



Washington Systems Center - Storage

Accelerate with IBM Storage:

Spectrum Discover v2.0.2 Update

Isom L. Crawford Jr., PhD.
North America Software Defined Infrastructure
Subject Matter Expert Team
Isom.Crawford@ibm.com

Norman Bogard
Senior IT Specialist
Spectrum Storage Technical Leadership Team
bogey@us.ibm.com

Dan Thompson
Spectrum Storage Technical Specialist
danthomp@us.ibm.com



Accelerate with IBM Storage Webinars

The Free IBM Storage Technical Webinar Series Continues in 2020...

Washington Systems Center – Storage 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.

Located in the Accelerate with IBM Storage Blog:

<https://www.ibm.com/developerworks/mydeveloperworks/blogs/accelerate/?lang=en>

Also, check out the WSC YouTube Channel here:

https://www.youtube.com/channel/UCNuks0go1_ZrVVF1igOD6Q

2020 Upcoming Webinars:

January 14 – Brocade / IBM b-type SAN Modernization

Register Here: <https://ibm.webex.com/ibm/onstage/g.php?MTID=ed703072fdd739b7e1384ed84869aed8d>

January 21 - Cisco / IBM c-type SAN Analytics

Register Here: <https://ibm.webex.com/ibm/onstage/g.php?MTID=eabc0e050b2ff8bdcea4e012c30dfbf71>



Session Objectives

Spectrum Discover v2.0.2 Update

- **Spectrum Protect Data Source**
 - Single Helm Chart Deployment
 - Collection Admin
 - Application Catalog
 - Spectrum Scale Live Events

Spectrum Protect Data Source

What is supported?

- Spectrum Protect versions 8.X and newer
- Backup/archive client data
- All Spectrum Protect server platforms (e.g. Windows, Linux, AIX, etc)
- Tagging, policies*, collections, searching, reporting

What is not supported?

- Versions 7.X and older**
 - **Test coverage limitation. Can be considered on a case-by-case basis
- Protect-VE (virtual environments)
- Spectrum Protect Plus
- *Content search of backup/archive data

Create a Spectrum Protect Connection

- New connection type 'IBM Spectrum Protect'
- Required inputs:
 - Connection name
 - Server name/IP
 - ODBC Port (51500 by default)
 - Database instance user (e.g. tsminst1)
 - Database instance user password
- Spectrum Discover uses an ODBC connection to pull information from the DB2 database used by Spectrum Protect

Data Connections
Add Data Source Connection

Connection Name The field can't be empty.

Connection Type

Select a Collection

Schedule Data Scan

Spectrum Protect Server IP Invalid input.

ODBC Port Invalid input. Must be integer less than 65536.

Instance User Invalid input.

Instance User Password Invalid input.

Site (Optional)

| Source Name | Platform | Cluster | Data source | Site | Online | Scan Status |
|-------------|----------------------|-------------------------------------------------|-------------|------|--------|-------------|
| sp2 | IBM Spectrum Protect | 88.f5.a9.20.73.e0.e9.11.a5.39.00.50.56.b2.4b.98 | MODEVVM49 | | ● | |

Info harvested from Protect Server

Scan Spectrum Protect Connection

Data Source Connections

| Data Source Connections | | | | | | | Cancel |
|-------------------------|----------------------|-------------------------------------------------|-------------|------|--------|-------------|-----------|
| Source Name | Platform | Cluster | Data source | Site | Online | Scan Status | Next Scan |
| sp2 | IBM Spectrum Protect | 88.f5.a9.20.73.e0.e9.11.a5.39.00.50.56.b2.4b.98 | MODEVVM49 | | ● | | |
| spwin | | | | | | | |

Data Source Connections

| Data Source Connections | | | | | | | Cancel |
|-------------------------|----------------------|-------------------------------------------------|-----------------|------|--------|---------------------------------------------------------------|-----------|
| Source Name | Platform | Cluster | Data source | Site | Online | Scan Status | Next Scan |
| sp2 | IBM Spectrum Protect | 88.f5.a9.20.73.e0.e9.11.a5.39.00.50.56.b2.4b.98 | MODEVVM49 | | ● | | |
| spwin | IBM Spectrum Protect | eb.57.94.e0.c5.28.11.e9.a4.49.00.50.56.b2.d7.5a | WIN-RUT6M1TE74U | | ● | Scanning backup data on Spectrum Protect server (Step 2 of 4) | |

- Protect scans are done in four phases
 - Preparing to scan connection
 - Scanning backup data on Spectrum Protect server
 - Scanning archive data on Spectrum Protect server
 - Indexing new scan data

Searching

- Search experience is largely the same for Spectrum Protect.
- Search results can include both primary and Protect data

or start a visual exploration

| | | |
|------------------------------------|------------------------------------------|--------------------------------------|
| <input type="checkbox"/> Cluster | <input type="checkbox"/> Datasource | <input type="checkbox"/> Owner |
| <input type="checkbox"/> Platform | <input type="checkbox"/> Site | <input type="checkbox"/> Tier |
| <input type="checkbox"/> SizeRange | <input type="checkbox"/> TimeSinceAccess | <input type="checkbox"/> MgmtClass |
| <input type="checkbox"/> NodeName | <input type="checkbox"/> Fileset | <input type="checkbox"/> Filespace |
| <input type="checkbox"/> State | <input type="checkbox"/> COLLECTION | <input type="checkbox"/> TEMPERATURE |

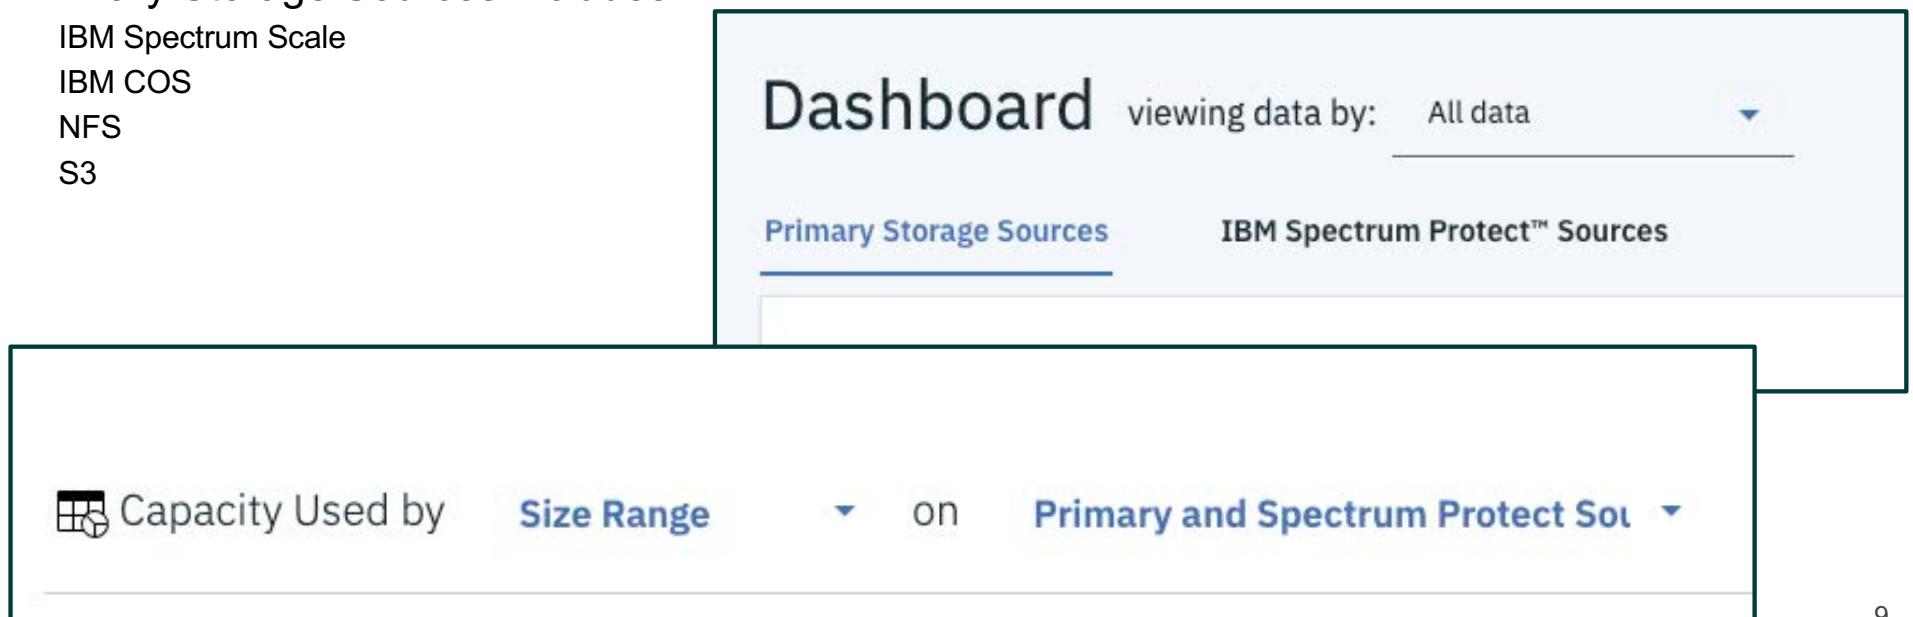
| | |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <p><input type="checkbox"/> PLATFORM</p> <p><input type="checkbox"/> Select all</p> <p><input type="checkbox"/> IBM Spectrum Protect (48,000,194)</p> <p><input type="checkbox"/> NFS (70,243)</p> <p><input type="checkbox"/> Spectrum Scale (1,000,004)</p> | <p><input type="checkbox"/> SIZERANGE</p> <p><input type="checkbox"/> Select all</p> <p><input type="checkbox"/> extra small (733,019)</p> <p><input type="checkbox"/> large (2,449)</p> <p><input type="checkbox"/> medium (43,093,675)</p> <p><input type="checkbox"/> small (5,241,297)</p> <p><input type="checkbox"/> Empty value (1)</p> |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

