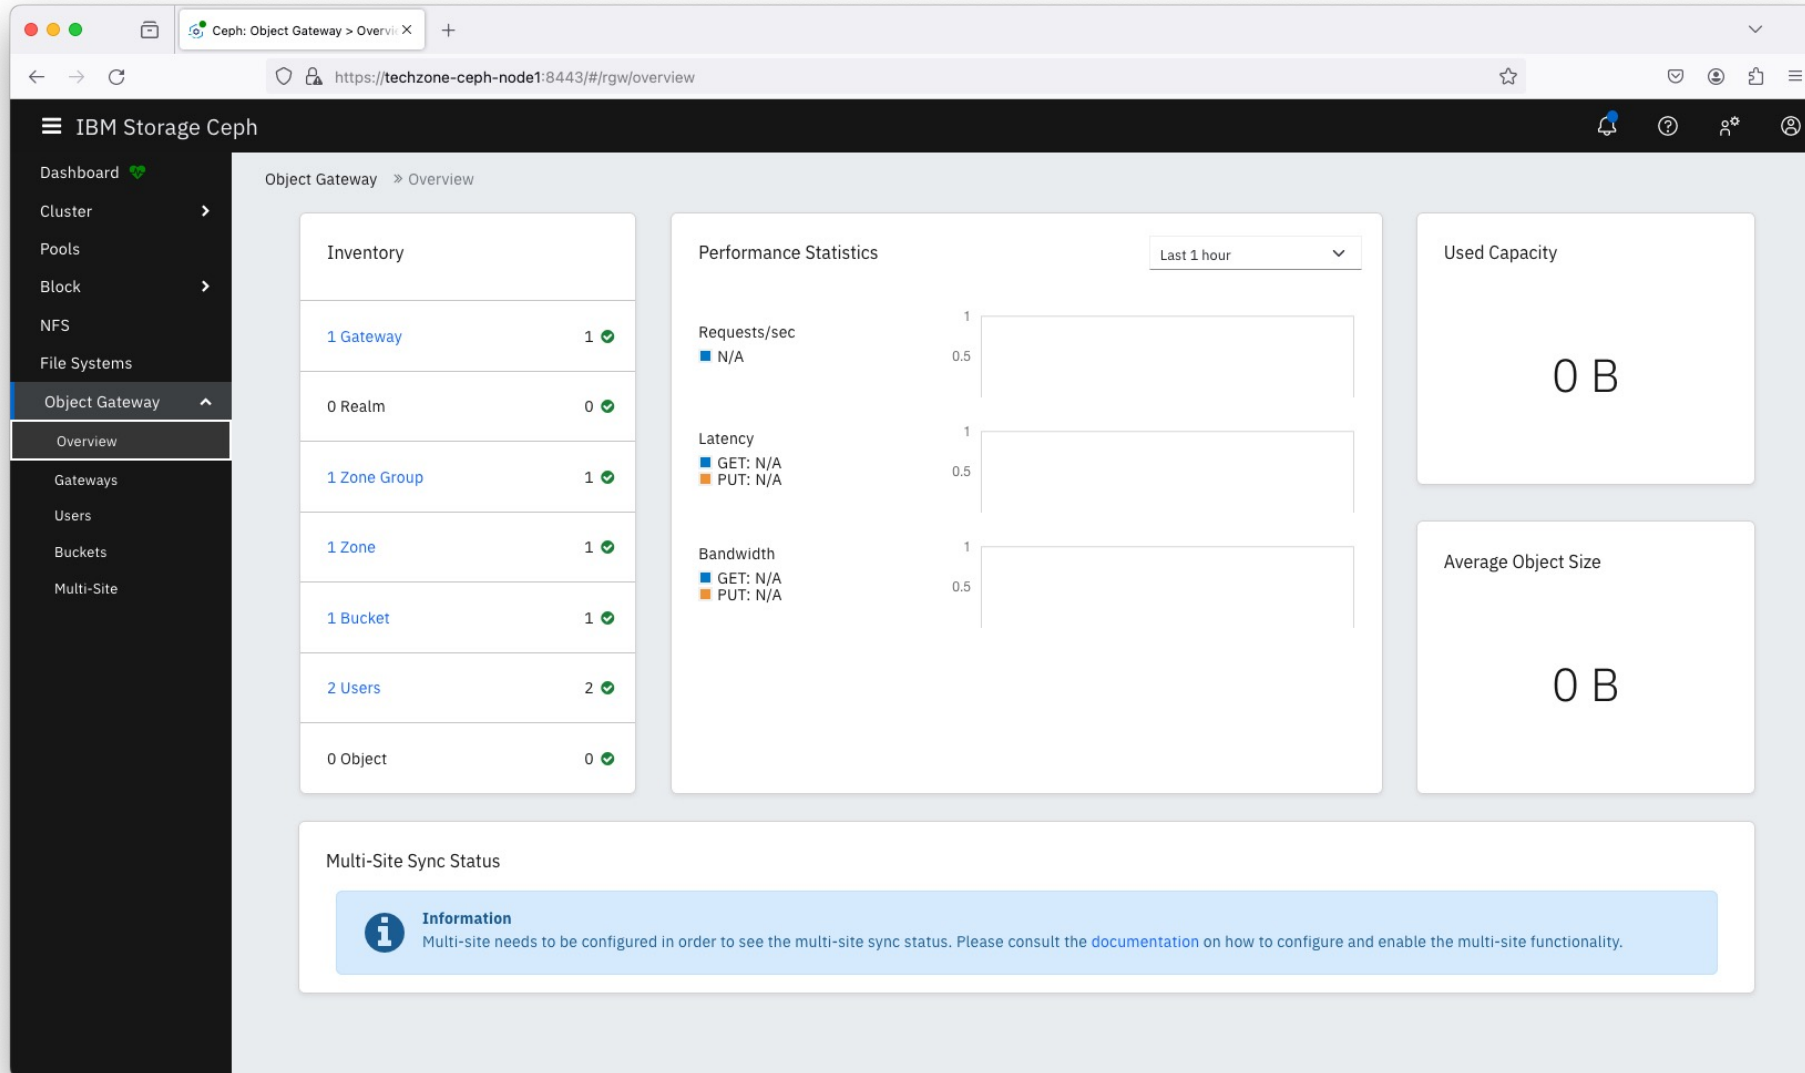


IBM Storage Ceph Administration

- Dashboard UI
- Command line
- Service specification file
- Performance visualization
- Multi tenancy