New Spectrum Protect Facets

| Facet | Description |
|-----------|---------------------------------------------------------------|
| MgmtClass | Management Class (how many copies to keep, for how long, etc) |
| NodeName | Name of the client node |
| FileSpace | File space name for a client node |
| State | Backup state (e.g. ACTIVE/INACTIVE), Archive |

New Dashboard Look

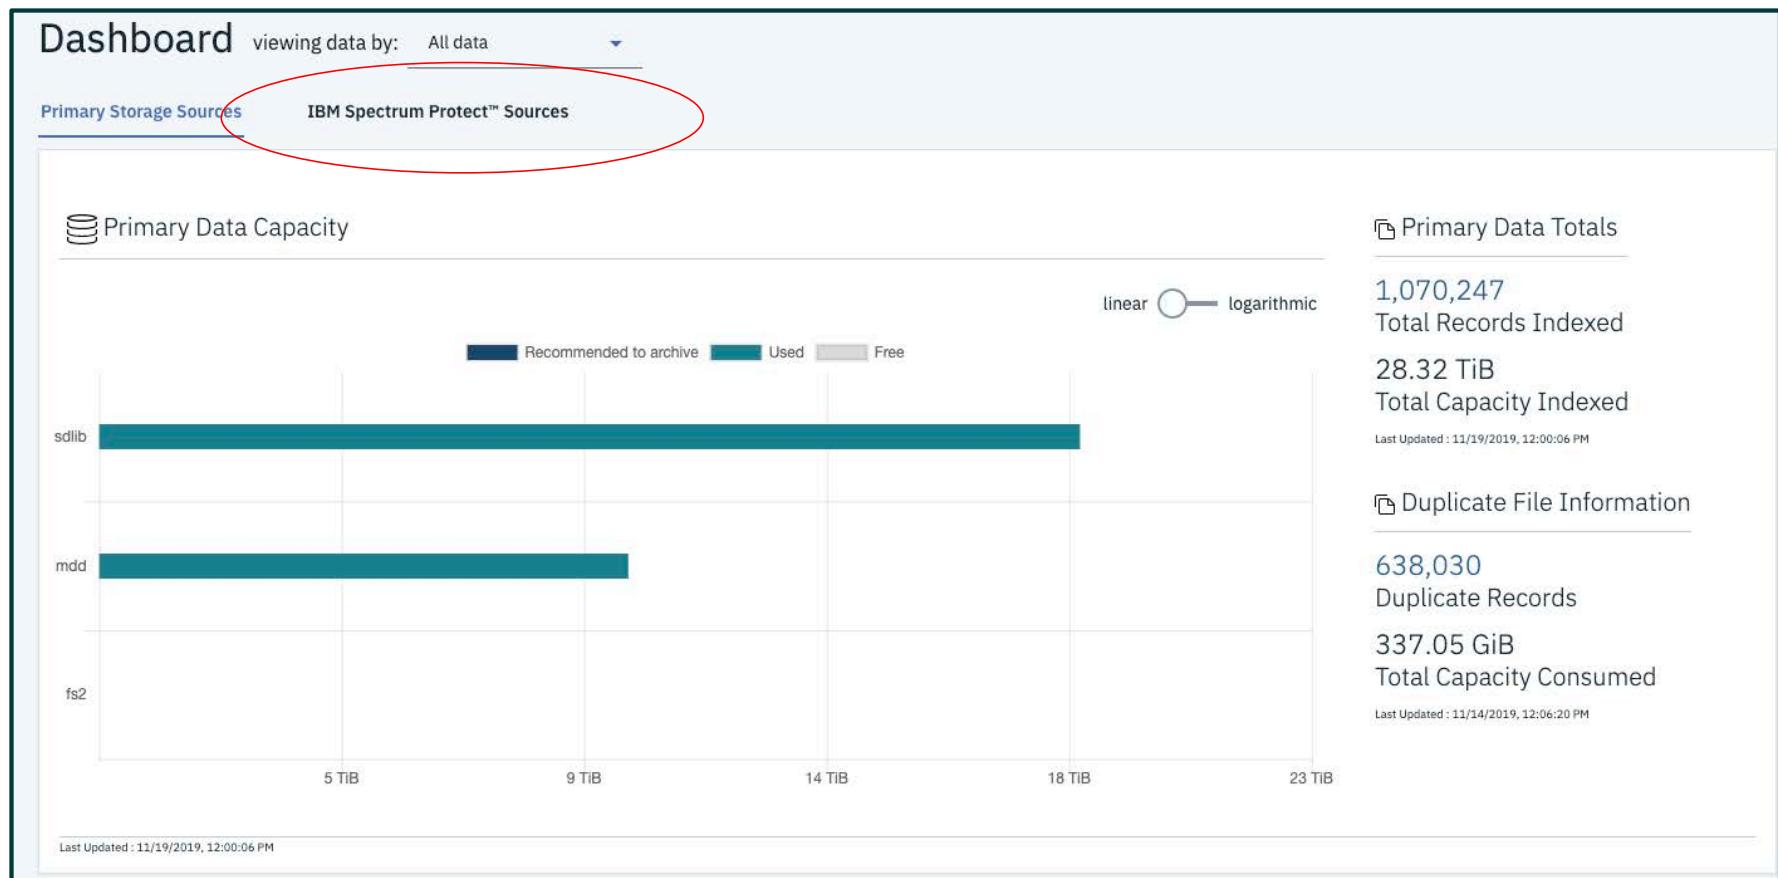
- Spectrum Protect data sources are broken out separately in the visualizations to better organize the experience
- We now have two groupings, Primary Storage Sources and IBM Spectrum Protect Sources
- Primary Storage Sources includes
 - IBM Spectrum Scale
 - IBM COS
 - NFS
 - S3



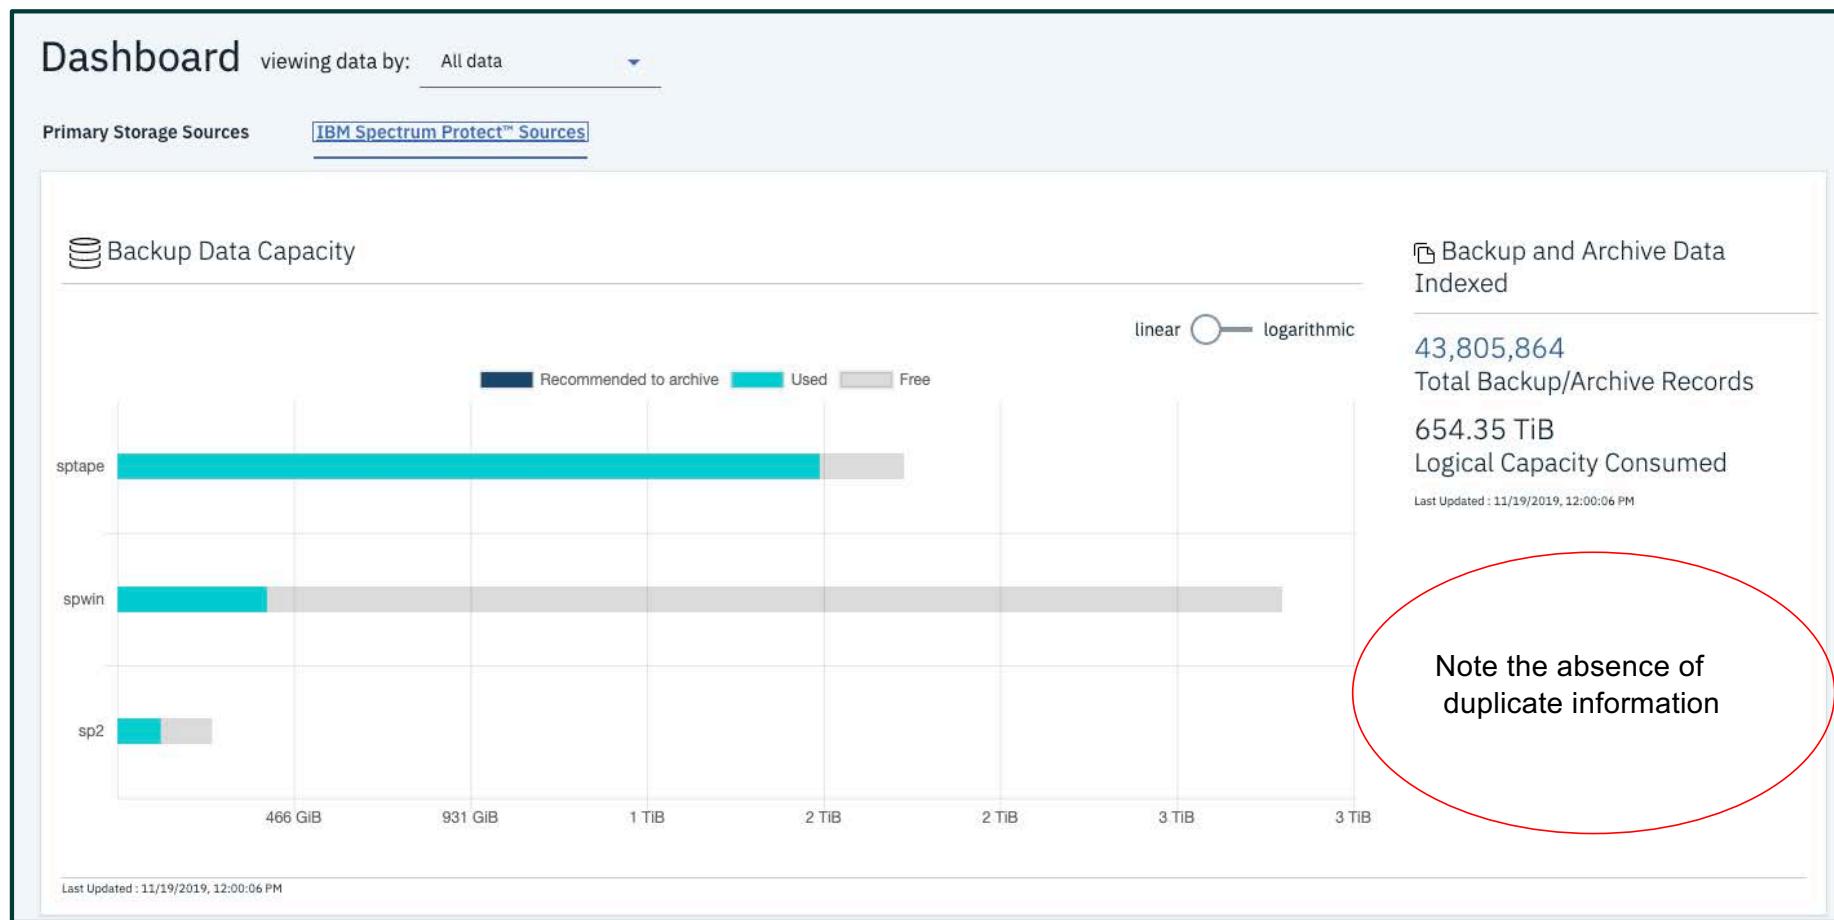
Washington Systems Center - Storage

New Dashboard Look - 1

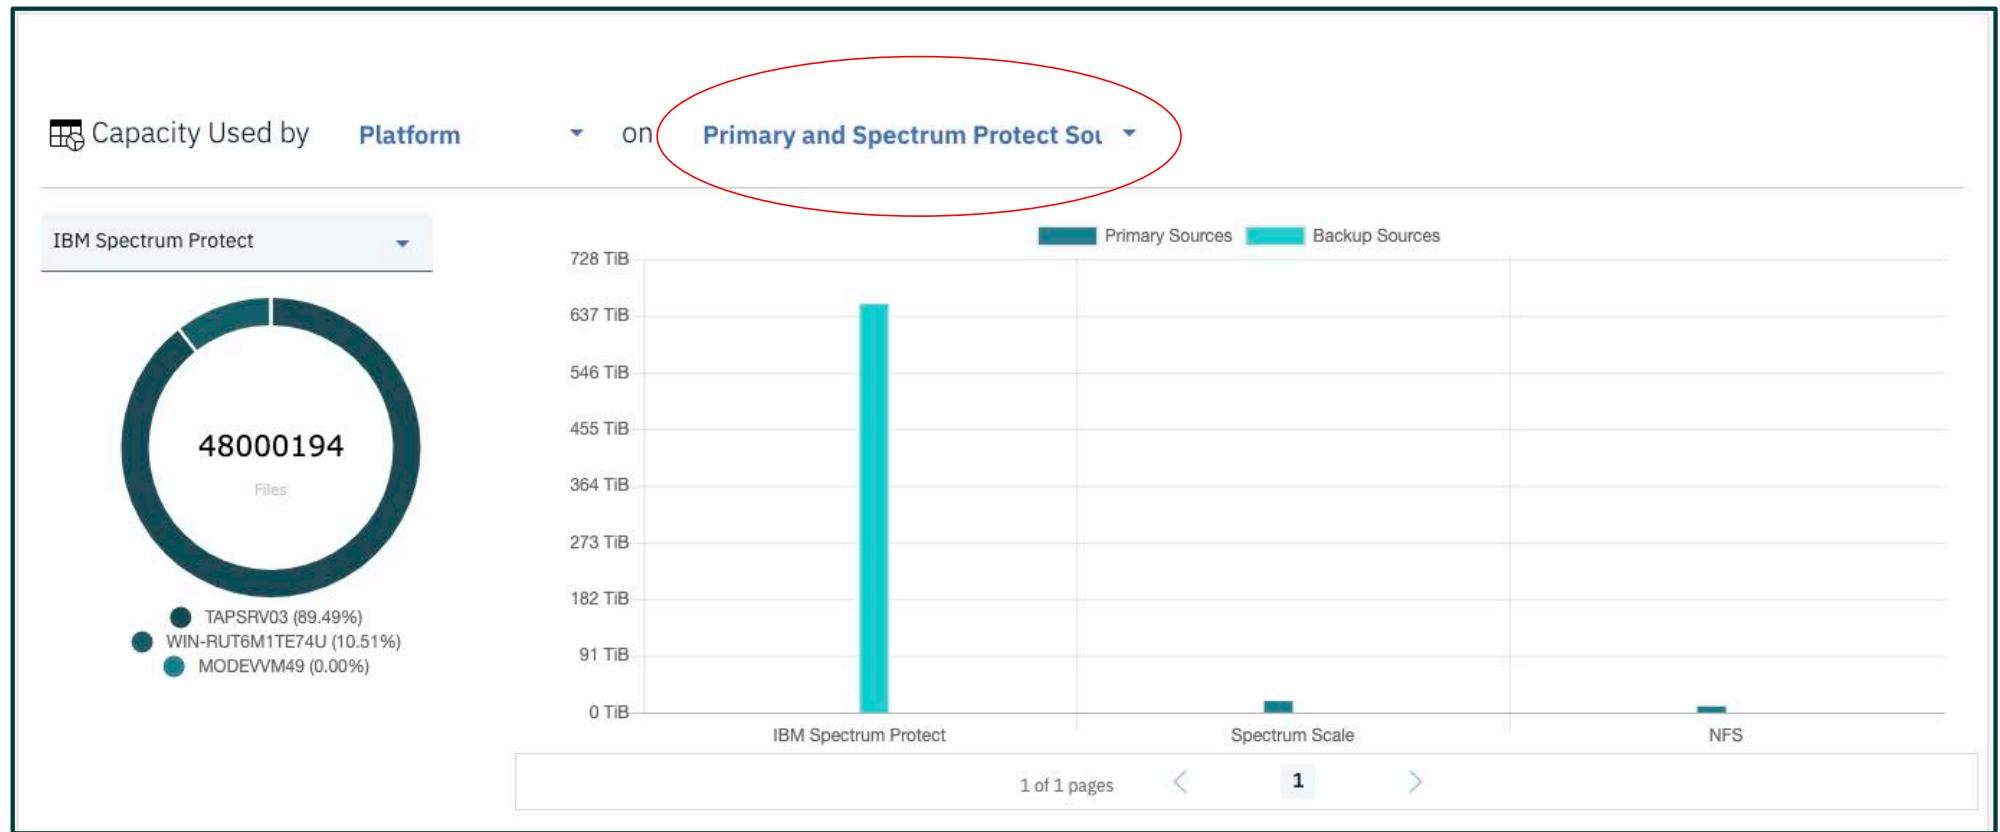
IBM Storage & SDI



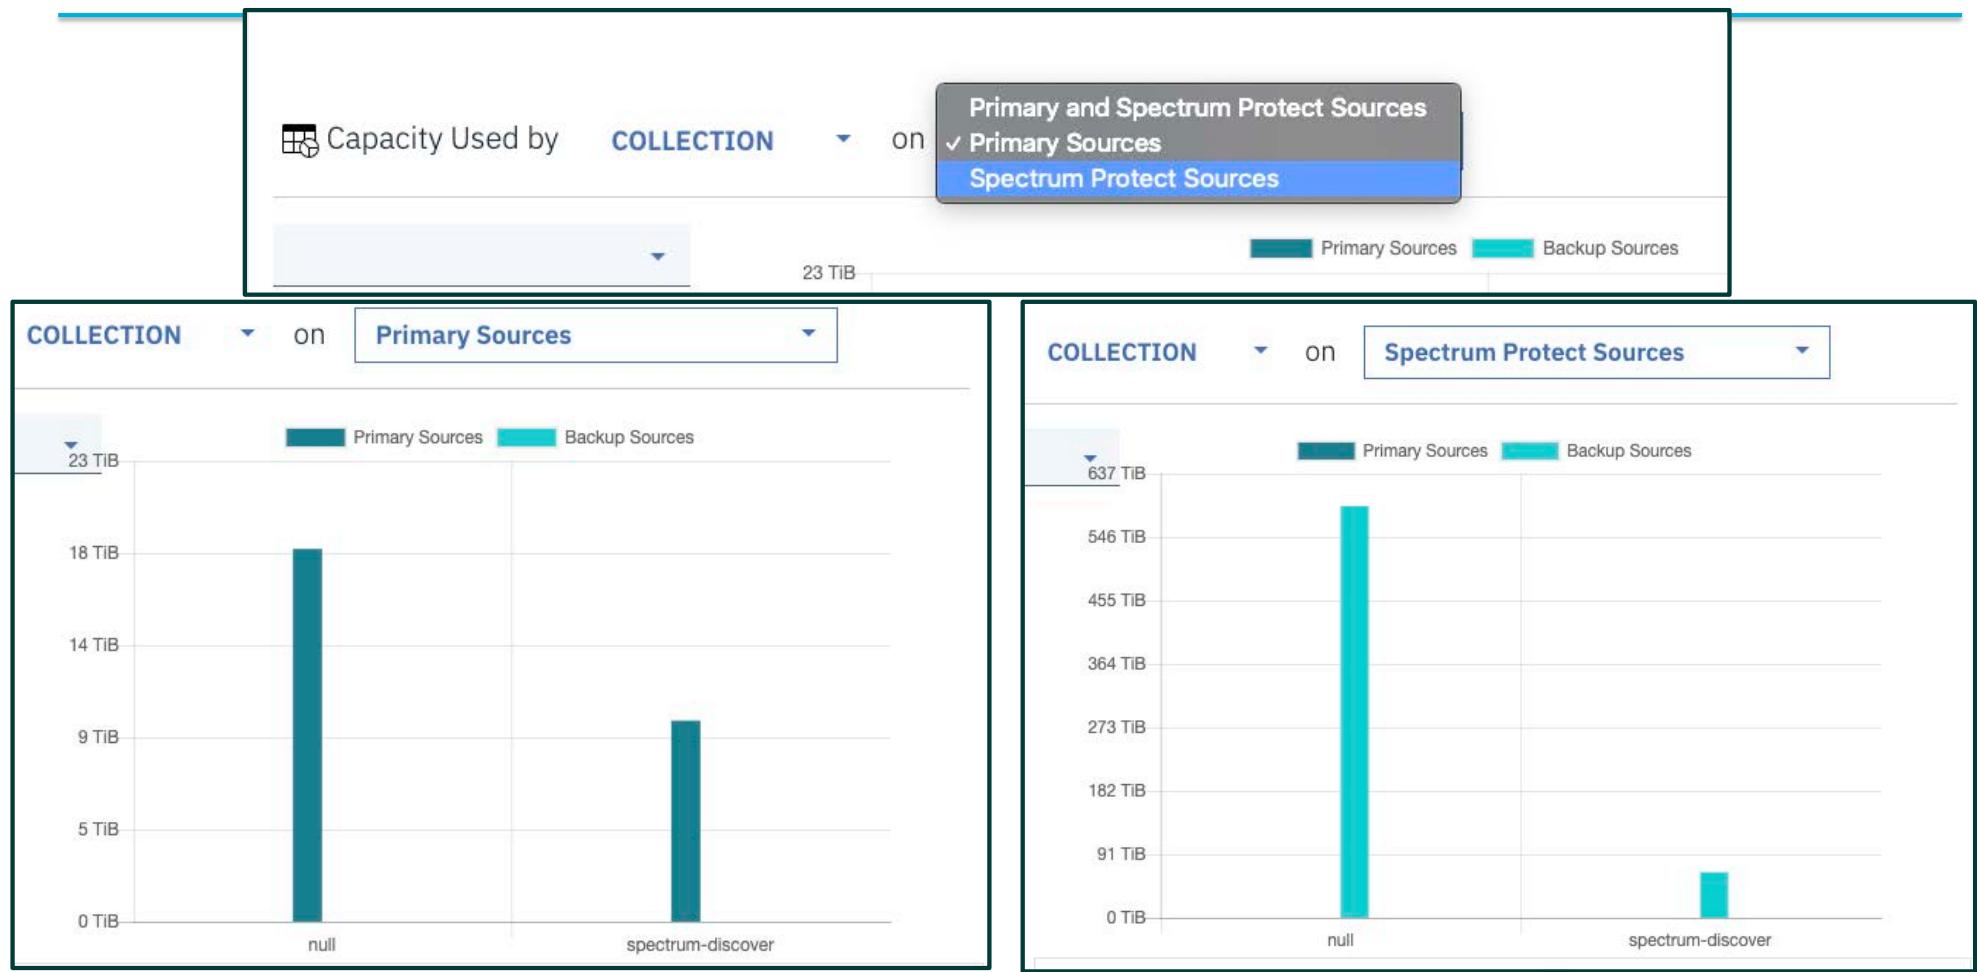
New Dashboard Look - 2



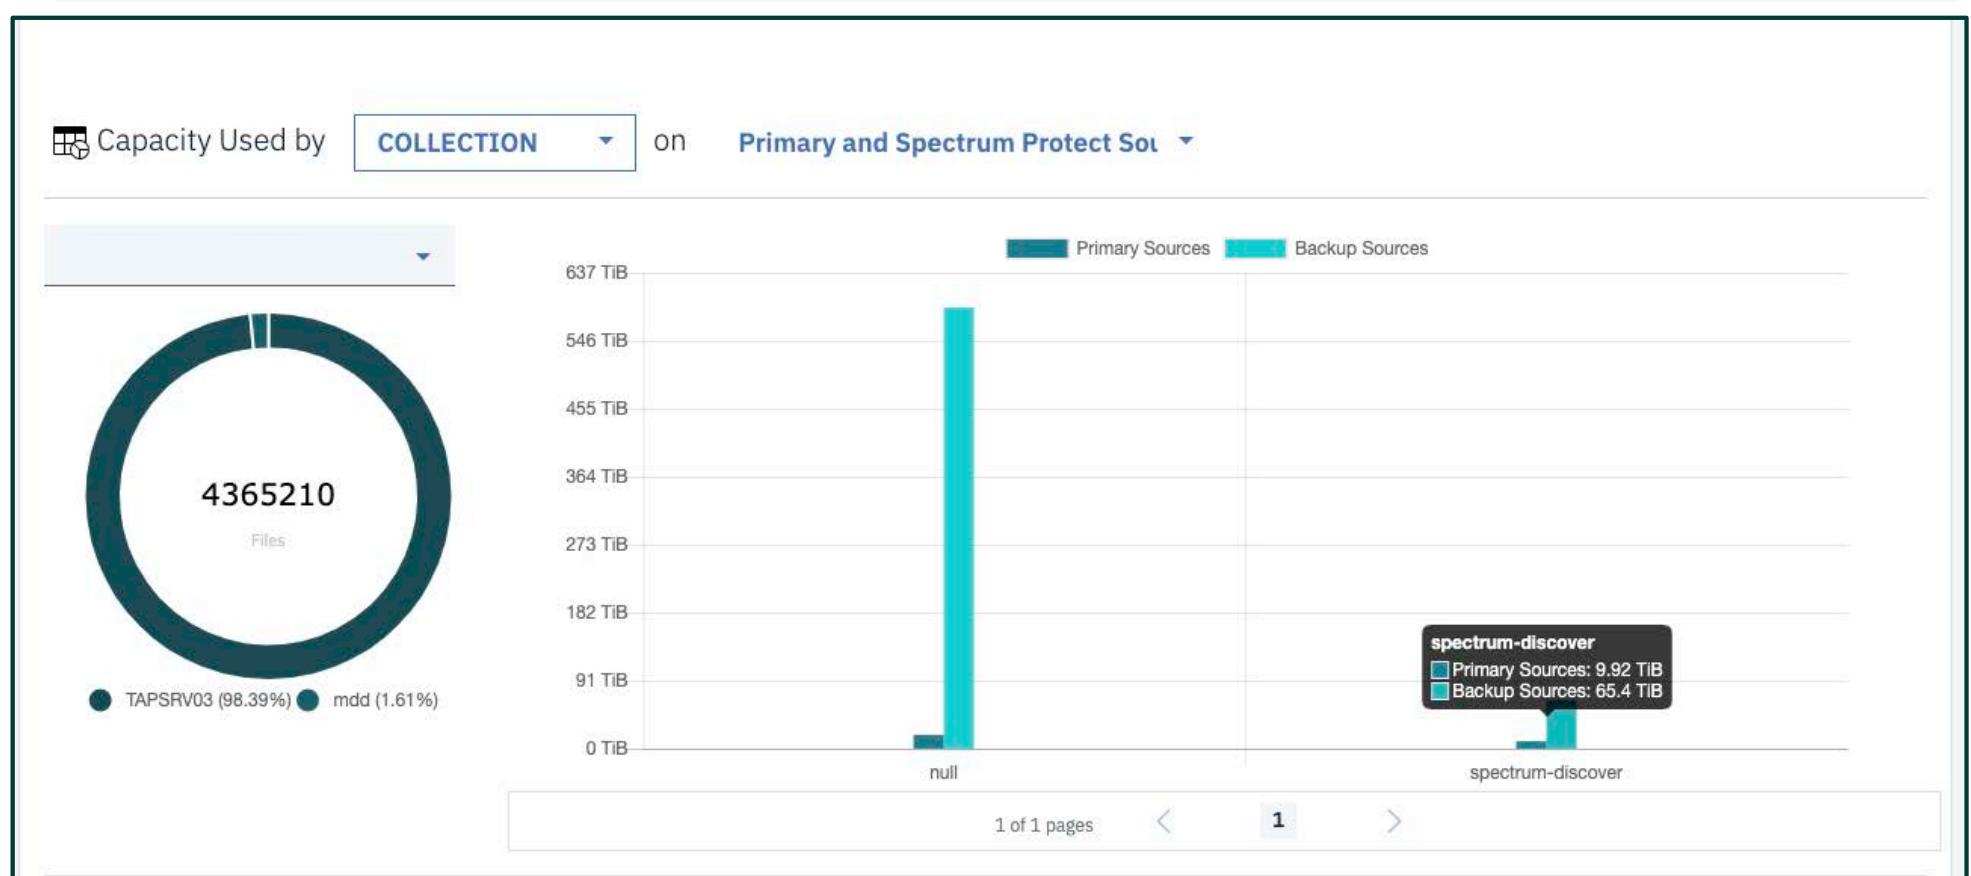
New Dashboard Look - 3



New Dashboard Look - 4



New Dashboard Look - 5



Session Objectives

Spectrum Discover v2.0.2 Update

- Spectrum Protect Data Source
- **Single Helm Chart Deployment**
- Collection Admin
- Application Catalog
- Spectrum Scale Live Events

Single Helm Chart Deployment

Changes

- Single namespace
- Same component naming convention across applications
- All pods have new labels for simpler lookups
- All containers and helm chart re-built at time of OVA build

Previously

1. Separate namespace per application
2. RBAC permissions to access components in other namespaces
3. Service calls using FQDN - <service>.
 - <namespace>.svc.cluster.local
4. Two-step process to access pod logs:
 - kubectl get pods -n policyengine
 - kubectl logs <pod-name> -n policyengine
5. Not always clear which changes were in a build as the changelog may have commits that have not yet been re-built to Artifactory

Now in v2.0.2

1. Single namespace for all of Spectrum Discover
 - spectrum-discover
2. Fewer permissions required
 - Everything is within a single namespace
3. Service calls using just service names
 - <service>
4. Single call to access pod logs
5. All applications re-built at time of OVA build