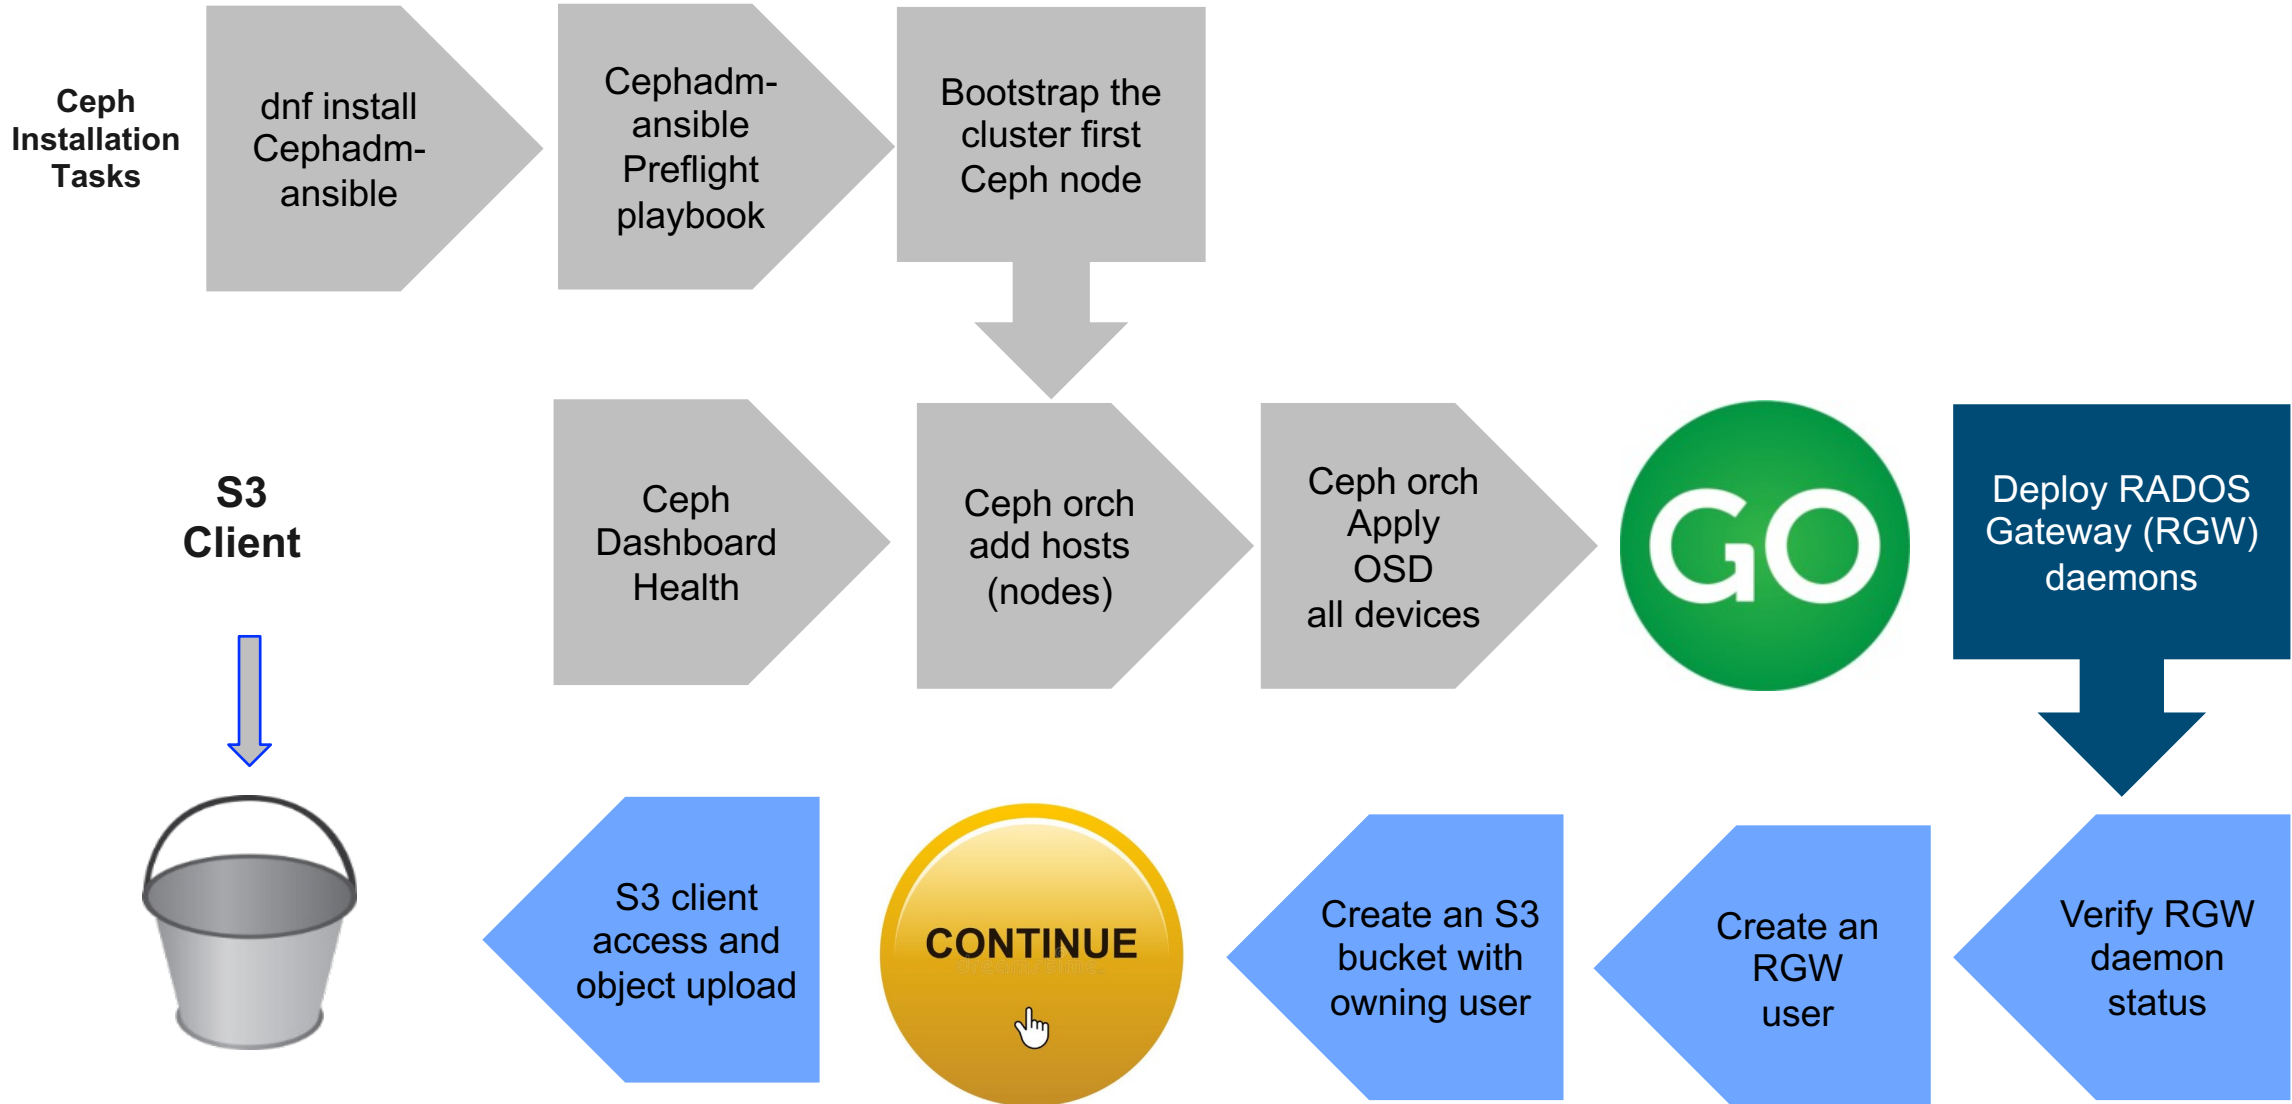


Dashboard UI -> simple control



IBM Storage Ceph RGW configuration tasks

LIVE



Ceph command line -> discrete control

```
[root@ceph-poc-node1 ~]# radosgw-admin user create --uid=jerrod \  
--email="jcarr@us.ibm.com" \  
--display-name="Carr, Jerrod"  
  
. . . output omitted . . .  
  
keys": [  
  {  
    "user": "jerod",  
    "access_key": "158FPSP3BTYBWTJXZ65N",  
    "secret_key": "XXjdDUsAaBGE5UJpCqqw560BOIrQ2AwmeKyAovgI"  
  }  
]  
  
. . . output omitted . . .
```

Service specification file – automated control

```
[root@ceph-poc-node1 ~]# cat radosgw.yml

service_type: rgw
service_id: s3service
service_name: rgw.s3service
placement:
  count: 2
  hosts:
    - ceph-poc-node3
    - ceph-poc-node4
spec:
  rgw_frontend_port: 80

[root@ceph-poc-node1 ~]# ceph orch apply -i radosgw.yml
```

<https://www.ibm.com/docs/en/storage-ceph/7?topic=deployment-deploying-ceph-object-gateway-using-service-specification>

Simple RGW configuration in the Dashboard or command line

Create Service
✕

Type * ✓

Id * 🔒

Unmanaged

Placement ▾

Hosts ✎ There are no hosts.

Count ?

Port

SSL

Cancel
Create Service

```
[root@ceph-poc-node1 ~]# ceph orch apply rgw s3service
```

```
[root@ceph-poc-node1 ~]# ceph orch ls
```

NAME	PORTS	RUNNING
REFRESHED	AGE	PLACEMENT
. . . output omitted . . .		
rgw.s3service	?:80	2/2
103s	count:2	83s ago

<https://www.ibm.com/docs/en/storage-ceph/7?topic=ceph-object-gateway>

Controlled RGW configuration in the Dashboard or command line

Create Service
✕

Type * ✓

Id * ✓

Unmanaged

Placement

Hosts ✕ ✕

Count ? ✓

Port ✓

SSL

Cancel
Create Service

```
[root@ceph-poc-node1 ~]# ceph orch apply rgw
s3service \
--placement="2 ceph-poc-node3 ceph-poc-node4" --
port=8080
```

```
[root@ceph-poc-node1 ~]# ceph orch ls -service-
type rgw
```

NAME	REFRESHED	AGE	PLACEMENT	PORTS	RUNNING
rgw.s3service	?:8080	56s ago	ceph-poc-node3;ceph-poc-node4;count:2	2/2	72s

<https://www.ibm.com/docs/en/storage-ceph/7?topic=ceph-object-gateway>

Controlled RGW configuration in the Dashboard or command line

Create Service
✕

Type * ✓

Id * ✓

Unmanaged

Placement ✓

Hosts ceph-poc-node3 ✕ ceph-poc-node4 ✕

Count ⓘ ✓

Port ✓

SSL

Cancel
Create Service

```
[root@ceph-poc-node1 ~]# ceph orch apply rgw
s3service \
--placement="4 ceph-poc-node3 ceph-poc-node4" --
port=80
```

```
[root@ceph-poc-node1 ~]# ceph orch ls -service-
type rgw
```

NAME	REFRESHED	AGE	PLACEMENT	PORTS	RUNNING
rgw.s3service	?	80	4/4	56s ago	72s
ceph-poc-node3;ceph-poc-node4;count:4					

<https://www.ibm.com/docs/en/storage-ceph/7?topic=ceph-object-gateway>

IBM Storage Ceph Dashboard for object storage services

The screenshot displays the IBM Storage Ceph dashboard for the Object Gateway Overview. The interface includes a left-hand navigation menu, a main content area with several widgets, and a bottom section for Multi-Site Sync Status.

Navigation Menu:

- Dashboard
- Cluster
- Pools
- Block
- NFS
- File Systems
- Object Gateway**
 - Overview
 - Gateways
 - Users
 - Buckets
 - Multi-Site

Object Gateway Overview Widgets:

- Inventory:** A table listing components and their counts with status indicators.

Component	Count	Status
1 Gateway	1	✓
0 Realm	0	✓
1 Zone Group	1	✓
1 Zone	1	✓
1 Bucket	1	✓
3 Users	3	✓
0 Object	0	✓
- Performance Statistics:** Three line charts showing Requests/sec, Latency, and Bandwidth for the last 1 hour. All metrics are currently N/A.
- Used Capacity:** A card displaying 0 B.
- Average Object Size:** A card displaying 0 B.

Multi-Site Sync Status:

Information
Multi-site needs to be configured in order to see the multi-site sync status. Please consult the [documentation](#) on how to configure and enable the multi-site functionality.

What do I get with a simple RGW configuration?

Realm: { }

Zonegroup: {default}

Zone: {default}

Object gateway(s)



OSD Nodes

Placement Targets



Replica x3

Create an RGW storage account in the Dashboard or command line

Create User

User ID * ✓

Show Tenant

Full name * ✓

Email address ✓

Max. buckets ▾

▾

Suspended

Enabled

Bucket quota

Enabled

```
[root@ceph-poc-node1 ~]# radosgw-admin user create
--uid=john \
--email="jshubeck@us.ibm.com" \
--display-name="Shubeck, John"

... output omitted ...

keys": [
  {
    "user": "john",
    "access_key": "158FPSP3BTYBWTJXZ65N",
    "secret_key":
"XXjdDUaAaBGE5UJpCqqW560BOIrQ2AwmeKyAovgI"
  }
]

... output omitted ...
```

<https://www.ibm.com/docs/en/storage-ceph/7?topic=ceph-object-gateway>

Configuring HTTPS in the object gateway

Create Service
✕

Type * ✓

Id * ✓

Unmanaged

Placement ✓

Label ✓

Count ? ✓

Port ✓

SSL

Certificate ? ✓

Cancel
Create Service

```
[root@ceph-poc-node4 ~]# openssl req -x509 -newkey  
rsa:4096 -keyout key.pem -out cert.pem -sha256 -  
days 365 -nodes
```

... Output omitted

```
[root@ceph-poc-node4 ~]# copy/paste from cert.pem  
and key.pem into the "Certificate" text box in the  
Dashboard UI
```

<https://www.ibm.com/docs/en/storage-ceph/7?topic=ceph-object-gateway>

Configuring HTTPS in the object gateway using a specification file

```
[root@ceph-poc-node4 ~]# cat <<EOF >> /root/rgw-config.yml
service_type: rgw
service_id: s3ssl
service_name: rgw.s3ssl
placement:
  hosts:
    - ceph-poc-node4
spec:
  ssl: true
  rgw_frontend_ssl_certificate: |
  rgw_realm: default
  rgw_zone: default
  rgw_zonegroup: default
  ssl: true
  rgw_frontend_port: 443
  rgw_frontend_type: beast
  rgw_frontend_ssl_certificate: |
    -----BEGIN CERTIFICATE-----
$( cat /root/cert.pem | grep -v CERTIFICATE | awk '{ $1="    "$1 }' )
    -----END CERTIFICATE-----
    -----BEGIN RSA PRIVATE KEY-----
$( cat /root/key.pem | grep -v PRIVATE | awk '{ $1="    "$1 }' )
    -----END RSA PRIVATE KEY-----
EOF
[root@ceph-poc-node4 ~]# ceph orch apply -i /root/rgw-config.yml
```


RADOS Gateway Multi Site



RADOS Gateway Multi Site Configuration (SAMPLE)

Realm: atg

Zone group: us

Master Zone: herndon

Object gateway

Secondary Zone: dallas

Object gateway



Monitors



Monitors



OSDs



OSDs

IBM Storage Ceph Dashboard for multi site topology view

The screenshot displays the IBM Storage Ceph dashboard interface. The browser address bar shows the URL `https://atg-ws-ceph-11:8443/#/rgw/multisite`. The left sidebar contains a navigation menu with the following items: Dashboard, Cluster, Pools, Block, NFS, File Systems, Object Gateway (expanded), Overview, Gateways, Users, Buckets, and Multi-Site (highlighted with a green border). The main content area is titled "Object Gateway" and includes an information banner stating: "In order to access the import/export feature, the rgw module must be enabled [Enable the Object Gateway Module](#)". Below this banner are buttons for "+ Create Realm", "Import", and "Export". The "Topology Viewer" section shows a hierarchical tree structure of the multi-site topology:

- atg (default)
 - us (default, master)
 - herndon (default, master)
 - dallas (secondary-zone)

RADOS Gateway storage class and placement target



RADOS Gateway - Storage class and placement target

Realm: { }

Zonegroup: {default}

Zone: {default}

Object gateways




OSD Nodes

Placement Targets



Replica x3



Erasure Coding 4+2

RADOS Gateway storage pools – Ceph baseline pool



The screenshot shows the IBM Storage Ceph web interface. The left sidebar contains navigation options: Dashboard, Cluster, Pools (selected), Block, NFS, File Systems, and Object Gateway. The main content area is titled 'Pools' and has two tabs: 'Pools List' (active) and 'Overall Performance'. A '+ Create' button is visible. Below it is a table with columns: Name, Data Protection, Applications, PG Status, Usage, Read bytes, Write bytes, Read ops, and Write ops. The table contains one entry: 'device_health_metrics' with 'Data Protection' set to 'replica: x3', 'Applications' set to 'mgr_devicehealth', 'PG Status' set to '1 active+clean', and 'Usage' set to '0%'. The bottom of the table shows '0 selected / 1 total'.

Name	Data Protection	Applications	PG Status	Usage	Read bytes	Write bytes	Read ops	Write ops
> device_health_metrics	replica: x3	mgr_devicehealth	1 active+clean	0%			0/s	0/s

RADOS Gateway default storage pools – First RGW service instance

The screenshot displays the IBM Storage Ceph web interface. The left sidebar contains a navigation menu with the following items: Dashboard, Cluster, Hosts, Physical Disks, Monitors, Services, OSDs, Configuration, CRUSH map, Manager Modules, Logs, Monitoring, Pools, Block, NFS, File Systems, and Object Gateway. The main content area is titled 'Pools' and shows a 'Pools List' tab. A '+ Create' button is visible at the top left of the table. The table has the following columns: Name, Data Protection, Applications, PG Status, Usage, Read bytes, Write bytes, Read ops, and Write ops. There are five rows of data representing different storage pools.