 - Same time the change log is compiled

Single Helm Chart Deployment Continued...

Additional Benefits

- Simpler Deployment
 - Single Helm install call
- Fewer Kubernetes components
- Simpler RBAC

LISTING ALL SPECTRUM DISCOVER PODS

✓ `kubectl get pods -n spectrum-discover`

GETTING CONNECTION MANAGEMENT LOGS

✓ `kubectl logs -l app=spectrum-discover-connmgr -n spectrum-discover`

Session Objectives

Spectrum Discover v2.0.2 Update

- Spectrum Protect Data Source
- Single Helm Chart Deployment
- **Collection Admin**
- Application Catalog
- Spectrum Scale Live Events

Collection Admin

What is a collection?

- A name for a group of records in Spectrum Discover
 - Collection name
- Spectrum Discover can limit a user's visibility to one or more collections
 - Example: Users in one department, only manage the metadata records from their department
- Achieved by Tagging Spectrum Discover records with a collection name
 - Uses special COLLECTION tag
 - Users are assigned roles on collections

Collection Admin Role

- Admin role
 - Manages all users and groups
- Data Admin role
 - Manages all data, collections and user and group access to collections
- Collection Admin
 - Manages data within a list of collections and user and group access to those collections
 - A role between Data Admin and Data User
- Data User
 - Manages data within a set of collections

➤ Spectrum Discover Supports 2 Role Types:

- Domain level
 - A role is assigned to the user within the domain