Name	Data Protection	Applications	PG Status	Usage	Read bytes	Write bytes	Read ops	Write ops
> device_health_metrics	replica: x3	mgr_devicehealth	1 active+clean	0%	0/s	0/s	0/s	0/s
> default.rgw.meta	replica: x3	rgw	32 active+clean	0%	0/s	0/s	0/s	0/s
> default.rgw.log	replica: x3	rgw	32 active+clean	0%	0/s	0/s	0/s	0/s
> default.rgw.control	replica: x3	rgw	32 active+clean	0%	0/s	0/s	0/s	0/s
> .rgw.root	replica: x3	rgw	32 active+clean	0%	0/s	0/s	0/s	0/s

0 selected / 5 total

RADOS Gateway default storage pools – First bucket created

The screenshot shows the IBM Storage Ceph management interface. The left sidebar contains navigation options: Dashboard, Cluster, Pools, Block, NFS, File Systems, Object Gateway, Daemons, Users, and Buckets. The main content area is titled 'Pools' and has two tabs: 'Pools List' (selected) and 'Overall Performance'. A '+ Create' button is visible. A search bar contains the text 'rgw'. Below the search bar is a table of storage pools. The pool 'default.rgw.buckets.index' is highlighted with a red rectangular box. The table columns are: Name, Data Protection, Applications, PG Status, Usage, Read bytes, Write bytes, Read ops, and Write ops. The 'default.rgw.buckets.index' row shows 'default.rgw.buckets.index' as the name, 'replica: x3' for data protection, 'rgw' for applications, '32 active+clean' for PG status, '0%' for usage, and small line graphs for read/write bytes and '0/s' for read/write ops.

Name	Data Protection	Applications	PG Status	Usage	Read bytes	Write bytes	Read ops	Write ops
> default.rgw.meta	replica: x3	rgw	32 active+clean	0%			0/s	0/s
> default.rgw.log	replica: x3	rgw	32 active+clean	0%			0/s	0/s
> default.rgw.control	replica: x3	rgw	32 active+clean	0%			0/s	0/s
> default.rgw.buckets.index	replica: x3	rgw	32 active+clean	0%			0/s	0/s
> .rgw.root	replica: x3	rgw	32 active+clean	0%			0/s	0/s

0 selected / 5 found / 6 total

RADOS Gateway default storage pools – First object stored

The screenshot shows the IBM Storage Ceph interface for the 'Pools' section. The 'Pools List' tab is active, displaying a table of storage pools. The search bar contains 'rgw'. The table lists several pools, with 'default.rgw.buckets.index' and 'default.rgw.buckets.data' highlighted. The status for all pools is 'active+clean' with 32 PGs and 0% usage.

Name	Data Protection	Applications	PG Status	Usage	Read bytes	Write bytes	Read ops	Write ops
> default.rgw.meta	replica: x3	rgw	32 active+clean	0%	0/s	0/s	0/s	0/s
> default.rgw.log	replica: x3	rgw	32 active+clean	0%	0/s	0/s	0/s	0/s
> default.rgw.control	replica: x3	rgw	32 active+clean	0%	0/s	0/s	0/s	0/s
> default.rgw.buckets.index	replica: x3	rgw	32 active+clean	0%	0/s	0/s	0/s	0/s
> default.rgw.buckets.data	replica: x3	rgw	32 active+clean	0%	0/s	0/s	0/s	0/s
> .rgw.root	replica: x3	rgw	32 active+clean	0%	0/s	0/s	0/s	0/s

0 selected / 6 found / 7 total

Storage Class and Placement Target

```
[root@ceph-poc-node1 ~]# radosgw-admin zonegroup get
"placement_targets": [
  {
    "name": "default-placement",
    "tags": [ ],
    "storage_classes": [
      "STANDARD"
    ]
  }
]
```

Storage Class and Placement Target

```
[root@ceph-poc-node1 ~]# radosgw-admin zone get
```

```
"placement_pools": [  
  {  
    "key": "default-placement",  
    "value": {  
      "index_pool": "default.rgw.buckets.index",  
      "storage_classes": {  
        "STANDARD": {  
          "data_pool": "default.rgw.buckets.data"  
        }  
      },  
      "data_extra_pool": "default.rgw.buckets.non-ec",  
      "index_type": 0  
    }  
  }  
]
```

?

New pool for custom placement bucket data

Create Pool

Name * ✓

Pool type * ✓

PG Autoscale

Flags EC Overwrites

Applications

CRUSH

Erasure code profile ⓘ + 🗑️

Crush ruleset A new crush ruleset will be implicitly created.

Quotas

Max bytes ⓘ

Max objects ⓘ

```
[root@ceph-poc-node1 ~]# ceph osd pool create \
default.rgw.glacier.data erasure ec-2-plus-1
```

```
[root@ceph-poc-node1 ~]# ceph osd pool application
enable \
default.rgw.glacier.data rgw
```

New pool for custom placement bucket index

Create Pool

Name * ✓

Pool type * ✓

PG Autoscale

Replicated size *

Applications [+](#) No applications added

CRUSH

Crush ruleset ⓘ + 🗑️

Compression

Mode

Quotas

Max bytes ⓘ

Max objects ⓘ

```
[root@ceph-poc-node1 ~]# ceph osd pool create \
default.rgw.glacier.buckets.index replicated
```

```
[root@ceph-poc-node1 ~]# ceph osd pool application
enable default.rgw.glacier.buckets.index rgw
```

New pool for custom placement bucket Extra Pool (Multi-part upload)

Create Pool

Name * ✓

Pool type * ✓

PG Autoscale

Replicated size *

Applications

CRUSH

Crush ruleset

Compression

Mode

Quotas

Max bytes ?

Max objects ?

```
[root@ceph-poc-node1 ~]# ceph osd pool create \
default.rgw.glacier.buckets.non-ec replicated
```

```
[root@ceph-poc-node1 ~]# ceph osd pool application \
enable default.rgw.glacier.buckets.non-ec rgw
```

RADOS Gateway custom storage pools

The screenshot shows the IBM Storage Ceph management interface. The left sidebar contains navigation options: Dashboard, Cluster, Pools, Block, Images, Mirroring, iSCSI, NFS, File Systems, and Object Gateway. The main content area is titled 'Pools' and has two tabs: 'Pools List' (selected) and 'Overall Performance'. A search bar at the top right contains the text 'glacier'. Below the search bar is a table of storage pools. Three rows are visible, each with a colored border: a green border for 'default.rgw.glacier.buckets.non-ec', a red border for 'default.rgw.glacier.buckets.index', and a blue border for 'default.rgw.glacier.buckets.data'. The table columns include Name, Data Protection, Applications, PG Status, Usage, Read bytes, Write bytes, Read ops, and Write ops. The status for all pools is 'active+clean'.

Name	Data Protection	Applications	PG Status	Usage	Read bytes	Write bytes	Read ops	Write ops
default.rgw.glacier.buckets.non-ec	replica: x3	rgw	1 active+clean	0%			0 /s	0 /s
default.rgw.glacier.buckets.index	replica: x3	rgw	32 active+clean	0%			0 /s	0 /s
default.rgw.glacier.buckets.data	EC: 2+1	rgw	32 active+clean	0%			0 /s	0 /s

0 selected / 3 found / 11 total

Add the a new Placement ID and Storage Class

```
[root@ceph-poc-node1 ~]# radosgw-admin zonegroup placement add \  
--rgw-zonegroup default \  
--placement-id glacier \  
--storage-class GLACIER
```

```
[root@ceph-poc-node1 ~]# radosgw-admin zone placement add \  
--rgw-zone default \  
--placement-id glacier \  
--storage-class GLACIER \  
--data-pool default.rgw.glacier.data \  
--index-pool default.rgw.glacier.index \  
--data-extra-pool default.rgw.glacier.non-ec
```

1 <https://www.ibm.com/docs/en/storage-ceph/7?topic=administration-creating-storage-policies>

? <https://docs.ceph.com/en/latest/radosgw/placement/>

? https://access.redhat.com/documentation/en-us/red_hat_ceph_storage/5/html/object_gateway_guide/administration#creating-storage-policies-rgw

Storage Class and Placement Target

```
[root@ceph-poc-node1 ~]# radosgw-admin zonegroup placement add \  
--rgw-zonegroup default \  
--placement-id glacier \  
--storage-class GLACIER
```

```
[root@ceph-poc-node1 ~]# radosgw-admin zone placement add \  
--rgw-zone default \  
--placement-id glacier \  
--storage-class GLACIER \  
--index-pool default.rgw.glacier.index \  
--data-extra-pool default.rgw.glacier.non-ec \  
--data-pool default.rgw.glacier.data --compression lz4
```

1 <https://www.ibm.com/docs/en/storage-ceph/7?topic=administration-creating-storage-policies>

? <https://docs.ceph.com/en/latest/radosgw/placement/>

? https://access.redhat.com/documentation/en-us/red_hat_ceph_storage/5/html/object_gateway_guide/administration#creating-storage-policies-rgw

RADOS Gateway – Multi tenancy

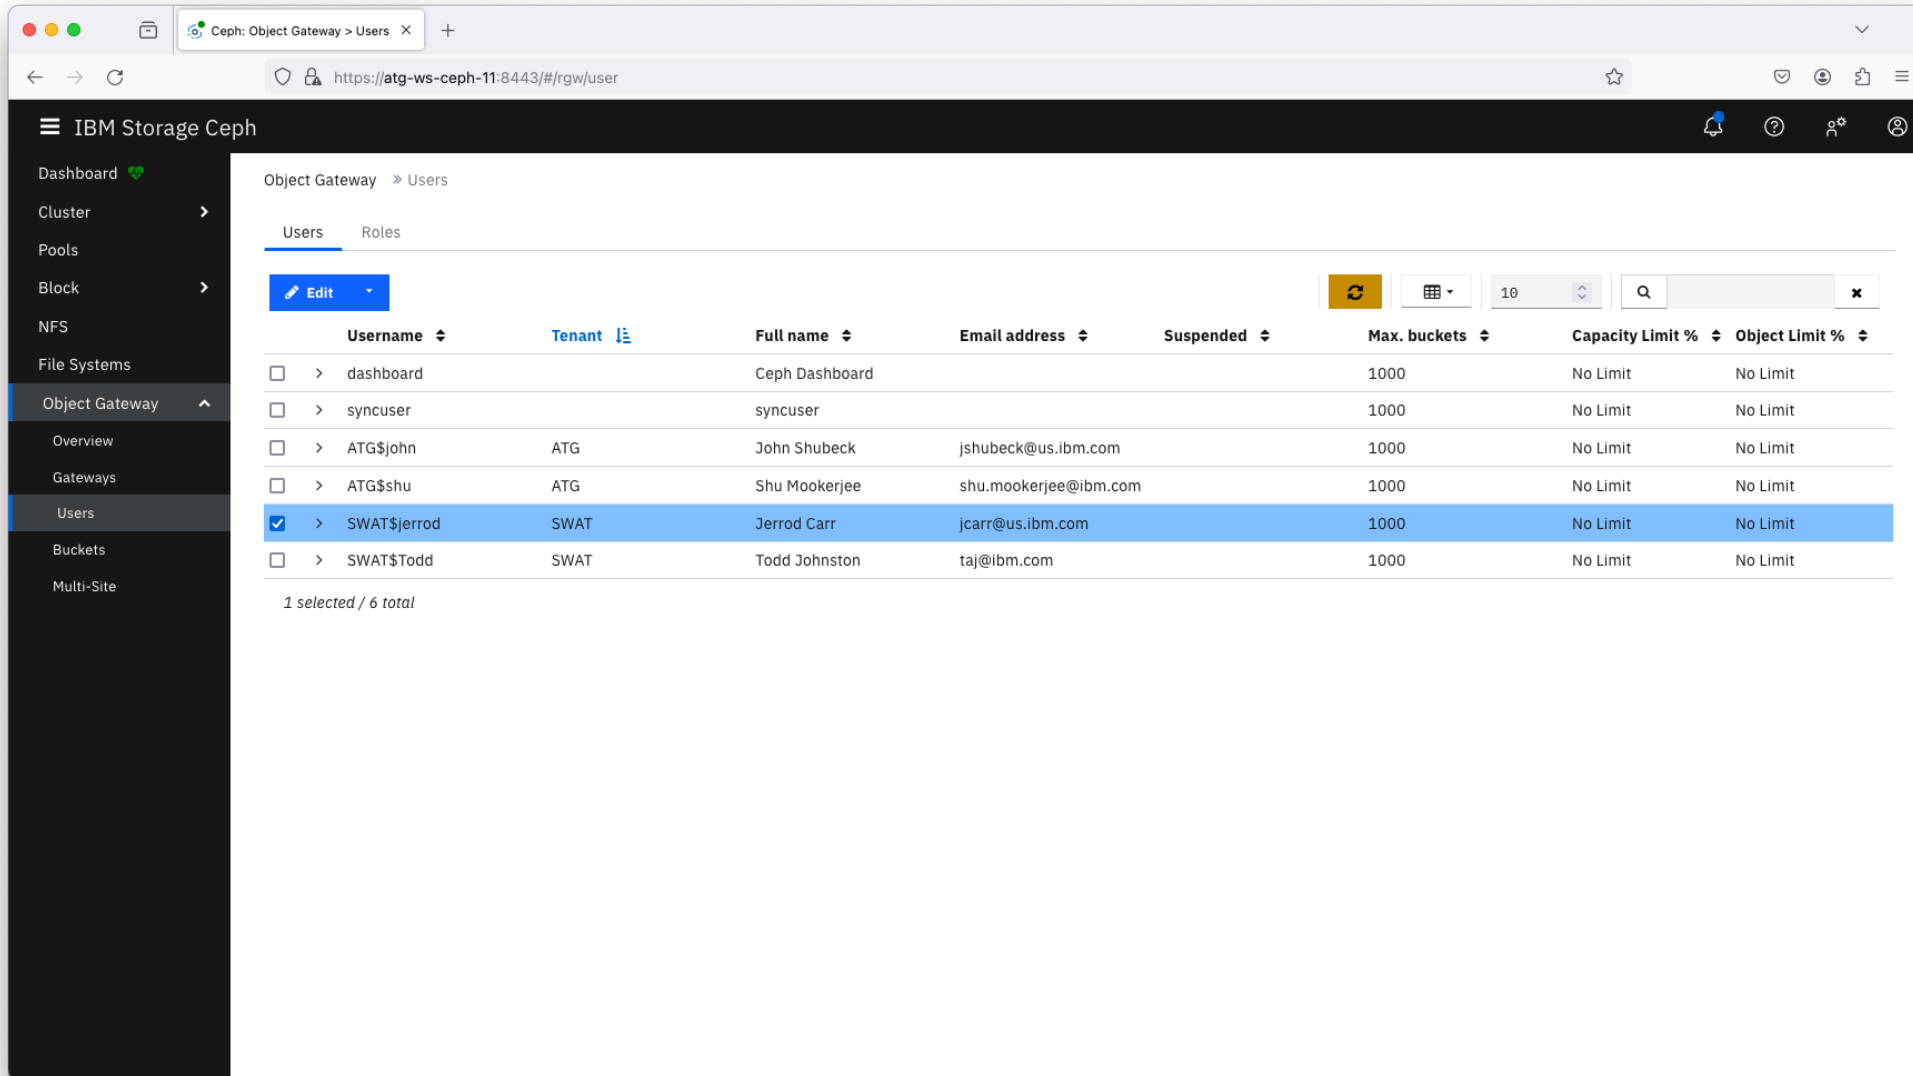


The screenshot shows the IBM Storage Ceph Object Gateway Users management interface. The browser address bar shows the URL `https://192.168.65.102:8443/#/rgw/user`. The left sidebar contains navigation options: Dashboard, Cluster, Pools, Block, NFS, File Systems, Object Gateway (selected), Daemons, Users, and Buckets. The main content area is titled "Object Gateway » Users" and features a "+ Create" button, a refresh icon, a table view selector, a page size dropdown set to 10, and a search bar. A table lists three users, with the first row highlighted by a red box:

	Username	Tenant	Full name	Email address	Suspended	Max. buckets	Capacity Limit %	Object Limit %
<input type="checkbox"/>	> ATG\$shawn	ATG	Houston, Shawn	shawn@us.ibm.com		1000	No Limit	No Limit
<input type="checkbox"/>	> dashboard		Ceph Dashboard			1000	No Limit	No Limit
<input type="checkbox"/>	> john		Shubeck, John	jshubeck@us.ibm.com		1000	No Limit	No Limit

Below the table, it indicates "0 selected / 3 total".

Multi Tenancy



The screenshot shows the IBM Storage Ceph Object Gateway Users management interface. The browser address bar indicates the URL is `https://atg-ws-ceph-11-8443/#/rgw/user`. The left sidebar contains navigation options: Dashboard, Cluster, Pools, Block, NFS, File Systems, Object Gateway (selected), Overview, Gateways, Users, Buckets, and Multi-Site. The main content area is titled "Object Gateway » Users" and has tabs for "Users" and "Roles". A table lists the users with columns for Username, Tenant, Full name, Email address, Suspended, Max. buckets, Capacity Limit %, and Object Limit %. The user "SWAT\$jerrod" is selected. Below the table, it shows "1 selected / 6 total".

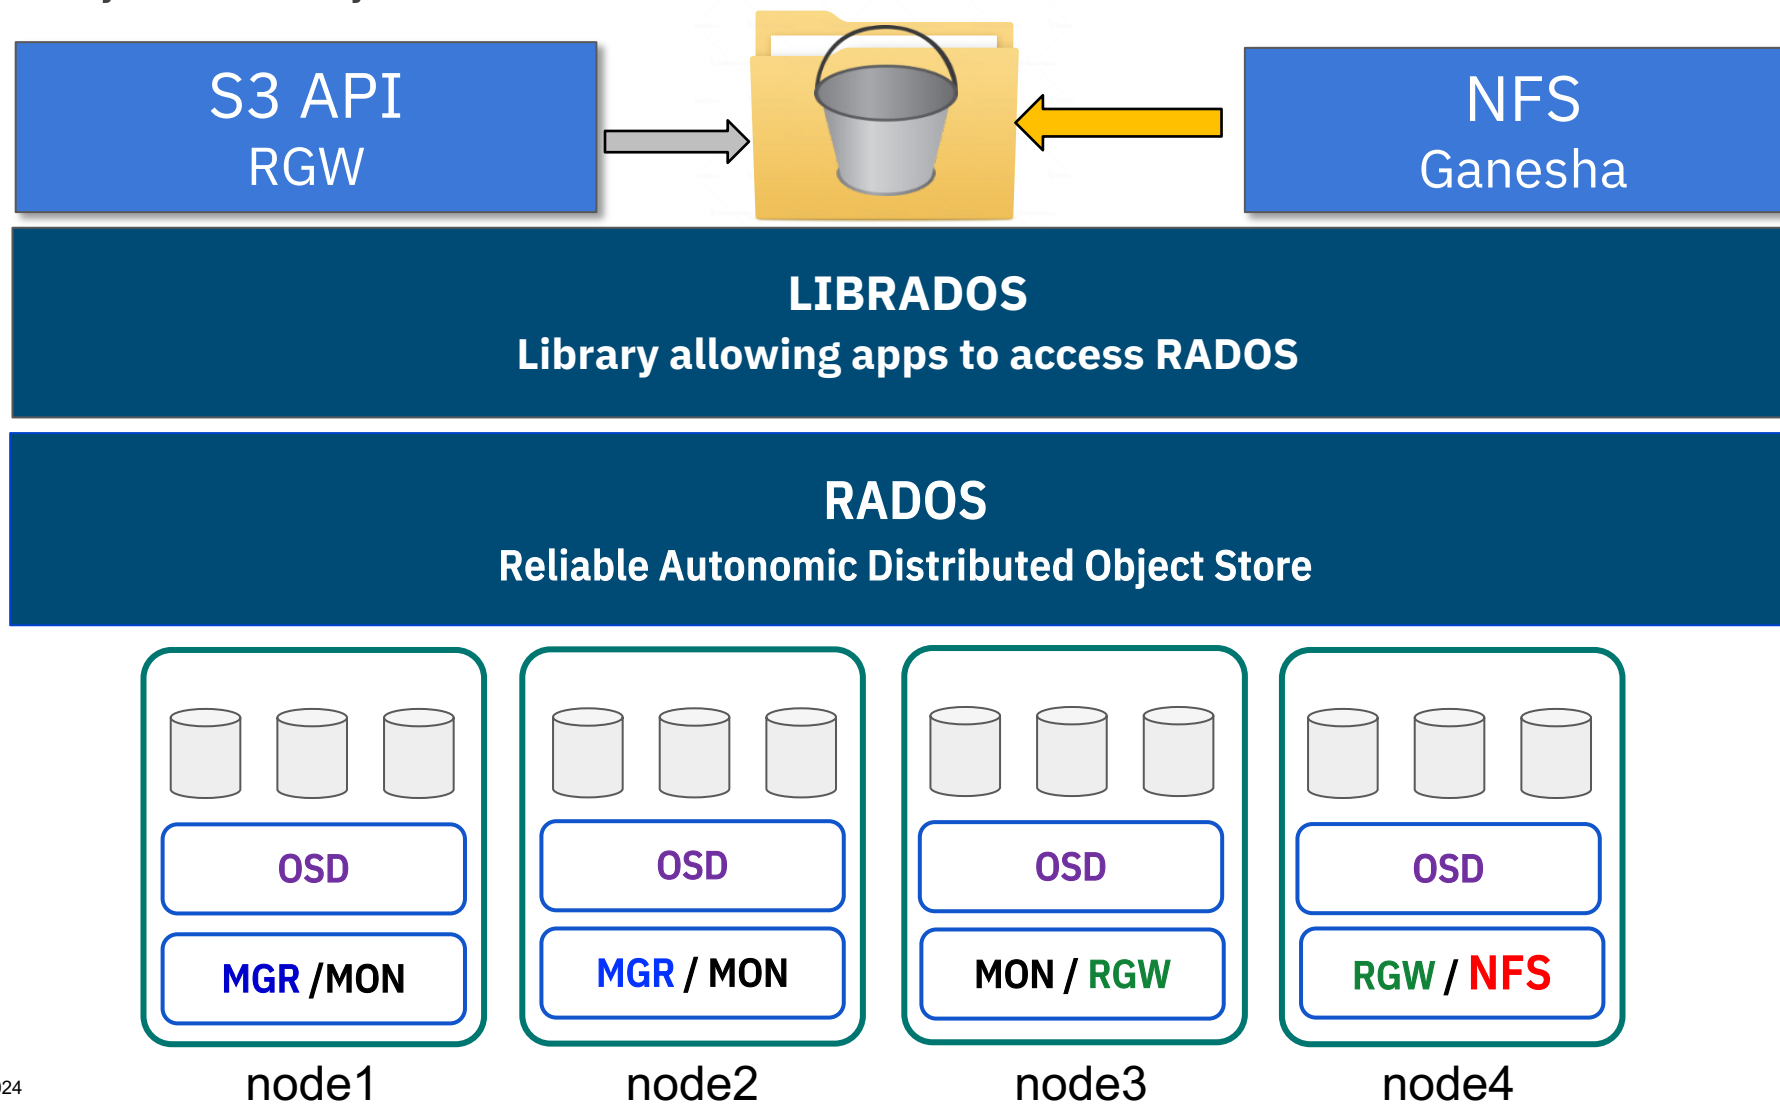
	Username	Tenant	Full name	Email address	Suspended	Max. buckets	Capacity Limit %	Object Limit %
<input type="checkbox"/>	> dashboard		Ceph Dashboard			1000	No Limit	No Limit
<input type="checkbox"/>	> syncuser		syncuser			1000	No Limit	No Limit
<input type="checkbox"/>	> ATG\$john	ATG	John Shubeck	jshubeck@us.ibm.com		1000	No Limit	No Limit
<input type="checkbox"/>	> ATG\$shu	ATG	Shu Mookerjee	shu.mookerjee@ibm.com		1000	No Limit	No Limit
<input checked="" type="checkbox"/>	> SWAT\$jerrod	SWAT	Jerrod Carr	jcarr@us.ibm.com		1000	No Limit	No Limit
<input type="checkbox"/>	> SWAT\$todd	SWAT	Todd Johnston	taj@ibm.com		1000	No Limit	No Limit