 - Collections do not apply
 - DataAdmin and Admin roles are assigned at this level
- Collection level
 - A role is assigned to a user/group on a collection
 - Role assignment comprises of 3 parts
 - user-id/role-id/collection-id
 - Data User roles and Collection Admin roles are assigned this way

➤ If a user is assigned multiple roles

- ✓ Role with the highest privilege applies

Overview of Roles

Admin Role

- Manage users/groups and access
 - Does not access/manage the data in the system
 - Can edit local users and local groups
 - Details, roles, & groups
 - Can create collections
 - Cannot manage collection policies
 - Can edit an external (LDAP/COS) user/group
 - Can only add roles to these members

Data Admin Role

- Manage all data within the system
- Controls user's access to data on the system
 - Can edit users and groups
 - Roles & collections only
- Restricted to Collection Admin and Data User roles only
 - Can create
 - Collections
 - Collection policies
 - Can assign collections to connections

Collection Admin Role

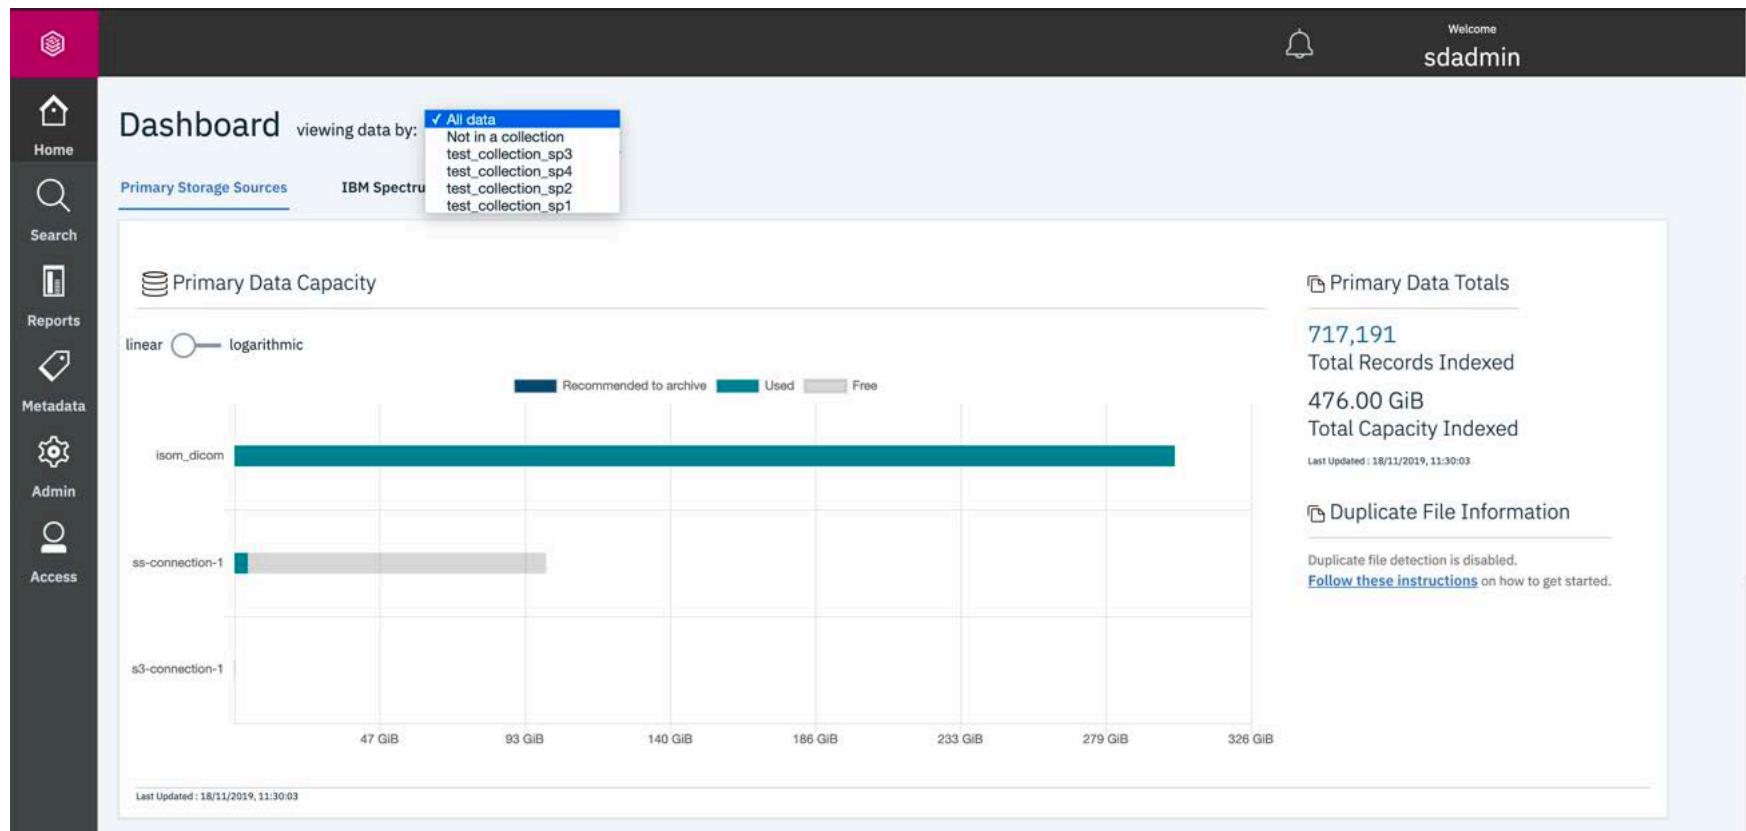
- Manage the Data and Data Users within a set of assigned collections
 - Can edit users and groups
 - Add/remove Data User roles for the collections they administer only
 - Can only see the collections they administer
 - Same applies to users/groups roles
 - Cannot edit a user's groups
 - Does not know of other collections in the group
 - Cannot create collections or collection policies