Multi protocol access

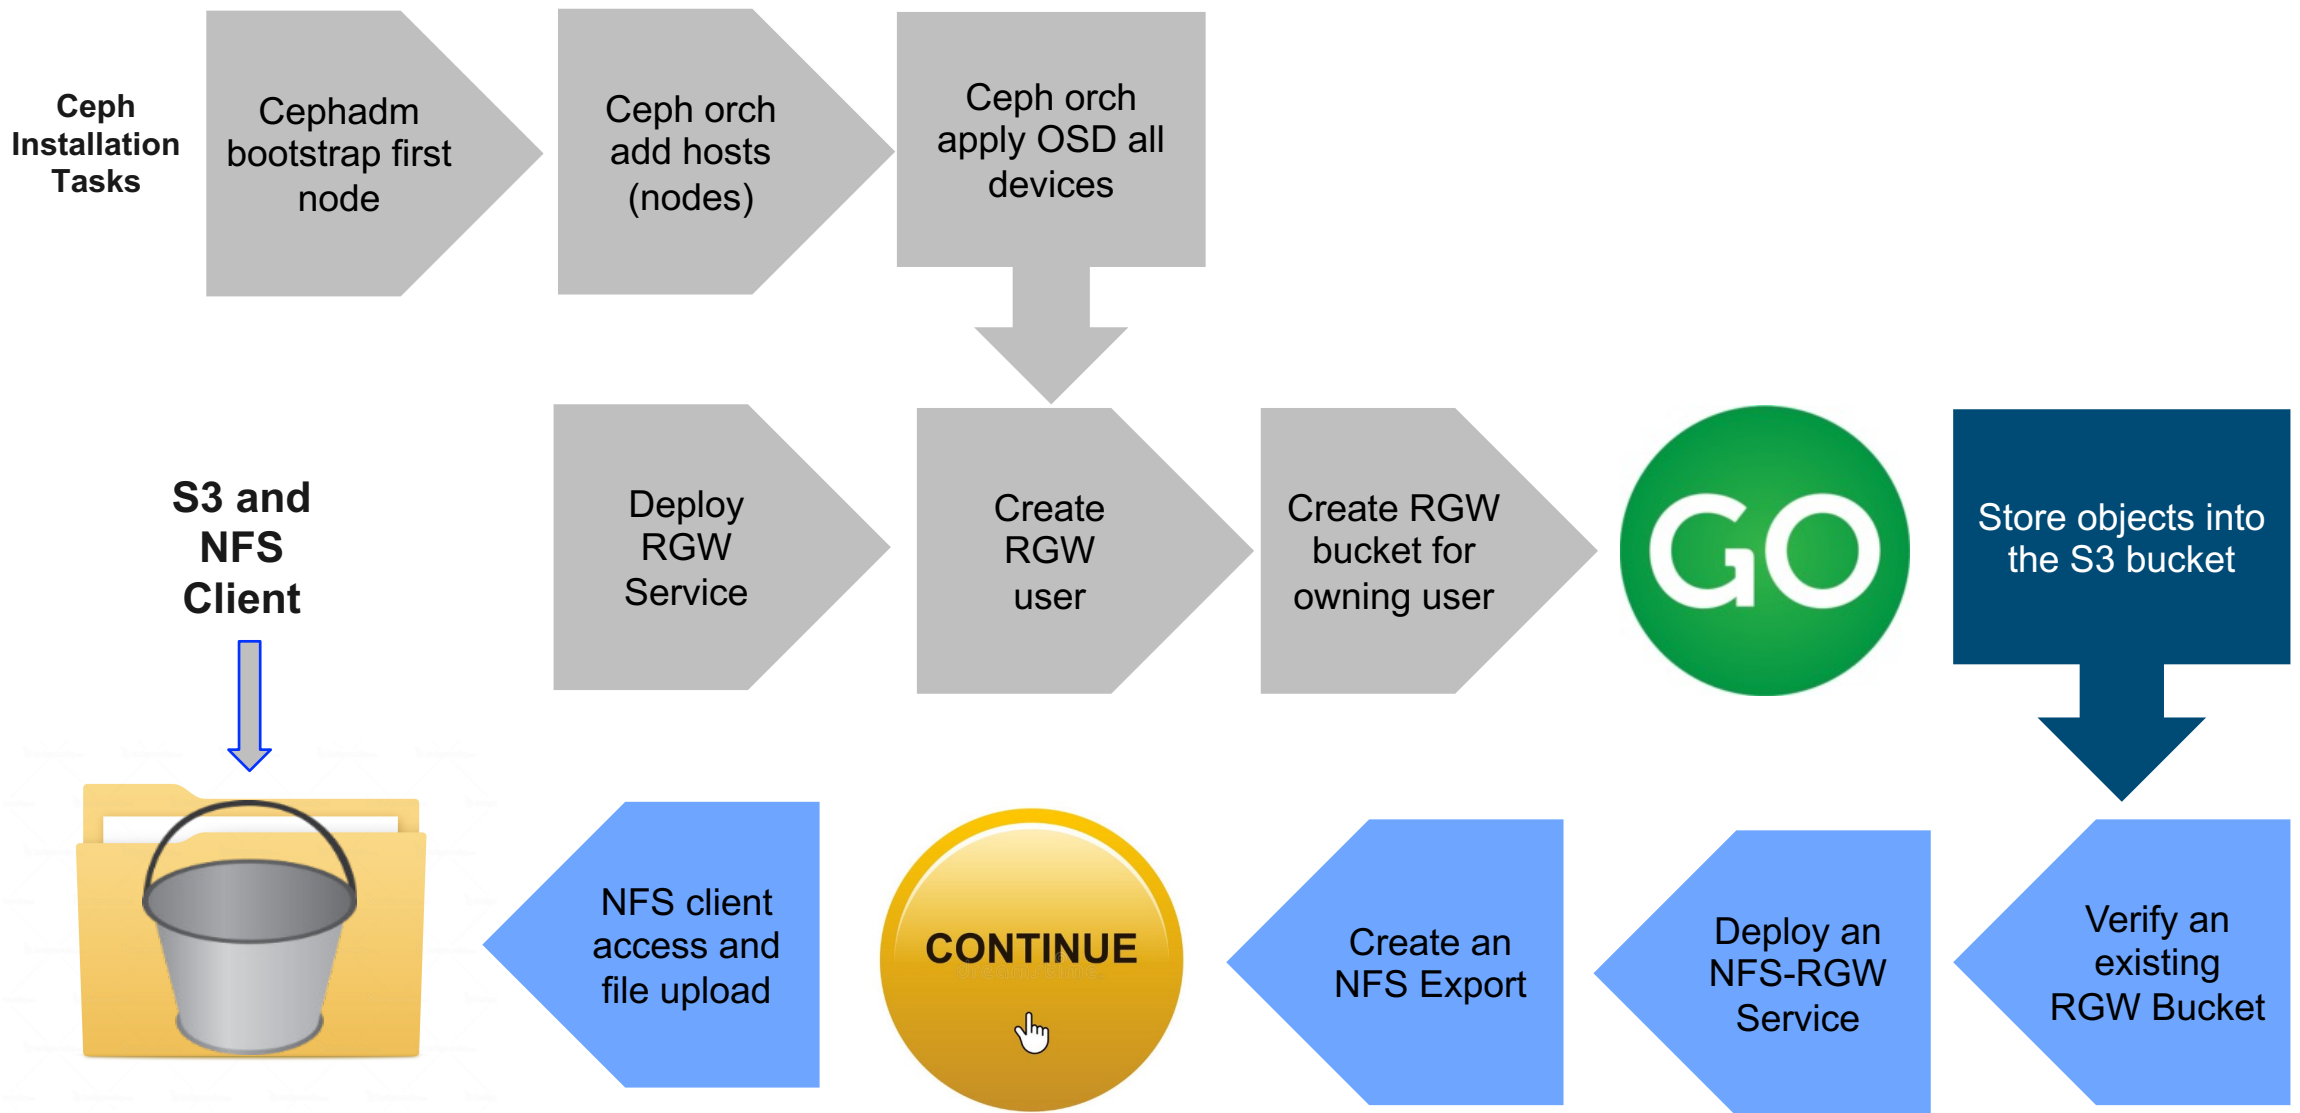


IBM Storage Ceph – Overview of the RADOS Gateway

RADOS Gateway (RGW) layout



IBM Storage Ceph multiprotocol configuration tasks



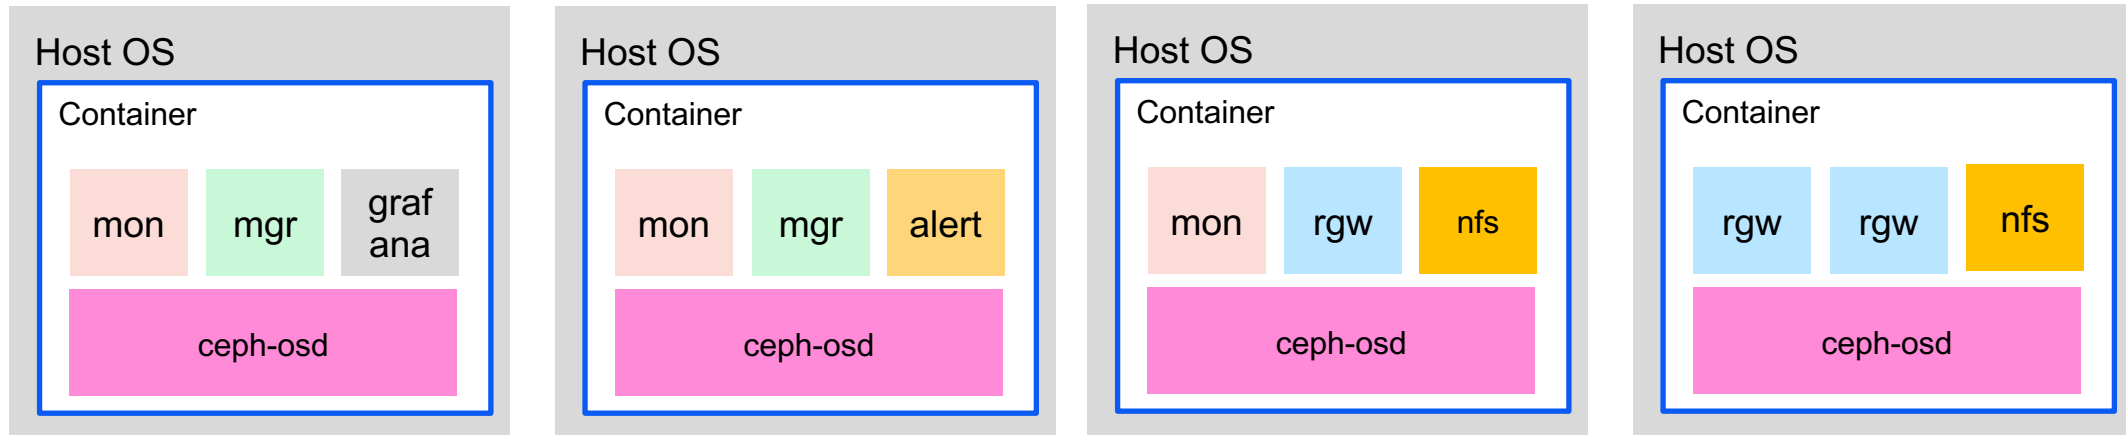
RADOS Gateway – Multiprotocol

The screenshot displays the IBM Storage Ceph web interface for NFS configuration. The left sidebar shows navigation options: Dashboard, Cluster, Pools, Block, NFS (selected), File Systems, and Object Gateway. The main content area is titled 'NFS' and features a '+ Create' button, a table view selector, a page size dropdown set to '10', and a search bar. Below these controls is a table with the following data:

Path	Pseudo	Cluster	Storage Backend	Access Type
john-bucket-1	/john-bucket-1	nfs-s3	Object Gateway	RW