Data User Role

- Can search & create policies
 - Collections assigned to only
 - Cannot manage users
 - Cannot create or edit collections
 - Can only see collections assigned to

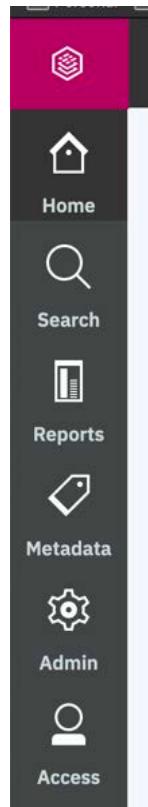
Collection Admin Metadata Management

- Admin user now has dropdown for view by collection

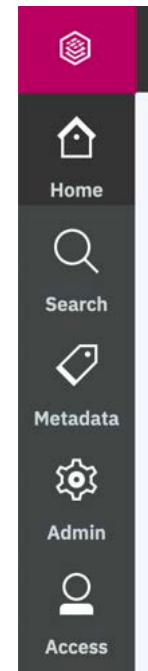


Sidebar Views

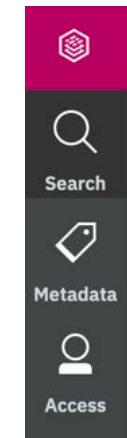
Data Admin



Collection Admin

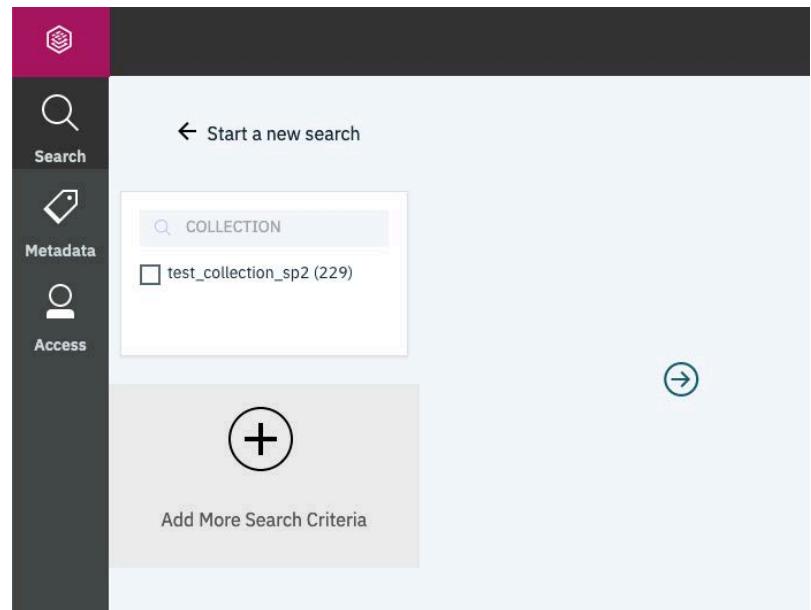


Data User



Metadata Management - Search

- Data Admin
 - No change
- Collection Admin
 - Results restricted to data within collections assigned
 - Transparent to user
- Data User
 - Results restricted to data within collections assigned
 - Transparent to user



Metadata Management - Policies

- Data Admin
 - Can view/edit all policies/tags
- Collection Admin
 - Can view/edit policies that apply to collections assigned
 - Can view/edit policies that is owner of
 - Can view any tag
 - Can only create/modify characteristics tags
- Data User can:
 - Can view/edit policies that is owner of
 - Can view any tag
 - Can only create/modify characteristics tags

✓ There is now a collections dropdown in policy page to manage collections

Edit policy

Inactive Active

Name: test_one_dicom

Policy Type: CONTENT SEARCH

Collections:

Type search collection

spectrum-discover

test_collection_delete

test_collection_sp1

Application: contentsearchagent

Tag: dicommfg

Search Expression: 1 x Search Expression

Value: Value matching expression

+Add Row

Schedule:

Now Daily Weekly Monthly

Session Objectives

Spectrum Discover v2.0.2 Update

- Spectrum Protect Data Source
- Single Helm Chart Deployment
- Collection Admin
- **Application Catalog**
- Spectrum Scale Live Events

Application Catalog - Overview

- Goal is to allow customers to use existing applications or easily create their own
- Create an ecosystem where customers can use applications not only provided by IBM but also 3rd parties
- Application SDK available on pypi and github
 - <https://pypi.org/project/ibm-spectrum-discover-application-sdk>
 - https://github.com/IBM/Spectrum_Discover_App_Catalog/application_sdk
- IBM provided applications available on dockerhub and source on github
 - Example Application
 - https://github.com/IBM/Spectrum_Discover_App_Catalog/example_application
 - <https://hub.docker.com/r/ibmcom/spectrum-discover-example-application>
 - Exif Header Extractor
 - https://github.com/IBM/Spectrum_Discover_App_Catalog/exif_header_extractor
 - <https://hub.docker.com/r/ibmcom/spectrum-discover-exif-header-extractor>
- More coming soon!
 - ✓ [spectrum-discover-lidar-metadata-extractor](#)
 - ✓ [spectrum-discover-spectrum-scale-get-xattrs](#)
 - ✓ [spectrum-discover-cos-get-x-amz-meta](#)
- Code available under MIT license to allow sharing and collaborating for applications

Additions to Existing Endpoints

- Refreshes the database for any new (or deleted applications) from dockerhub
 - Caches to the db2wh for faster lookup times. Refreshes every hour
 - `tcurl -X PUT https://localhost/db2whrest/v1/summary_tables/applicationcatalog/start`

The screenshot shows the 'Discover database' tab of the IBM Storage & SDI interface. The left sidebar includes icons for Home, Search, Reports, Metadata, Admin, and Access. The main content area is divided into three sections:

- Metadata Summarization Database:**
 - Auto-refresh: On
 - Last table refresh: 11/19/2019, 11:00:02 PM
 - Last refresh duration: 0:00:00
 - Next scheduled refresh: Every 30 minutes
 - Run table refresh** button
- Duplicate Record Database:**
 - Auto-refresh: Off
 - Last table refresh: Invalid Date
 - Last refresh duration: Not available
 - Run table refresh** button
- Application Catalog Cache:**
 - Auto-refresh: On
 - Last table refresh: 11/19/2019, 11:00:03 PM
 - Last refresh duration: 0:00:01
 - Next scheduled refresh: Every hour at :00
 - Run table refresh** button

New Endpoints

- publicregistry – Queries the database for available applications
 - tcurl -X GET https://localhost/api/application/appcatalog/publicregistry | jq
 - tcurl -X GET https://localhost/api/application/appcatalog/publicregistry/ibmcom/spectrum-discover-example-application | jq
- image – Acts upon docker images. Downloads or Deletes to local docker cache
 - tcurl -X POST https://localhost/api/application/appcatalog/image/ibmcom/spectrum-discover-example-application | jq
 - tcurl -X DELETE https://localhost/api/application/appcatalog/image/ibmcom/spectrum-discover-example-application | jq
- helm – Creates/deletes a Kubernetes pod and can scale the replicas
 - tcurl -X GET https://localhost/api/application/appcatalog/helm | jq
 - tcurl_json -X POST https://localhost/api/application/appcatalog/helm -d@example.json | jq
 - tcurl_json -X PATCH https://localhost/api/application/appcatalog/helm/quieting-lambkin-example-application -d@replicas.json | jq
 - tcurl -X DELETE https://localhost/api/application/appcatalog/helm/quieting-lambkin | jq