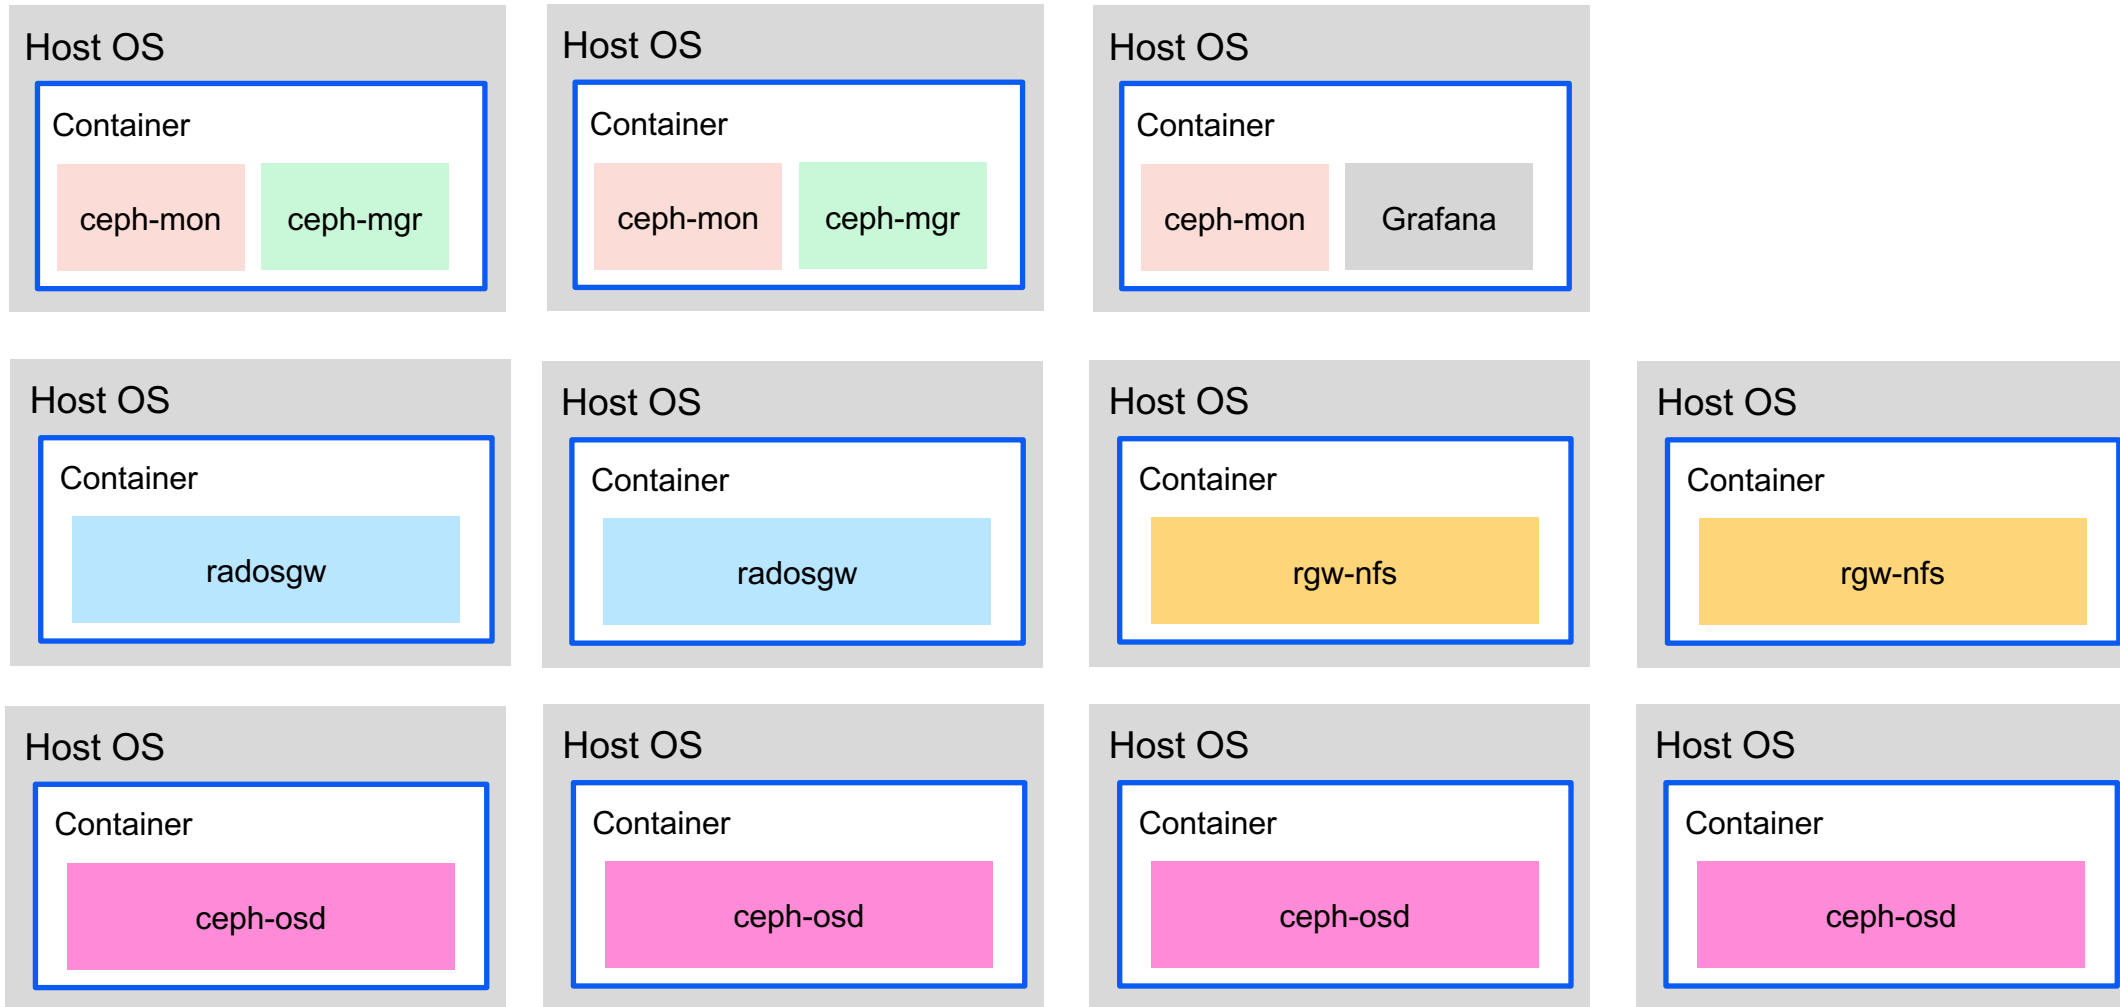
Below the table, it indicates '0 selected / 1 total'.

IBM Storage Ceph deployment – colocated daemons



Note: For illustrative purposes only; not a reference design.

IBM Storage Ceph deployment – non colocated daemons



Note: For illustrative purposes only; not a reference design.

S3 API command set and fidelity

IBM Storage Ceph delivers support for the S3 API that is mature and complete

Bucket Operations

PUT/GET/HEAD/DELETE

- Bucket
- Bucket ACL
- Bucket CORS
- Bucket Lifecycle
- Bucket Location
- Bucket Notification
- Bucket Policy
- Bucket Request Payment
- Bucket Status
- Bucket Website
- Bucket Versioning

Object Operations

PUT/GET/HEAD/COPY/DELETE

- Object
- Object Tagging
- Object ACL
- Object Legal Hold
- Object Lock
- Object Retention

Multipart Operations

PUT/GET/POST/DELETE

- Abort Multipart Upload
- Complete Multipart Upload
- Initiate Multipart Upload
- List Multipart Uploads
- List Parts
- Upload Part
- Upload Part Copy

Advanced Operations

- S3 Object Lock
- Secure Token Service (STS)
- Bucket quotas
- Bucket notifications
- Storage glass, Placement target
- **Server Side Encryption (SSE-S3)**

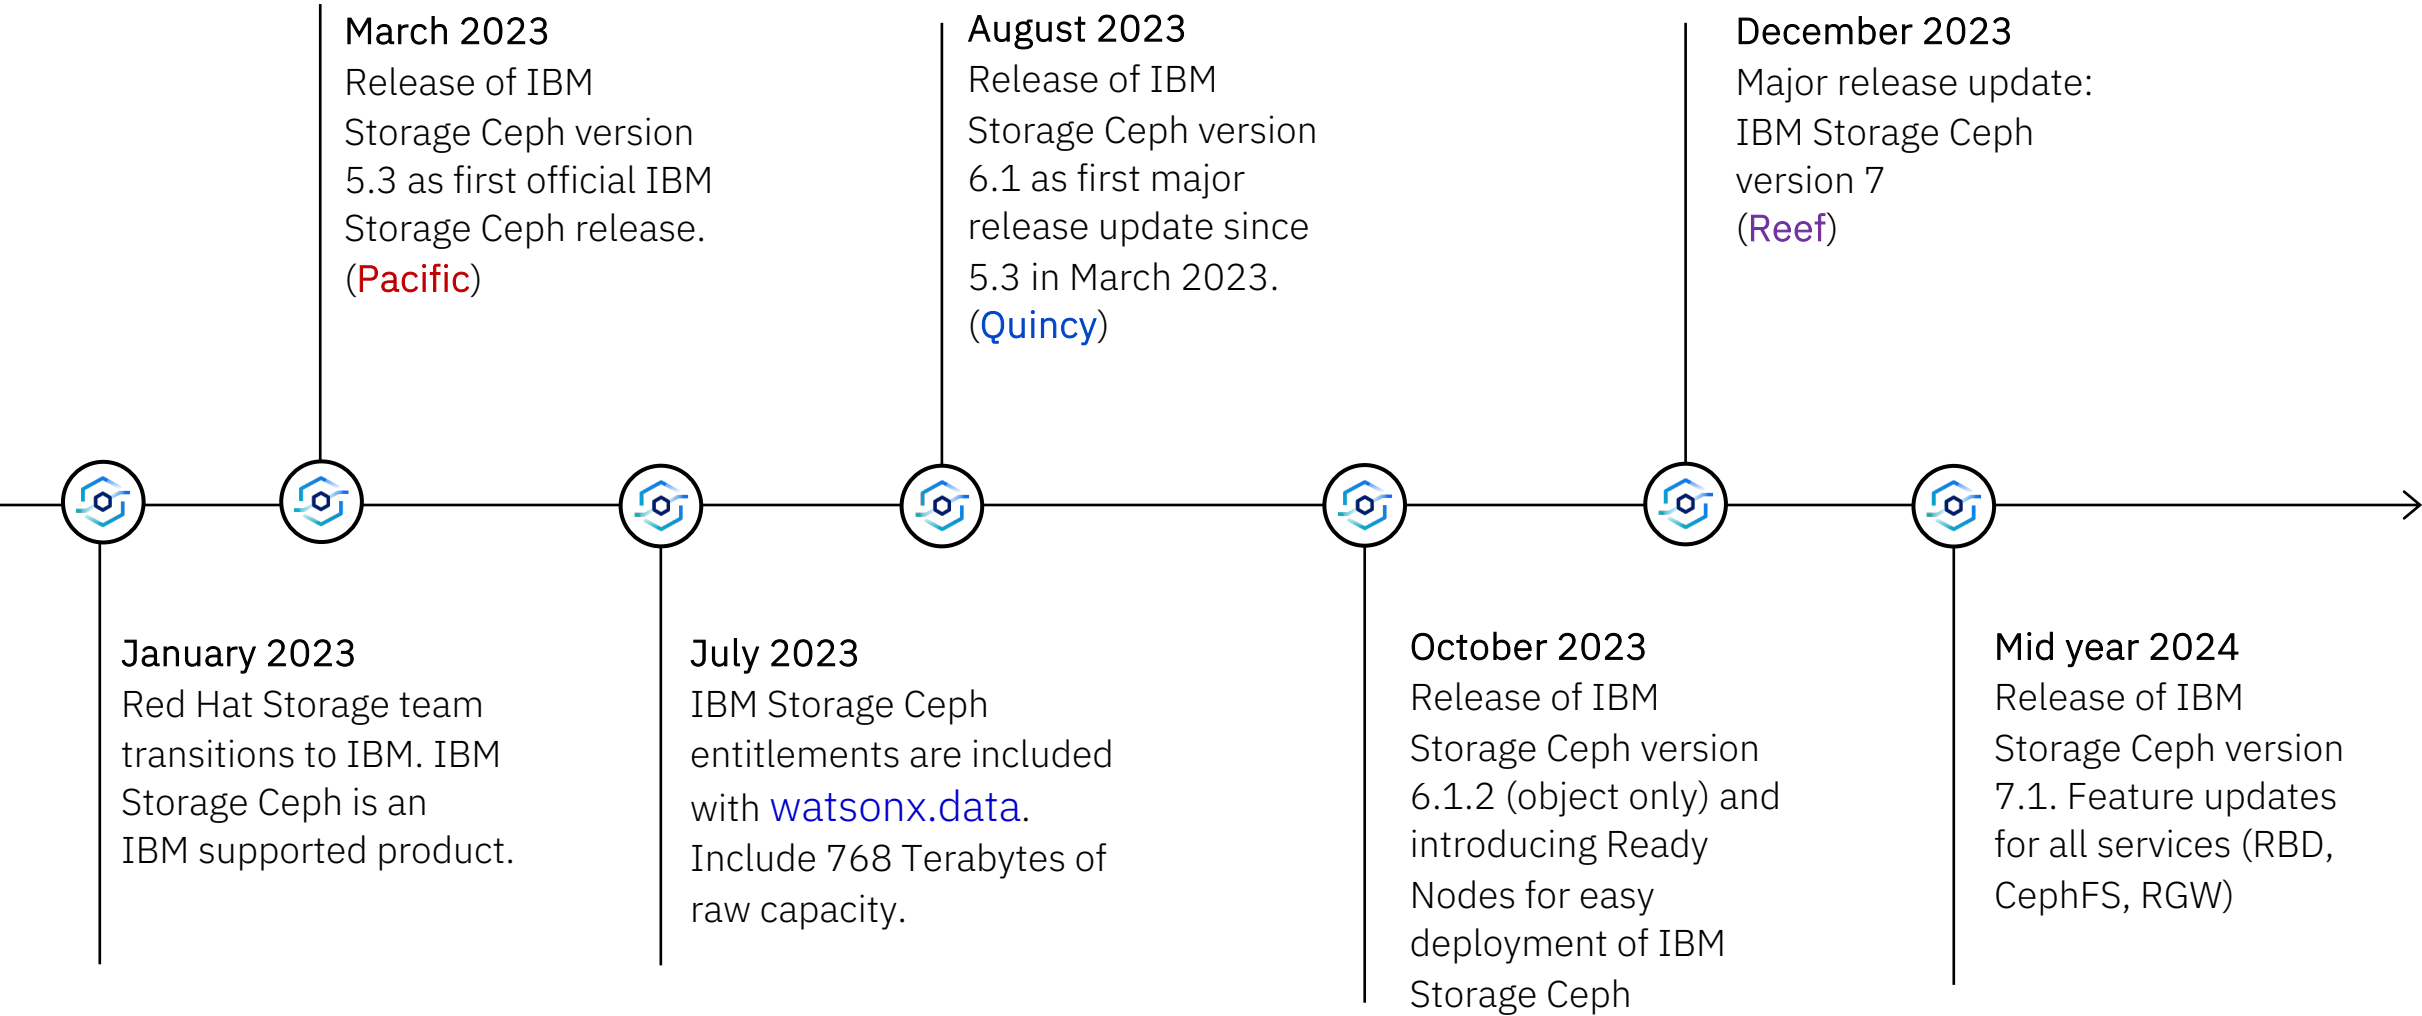
Data and AI Operations

- **Policy based rate limiting ingest**
- Granular bucket replication
- Policy based tiering to cloud (AWS, Azure, IBM)
- S3 Select
- **Policy based tiering to tape (future)**

IBM Storage Ceph 6 and Storage Ceph 7 updates for RADOS Gateway



IBM Storage Ceph 2023 - 2024 timeline



IBM Storage Ceph updates



IBM Storage Ceph 6 - Ceph release 17 (Quincy)

GA: August 2023

Block (RBD)

- SSD Write-back cache (GA)

File (CephFS)

- NFS for CephFS [Tech Preview]

Object (RGW)

- Policy based data archiving to AWS(GA)/Azure(TP)
- Policy based Bucket Granularity Replication(GA)

Management

- Day-2 management UI enhancements in the dashboard
- N+2 Upgrades

Deployment

- Support for deploying Ceph on VMware

IBM Storage Ceph updates



IBM Storage Ceph 7.0 – Ceph release 18 (Reef)

GA: December 2023

Block (RBD)

- NVMe over TCP (NVMe/TCP) protocol support [Tech Preview]

File (CephFS)

- NFS for CephFS (GA)

Object (RGW)

- Object Lock Worm Certification
- Archive Zone GA support
- S3-Select GA support CSV, Parquet File Formats.
- Erasure Coding on just 4 nodes

Management

- RGW Management Enhancements
- Enhanced CephFS Volume Management

Replication and Erasure Coding

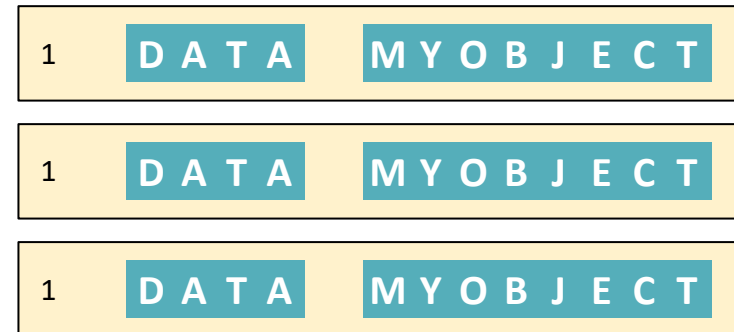
Replicas

- Replica 2x (SSD pool): min 3 nodes
- Replica 3x (HDD pool): min 4 nodes
- **Replica 4x (Metro cluster)**

Erasure Coding (EC Profiles)


- EC 4+2 : minimum 7 nodes
- EC 8+3 : minimum 12 nodes
- EC 8+4 : minimum 13 nodes
- **EC 2+2 : minimum 4 nodes**

REPLICATION



Total = 3.0x

ERASURE CODING




Total = 1.5x

IBM Storage Ceph buying and delivery options



1 Software only



IBM Storage Ceph



2 Ready Node HDD





IBM Storage Ceph System

3 Ready Node NVMe

Future - exploring...

4 Appliance



IBM Storage Ready Node with NVMe




Descriptions	Specifications
CPU	2 x Intel® Xeon® Gold 6438N 2G, 32C/64T, 16GT/s, 60M Cache, Turbo, HT (205W) DDR5-4800
Memory	16 x 32GB RDIMM (512GB)
OS Disks	BOSS-N1 controller card + with 2 M.2 480GB (RAID 1)
NVMe Device Density	3.84TB, 7.68, 15.36TB NVMe
NVMe Device Quantity	8, 16, or 24 Disks
Network	2 x 10GbE (SFP+ Optical Transceivers Included)
Network	2 x 100GbE (QSFP28 Optical Transceivers Included)
Dimensions	3.41" H x 18.97" W x 29.85" D (2U Rack Height)
Software Support	IBM Storage Ceph 7

More advanced features for the next Accelerate with ATG



Functionality

RGW policy-based data transition to public cloud



Archive or migrate

Object lifecycle transition into AWS compatible S3 cloud endpoints.

Data lifecycle

Extends current lifecycle transition and storage class model


Policy-based

Move data which meets policy criteria to an AWS compatible S3 bucket for archive, for cost, and manageability reasons

In overall

Cloud archiving provides a cost-effective, scalable, secure, and accessible way to store and preserve data

IBM Storage Ceph Object S3 Lifecycle Management



Lifecycle management

Ability for transitioning and/or expiration of objects.

Example policy:
Transition object to cold tier after 30 days.
Delete after 365 days

Granular object filtering

Granularly filter which objects in a bucket are susceptible of lifecycle management by object prefix or tags.

S3 Lifecycle Transition

Can be defined between arbitrary storage classes(tiers) in a cluster or to other S3 compatible endpoints (Including AWS),

Transition to cloud is a one-way only process.

S3 Lifecycle Expiration

Includes:

Current/Non-current
Delete Marker Expiration
Abort MultiPart Upload

Object storage for ML/analytics: S3select

S3select support general available

Support for all three defined S3select data formats:

- CSV
- JSON
- Parquet

Data analytics apps consume object

The data analytics market is a heavy consumer of object storage.

IBM Storage Ceph is first-class citizen in on-premise data analytics object Storage solutions.


Achieve reduced query times

IBM Storage Ceph keeps improving its features and integrations with market-leading tools like Presto, Trino and other apps.

Analytical tools use Parquet.

IBM Storage Ceph provides improved performance for these tools pushing down the queries into the RGWs.

NFS to RGW backend Tech Preview



IBM Storage Ceph RGW NFS gateway

NFS with RGW backend integration.

Allows for object access through NFS protocol.

This can be useful for easy ingests of object data from legacy applications which do not support the S3 object API.

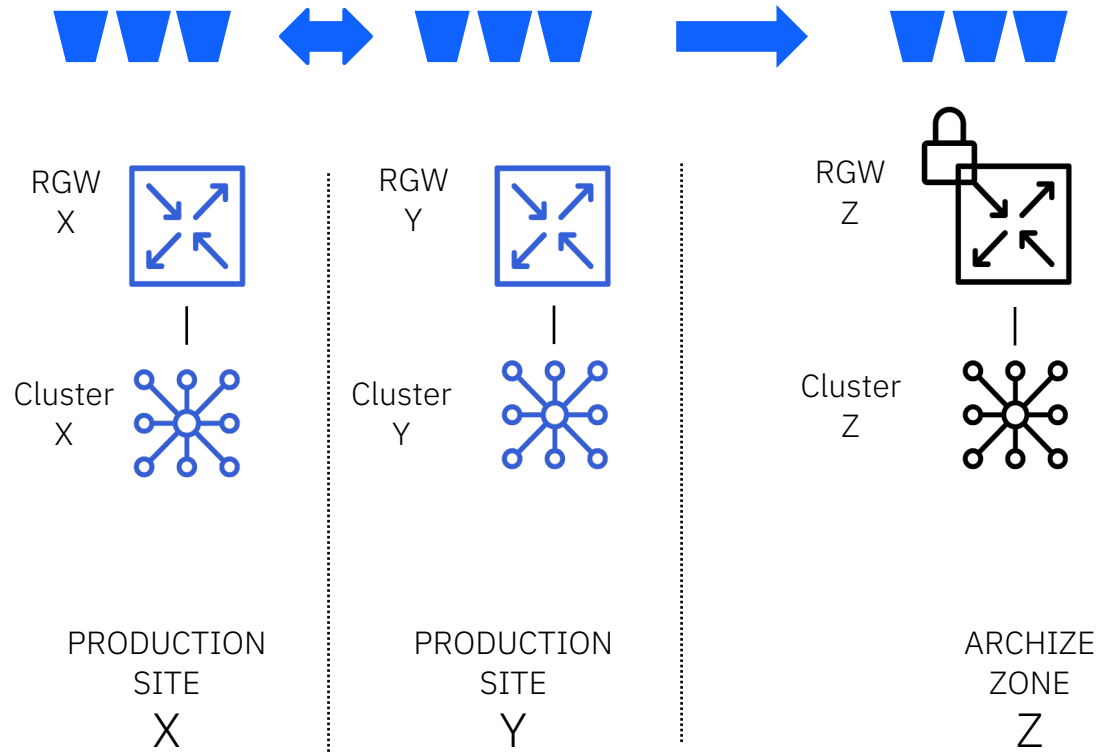
Use case examples for Data Scientists

A method to easily ingest existing business data from Windows and Linux clients into the Ceph object store.

An easy way to export results from analytics jobs and share results

Specifically, to users or applications that are unable to use the S3 API natively.

Object archive zone Tech Preview



IBM Storage Ceph archive zone

The archive zone receives all objects from the production zones.

It keeps every version for every object, providing the user with an object catalogue that contains the full history of the object.

Archive zone provides immutable objects that cannot be deleted nor modified from RGW endpoints.

Archive zone purpose and benefit

Ability to recover data from the archive zone.