Examples

```
[moadmin@darthrevan ~]$ tcurl -X GET https://localhost/api/application/appcatalog/publicregistry | jq
% Total    % Received % Xferd  Average Speed   Time   Time  Current
                                         Dload  Upload   Total Spent  Left  Speed
100 1159 100 1159  0     0  --:--:-- --:--:-- 1828
{
  "success": "true",
  "message": "Successfully retrieved all applications.",
  "output": [
    {
      "repo_name": "ibmcom/spectrum-discover-example-application",
      "star_count": 0,
      "pull_count": 2757,
      "created": "2019-11-14T16:46:23.787Z",
      "application_name": "example",
      "company_name": null,
      "company_url": null,
      "description": "Performs a character count on a specified file.",
      "filetypes": "all",
      "icon_url": null,
      "maintainer": null,
      "parameters": null,
      "version": "0.0.1",
      "license": "mit",
      "installed": 0,
      "installed_version": null,
      "instances": null,
      "additional_info": null,
      "update": "2019-11-19T20:07:28.195Z"
    },
    {
      "repo_name": "ibmcom/spectrum-discover-exif-header-extractor",
      "star_count": 0,
      "pull_count": 2735,
      "created": "2019-11-14T16:48:16.224Z",
      "application_name": "exif_header_extractor",
      "company_name": null,
      "company_url": null,
      "description": "Extracts exif header information for jpeg, tiff files",
      "filetypes": "jpg,jpeg,tiff",
      "icon_url": null,
      "maintainer": null,
      "parameters": null,
      "version": "0.0.1",
      "license": "mit",
      "installed": 0,
      "installed_version": null,
      "instances": null,
      "additional_info": null,
      "update": "2019-11-19T20:07:28.209Z"
    }
  ]
}
```

```
[moadmin@darthrevan ~]$ tcurl -X POST https://localhost/api/application/appcatalog/image/ibmcom/spectrum-discover-example-application | jq
% Total    % Received % Xferd  Average Speed   Time   Time  Current
                                         Dload  Upload   Total Spent  Left  Speed
100 118 100 118  0     0  31    0  0:00:03 0:00:03  --:--:-- 31
{
  "success": "true",
  "message": "Successfully pulled ibmcom/spectrum-discover-example-application application image."
}

[moadmin@darthrevan ~]$ tcurl -X DELETE https://localhost/api/application/appcatalog/image/ibmcom/spectrum-discover-example-application | jq
% Total    % Received % Xferd  Average Speed   Time   Time  Current
                                         Dload  Upload   Total Spent  Left  Speed
100 119 100 119  0     0  185   0  --:--:-- --:--:-- 185
{
  "success": "true",
  "message": "Successfully deleted application image: ibmcom/spectrum-discover-example-application"
}

[moadmin@darthrevan ~]$ cat example.json
{
  "application_name": "example",
  "version": "0.0.2",
  "description": "description",
  "repo_name": "ibmcom/spectrum-discover-example-application"
}

[moadmin@darthrevan ~]$ tcurl_json -X POST https://localhost/api/application/appcatalog/helm -d@example.json | jq
% Total    % Received % Xferd  Average Speed   Time   Time  Current
                                         Dload  Upload   Total Spent  Left  Speed
100 282 100 143 100 139   47   45  0:00:03 0:00:03  --:--:-- 47
{
  "success": "true",
  "message": "Successfully deployed application.",
  "chart_name": "example-application",
  "deployment_name": "running-uakari"
}

[moadmin@darthrevan ~]$ tcurl -X DELETE https://localhost/api/application/appcatalog/helm/running-uakari | jq
% Total    % Received % Xferd  Average Speed   Time   Time  Current
                                         Dload  Upload   Total Spent  Left  Speed
100    71 100    71  0     0   90    0  --:--:-- --:--:-- 90
{
  "success": "true",
  "message": "release \"running-uakari\" deleted\n"
}
```

Examples Continued...

```
[moadmin@darthrevan ~]$ cat replicas.json
{
  "replicas": 10
}

[moadmin@darthrevan ~]$ tcurl_json -X PATCH https://localhost/api/application/appcatalog/helm/running-uakari-example-application -d@replicas.json | jq
% Total    % Received % Xferd  Average Speed   Time   Time   Time  Current
          Dload  Upload   Total Spent    Left Speed
100  588  100  570  100    18  1965      62 --:--:-- --:--:-- --:--:-- 1972
{
  "success": "true",
  "message": "Successfully patched deployment.",
  "output": {
    "kind": "Scale",
    "apiVersion": "autoscaling/v1",
    "metadata": {
      "name": "running-uakari-example-application",
      "namespace": "spectrum-discover",
      "selfLink": "/apis/apps/v1/namespaces/spectrum-discover/deployments/running-uakari-example-application/scale",
      "uid": "d308e2e8-0b0e-11ea-891b-525400e51809",
      "resourceVersion": "5107383",
      "creationTimestamp": "2019-11-19T20:54:48Z"
    },
    "spec": {
      "replicas": 10
    },
    "status": {
      "replicas": 1,
      "selector": "app=example-application,release=running-uakari"
    }
  }
}
```

Creating a policy

Add new policy

Inactive Active

| | |
|-----------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------|
| Name example | Policy Type DEEP-INSPECT |
| Collections Type search collection | |
| Filter datasource like 'scale0' | |
| Application example-application | |
| Parameter extract_tags | Values char_count |
| |  |
| +Add tag | |
| Schedule | |
| <input checked="" type="radio"/> Now <input type="radio"/> Daily <input type="radio"/> Weekly <input type="radio"/> Monthly | |

Debugging

Much like debugging any other Kubernetes pod or docker container

```
kubectl describe pod running-uakari-example-application-bcf68949c-2lnrh
kubectl logs running-uakari-example-application-bcf68949c-2lnrh
kubectl exec running-uakari-example-application-bcf68949c-2lnrh -it bash
```

```
docker ps | grep running-uakari
docker logs -f 8a0df69547df
```

Session Objectives

Spectrum Discover v2.0.2 Update

- Spectrum Protect Data Source
- Single Helm Chart Deployment
- Collection Admin
- Application Catalog
- **Spectrum Scale Live Events**

Spectrum Scale Live Events support

Captures filesystem event notifications, which are delivered to Discover in near real-time using Kafka.

Provides a much more current representation of the state of the datasource/filesystem, reducing the need for scans.

What's new in 2.0.2?

Filesystem activity is captured and delivered using a clustered watch with a secondary sink (Kafka topic).
[Spectrum Scale version 5.0.3.1+]

The clustered watch provides a greater level of flexibility and robustness. Previously, events were delivered using a custom application '*tsspectrumdiscover*' running on a single Scale node.

tsspectrumdiscover will no longer be supported starting with 2.0.2. Scale systems running older versions of Scale will need to upgrade to take advantage of live events.

Scale Live event considerations and limitations

Requirements for enabling live events on Spectrum Scale:

Spectrum Scale live events will work starting with Scale 5.0.3.1, but due to a performance issue on the Scale side, Discover officially states beta support for live events starting with Scale version 5.0.4.1. We will recommend users upgrade to this level before enabling live events, especially if they encounter any issues keeping the Scale watch alive.

User must enable watch folder on Scale. A minimum of 3 Scale nodes are required to act as brokers.

The Scale nodes acting as brokers must meet a minimum local space requirement of 20GB in order to successfully enable the watch with a secondary sink.

When enabling live events using the GUI, the watch will be on the filesystem specified in the datasource connection.

Enabling Scale Live events

Can be enabled using the GUI or command line on the Scale node.

To enable when creating a Scale data source connection, simply check the box:

| | |
|---------------------------------------------------------------|----------------|
| Connection Name | ovalsquid |
| Connection Type | Spectrum Scale |
| <input checked="" type="checkbox"/> Enable live events (beta) | |
| <input type="checkbox"/> Select a Collection | |
| <input type="checkbox"/> Schedule Data Scan | |

| | |
|-------------------|------------------|
| User | root |
| Password | |
| Working Directory | /gpfs/gpfs0/work |
| Scan Directory | /gpfs/gpfs0 |
| Site (Optional) | Tucson |
| Cluster | ovalsquid.bvt |

Disabling Scale Live events

To disable live events from the GUI, just edit the connection and uncheck the 'Enable live events (beta)' option and apply the update.

Connection Name
ovalsquid

Connection Type
Spectrum Scale

Enable live events (beta) **Deselect this**

Select a Collection

Schedule Data Scan

User i
root

Password i
.....

Working Directory i
/gpfs/gpfs0/work

Scan Directory i
/gpfs/gpfs0

Site (Optional) i
Tucson

Cluster i
ovalsquid.bvt

Manually enabling/disabling live events

1. Create a connection to the SS system.
2. Copy over the Discover Kafka credentials required to establish a connection for the clustered watch external sink to location on Scale system.
 1. Get a token, then use token in command to get TLS cert:
https://www.ibm.com/support/knowledgecenter/SSY8AC_2.0.1/com.ibm.spectrum.discover.v2r01.doc/api_agents_gettlscertificate.html
 2. Copy output of PEM file to a directory (example: /home/{yourdir}/kafka_client.pem) on the Scale system
3. To set up Scale live events, the Scale administrator can run the following mmwatch command on the connected Scale system to capture the live events from a specific filesystem into Spectrum Discover.

```
# mmwatch <filesystem> enable --event-handler kafkasink --sink-brokers "Spectrum.Discover.NodeAddress:9093" --sink-topic scale-le-connector-topic --sink-auth-config <auth.file> --events IN_ATTRIB,IN_CLOSE_WRITE,IN MODIFY,IN_CREATE,IN_DELETE,IN_MOVED_FROM,IN_MOVED_TO
```

where <filesystem> is the filesystem to set the watch for, and <auth.file> is the full path to the configuration file with the settings to connect to the Spectrum Discover Kafka brokers.

auth.file contents:

```
-----  
SINK_AUTH_TYPE:CERT  
CA_CERT_LOCATION:/home/{yourdir}/kafka_client.pem  
CLIENT_PEM_CERT_LOCATION:/home/{yourdir}/kafka_client.pem  
CLIENT_KEY_FILE_LOCATION:/home/{yourdir}/kafka_client.pem
```

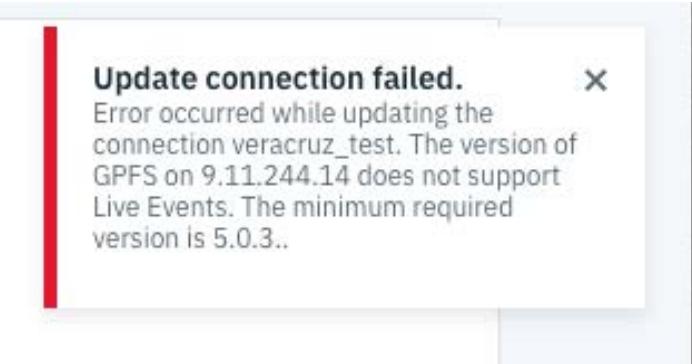
To stop the watch:

```
# mmwatch <filesystem> disable --watch-id <watch-id>
```

where <watch-id> is the ID of the established watch and can be obtained from output of `mmwatch all list`

Possible issues enabling or running live events

Cannot enable live events because Scale is not at the required level.