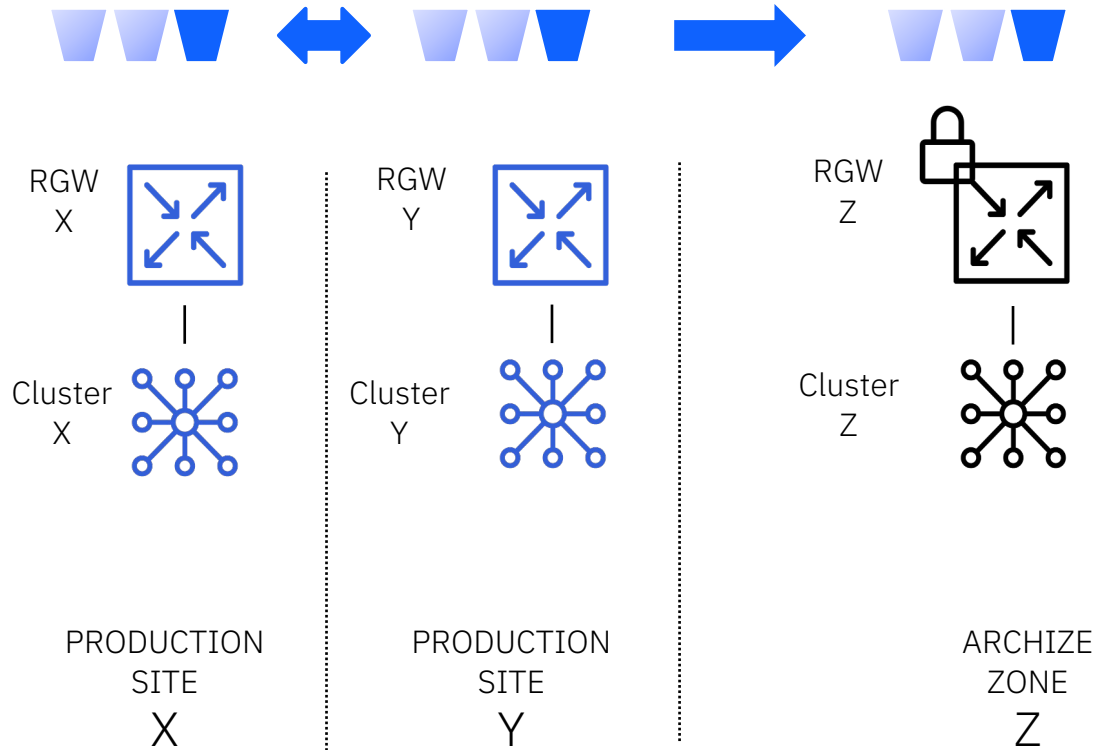
Enables for recovery of any version of any object that existed on production sites.

In case of data loss, ransomware or disaster recovery, still all valid versions of all objects can be recovered easily.

Also suitable for compliancy related use cases.

Functionality

Object archive zone with bucket granularity



Object archive zone bucket granularity

Allows clients to enable or disable replication to the archive zone on a per object bucket case.

Distinctions can be made based on a single bucket granular level.

Object archive zone granular choice

Goal is to reduce data storage in the archive zone.

In example, a set of test/development buckets are probably non-business critical.

System administrators may then decide to disable replication to the archive zone for these types of object data buckets.

Learning resources



IBM Storage Ceph resources

IBM website

IBM Storage Ceph enterprise downstream
<https://www.ibm.com/products/ceph>


IBM Seismic

Ceph Sales Kit on Seismic
<http://seismic.ceph.blue>

Cohasset Inc. Report

Ceph Object Lock (WORM) certification

[Report download link](#)



<https://www.ibm.com/downloads/cas/PJZN8VE3>

IBM Ceph easy demo videos


IBM Storage Ceph video demos
<http://easy.ceph.blue>

IBM Storage Ceph

IBM Storage Ceph Documentation
<http://docs.ceph.blue>

IBM Community Blog

Blogpost about IBM Storage Ceph 7.0
[Link to blogpost](#)



<https://community.ibm.com>

IBM Storage Ceph Videos at https://mediacenter.ibm.com

The screenshot shows a web browser window with the URL <https://mediacenter.ibm.com/research/search?keyword=ATG>. The page features the IBM logo and navigation menus. A search bar contains the text "ATG", and the results indicate "Found 125 results".

On the right side of the page, there is a section titled "2 Channels found" with two channel options:

- ADVANCED TECHNOLOGY GROUP - STORAGE
- ADVANCED TECHNOLOGY GROUP - ISV ON POWER

A blue arrow points to the "ATG" text in the search bar, and a green arrow points to the "ADVANCED TECHNOLOGY GROUP - STORAGE" channel option. Below the channel list is a link that says "Go to Channel Results".


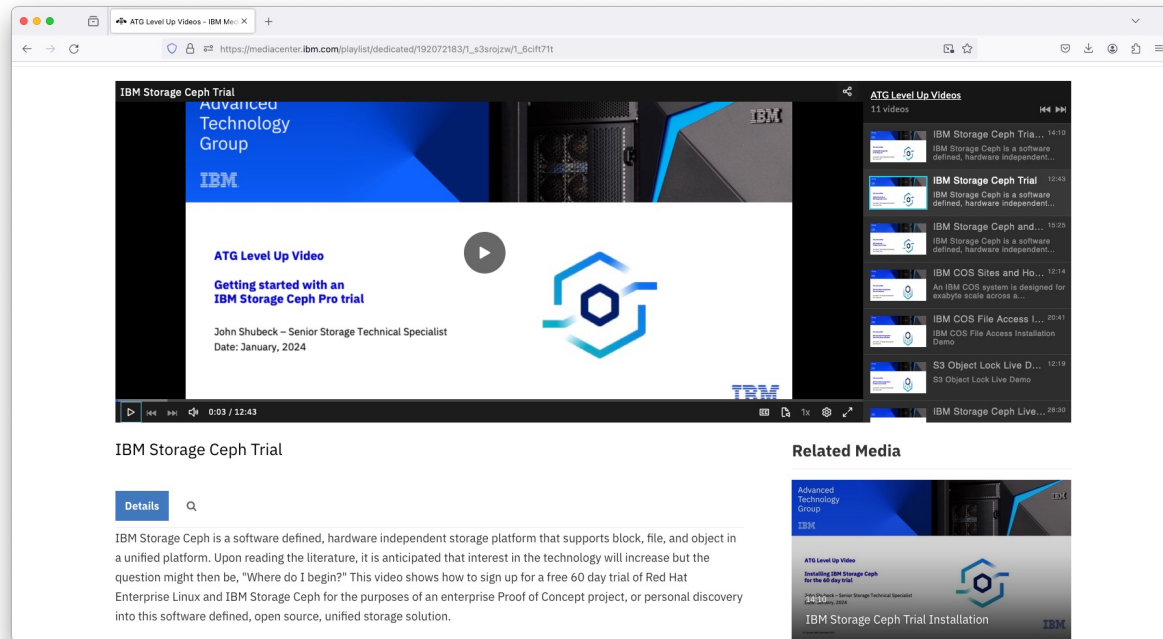
The main search results list includes:

- IBM DS8900F R9.3.2 Update - 03092023**
Listen to this replay with the ATG Storage team to hear about the updates that were delivered with
Keyword found in: 3 Tags [Show More](#)
- IBM Storage Ceph Trial Installation**
IBM Storage Ceph is a software defined, hardware independent storage platform that supports block, file, and object in ...
Keyword found in: 2 Tags | 1 Caption [Show More](#)
- IBM Storage Ceph and S3 Object Lock**
IBM Storage Ceph is a software defined, hardware independent storage platform that supports block, file, and object in ...
Keyword found in: 2 Tags | 2 Captions [Show More](#)

A "Cookie Preferences" button is visible in the bottom right corner.


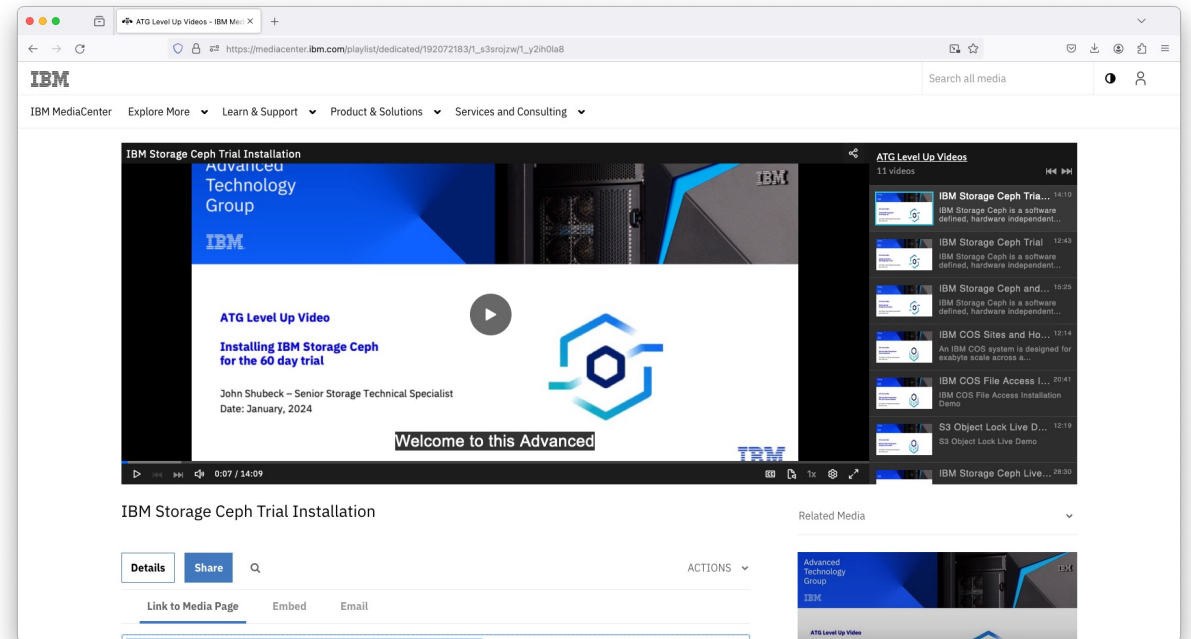
IBM Storage Ceph Trial videos

How to sign up and get subscriptions




https://mediacenter.ibm.com/media/IBM+Storage+Ceph+Trial/1_6cift71t

How to install a POC cluster




https://mediacenter.ibm.com/media/IBM+Storage+Ceph+Trial+Installation/1_y2ih0la8

IBM Storage Ceph Redbooks



<https://www.redbooks.ibm.com/abstracts/redp5721.html>



<https://www.redbooks.ibm.com/abstracts/redp5715.html>

IBM TechZone for IBM Storage Ceph Test Drive

The screenshot shows a web browser window displaying the IBM TechZone page for "North America ATG Storage - IBM Storage Ceph Test Drive - (VMware based) Overview". The page features a dark header with the word "Ceph" in yellow text, which is highlighted by a white box and a blue arrow. A large teal circle with the text "DO IT!" is overlaid on the page content. The page includes a navigation sidebar on the left with sections like Overview, Resources, Environments, Metadata, and Comments. The main content area contains introductory text about the test drive, a star rating system, and author information for John Shubeck and collaborators. A "Cookie Preferences" button is visible in the bottom right corner.

IBM Technology Zone | My library | Help

North America ATG Storage - IBM Storage Ceph Test Drive - (VMware based) Overview

☆☆☆☆☆ (0) Rate this resource

The IBM Storage Ceph (VMware based) of... al Ceph storage cluster.

The TechZone Ceph Test Drive is a virtual... is located in the Advanced Technology Group (ATG) lab in Herndon, Virginia.

The IBM Storage Ceph cluster is set up t... (GW). The key elements of the object gateway Realm, Zonegroups, Zones,

Placement Targets, and Storage Pools a... 3 API as well as the multi-protocol access via NFS is available.

The baseline configuration can also serv... ed lab environment but is treated with the operational care of a production

system. Therefore, after each demo, it co... nate, VMware snapshots can be used to restore the system to a known good

state.

*** NEW *** This demo has been updated f... 0 software.

If you need further assistance, please reach o...

Author
John Shubeck (jshubeck@us.ibm.com)

Collaborators
BethAnn Stugis (bethann.stugis@us.ibm.com), John Sing (jmsing@us.ibm.com), Carl Brown (carltb@us.ibm.com)

Cookie Preferences

Accelerate with ATG Survey

Please take a moment to share your feedback with our team!

You can access this 6-question survey via [Menti.com](#) with code 1708 6924 or

Direct link <https://www.menti.com/alwhyze7z1gz>

Or

QR Code



IBM