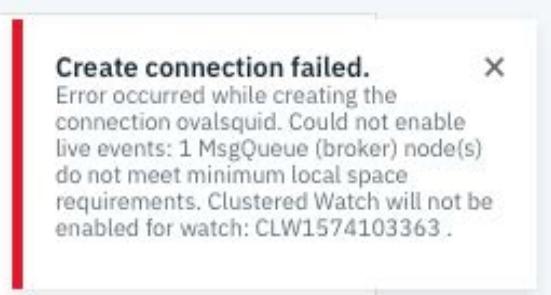


Solution: Have customer upgrade to Scale 5.0.3.1+(5.0.4.1 recommended).

Note that the output from the mmwatch command will be saved in the connection manager log. To view this log, you can use the alias *conlog* on the Discover node.

Possible issues enabling or running live events

Cannot enable live events due to not enough space on the brokers.



More detailed information can be found in the Spectrum Discover connection manager log.

Example:

```
2019-11-18 18:56:01,570 - connmgr - INFO - Live event enable: Starting watch with cmd /usr/lpp/mmfs/bin/mmwatch gpfs0 enable --event-handler kafkasink --sink-brokers "9.11.201.45:9093" --sink-topic scale-le-connector-topic --sink-auth-config /gpfs/gpfs0/work/kafka/auth.file --events IN_ATTRIB,IN_CLOSE_WRITE,IN MODIFY,IN CREATE,IN DELETE,IN MOVED_FROM,IN MOVED_TO
2019-11-18 18:56:12,217 - connmgr - INFO - [E] Broker node prt001st001 does not meet the minimum amount of local disk space required 20971520KB
2019-11-18 18:56:12,217 - connmgr - INFO - [W] Broker node prt002st001 meets the minimum amount of local disk space required 20971520KB but
2019-11-18 18:56:12,217 - connmgr - INFO - not the recommended amount of local disk space 41943040KB
2019-11-18 18:56:12,217 - connmgr - INFO - [W] Broker node prt003st001 meets the minimum amount of local disk space required 20971520KB but
2019-11-18 18:56:12,217 - connmgr - INFO - not the recommended amount of local disk space 41943040KB
2019-11-18 18:56:12,217 - connmgr - INFO - [E] 1 MsgQueue (broker) node(s) do not meet minimum local space requirements. Clustered Watch will not be enabled for watch: CLW1574103363
2019-11-18 18:56:12,217 - connmgr - INFO - Increase the amount of local disk space on the identified node(s) to meet the minimum of KB, recommended 41943040KB
2019-11-18 18:56:12,217 - connmgr - INFO - Or, if degraded performance is acceptable and all of the identified node(s) meet the degraded minimum of 10485760KB,
2019-11-18 18:56:12,217 - connmgr - INFO - run this command again with the '--degraded' option.
2019-11-18 18:56:12,217 - connmgr - INFO - mmwatch: Command failed. Examine previous error messages to determine cause.
2019-11-18 18:56:12,217 - connmgr - ERROR - Could not enable live events: 1 MsgQueue (broker) node(s) do not meet minimum local space requirements. Clustered Watch will not be enabled for watch: CLW1574103363
```

Solution: Customer increases space on Scale nodes, uses different nodes with more space, or manually starts watch folder with 'degraded' option.

Running Scale watch folder in degraded mode

Degraded option reduces the minimum required local storage for the Scale brokers to 10GB.

To enable this, a clustered watch must be manually enabled using the command line on one of the Scale nodes.

Here is an example:

```
# mmwatch gpfs0 enable --event-handler kafkasink --sink-brokers "thor.tuc.stglabs.ibm.com:9093" --sink-topic
scale-le-connector-topic --sink-auth-config auth.file --events
IN_ATTRIB,IN_CLOSE_WRITE,IN MODIFY,IN CREATE,IN DELETE,IN MOVED_FROM,IN MOVED_TO -degraded
```

Possible issues enabling or running live events

Filesystem activity is not being captured even though live events was successfully enabled.

Currently, errors that occur for a clustered watch initiated for live events are not surfaced to Spectrum Discover. Therefore, the state of the watch will need to be checked on the Scale nodes.

Solution: Check the state of the watch on Scale. Collect logs.

```
# mmwatch all status
Device          Watch ID      Watch State
delhaize       CLW1572947931  Active
               Node Name    Status
               prt001st001  HEALTHY
               prt005st001  HEALTHY
               prt006st001  HEALTHY

gpfs0          CLW1574112839  Active
               Node Name    Status
               prt001st001  DOWN
               prt005st001  DOWN
               prt006st001  DOWN
```

Log files of interest:

/var/adm/ras/mmwfcclient.log
/var/adm/ras/mmwatch.log

43

Possible issues enabling or running live events

Filesystem activity is not being captured even though live events was successfully enabled continued...

If the watch is still running:

- Validate the connection info is correct. The filesystem being watched should be the filesystem set in the connection
- Check for activity on scale-le-connector-topic on the Discover node:

```
[moadmin@thor ~]$ queryGroup.sh conn1 |grep scale-le-connector-topic
TOPIC                  PARTITION  CURRENT-OFFSET  LOG-END-OFFSET  LAG          CONSUMER-ID
HOST      CLIENT-ID
scale-le-connector-topic  0          3536241       3536241       0          spectrum-discover-producer-scale-le-75d4f6fd98-
grgtn-7cbf25a4-3d95-4ed0-881e-2c098fe61f09  /10.1.240.144  spectrum-discover-producer-scale-le-75d4f6fd98-grgtn
```

LOG-END_OFFSET should be incrementing. If it is not, then either there is no activity on the filesystem, or the events are not being delivered by the Scale system to Discover.

- If the watch was enabled manually, validate that the authentication credentials and *mmwatch* parameters are correct.

Possible issues enabling or running live events

Scale watch is failed on version <5.0.4.1

There is a known issue affecting the Kafka batch sizes prior to Scale version 5.0.4.1 which could lead to delivery timeouts of notifications from Scale to Discover. If notifications are not being received, and the watch state shows that it is DOWN, it could be due to this issue. To verify, check /var/adm/ras/mmwfclient.log for entries similar to the following:

```
2019-10-21_19:16:02: [E] WF Producer: t: scale-le-connector-topic a: 2 < Delivery timeout of 45 seconds for acknowledgement of all messages sent has been exceeded. 5692/10000 messages delivered. > b: modevvm42.tuc.stglabs.ibm.com:9093,modevvm43.tuc.s
2019-10-21_19:16:02: [E] WF Producer: t: scale-le-connector-topic a: 2 < Timeout hit. Elapsed time: 45.000000. > b: modevvm42.tuc.stglabs.ibm.com:9093,modevvm43.tuc.s
```

Common Scale commands to run

To get the Scale version:

```
# mmdiag
Current GPFS build: "5.0.3.3 ".
Built on Sep  5 2019 at 15:10:38
```

To get list of watches:

```
# mmwatch all list
Filesystem ces has no watchers.

Filesystem delhaize has 1 watcher(s):
  Watcher          WatchID/PID      Type      Start Time      Path
  -----
  (cluster) architeuthis1.tuc.stglabs.ibm.com CLW1572947931  filesystem Tue Nov  5 02:59:49 2019  /gpfs/delhaize

Filesystem gpfs0 has 1 watcher(s):
  Watcher          WatchID/PID      Type      Start Time      Path
  -----
  (cluster) architeuthis1.tuc.stglabs.ibm.com CLW1574112839  filesystem Mon Nov 18 14:34:55 2019  /ibm/gpfs0
```

Common Scale commands to run, cont

To show state of clustered watch on each node:

```
# mmwatch all status
```

| Device | Watch ID | Watch State |
|-------------|---------------|-------------|
| delhaize | CLW1572947931 | Active |
| | | |
| Node Name | | Status |
| prt001st001 | | HEALTHY |
| prt005st001 | | HEALTHY |
| prt006st001 | | HEALTHY |
| | | |
| gpfs0 | CLW1574112839 | Active |
| | | |
| Node Name | | Status |
| prt001st001 | | HEALTHY |
| prt005st001 | | HEALTHY |
| prt006st001 | | HEALTHY |

Washington Systems Center - Storage

Thank You!

Isom L. Crawford Jr., PhD.
North America Software Defined Infrastructure
Subject Matter Expert Team
Isom.Crawford@ibm.com

Norman Bogard
Senior IT Specialist
Spectrum Storage Technical Leadership Team
bogey@us.ibm.com

Dan Thompson
Spectrum Storage Technical Specialist
danthomp@us.ibm.com

