

Clustered Watch Folder

—

Lukas Rupprecht

Research Staff Member

(jump in for Deepavali Bhagwat, Lead Architect)



Clustered Watch Folder

Distributed, Scalable, Multi-Cluster File System
Event Notification

Setup and Watch File System accesses across
Clusters with a Single Command

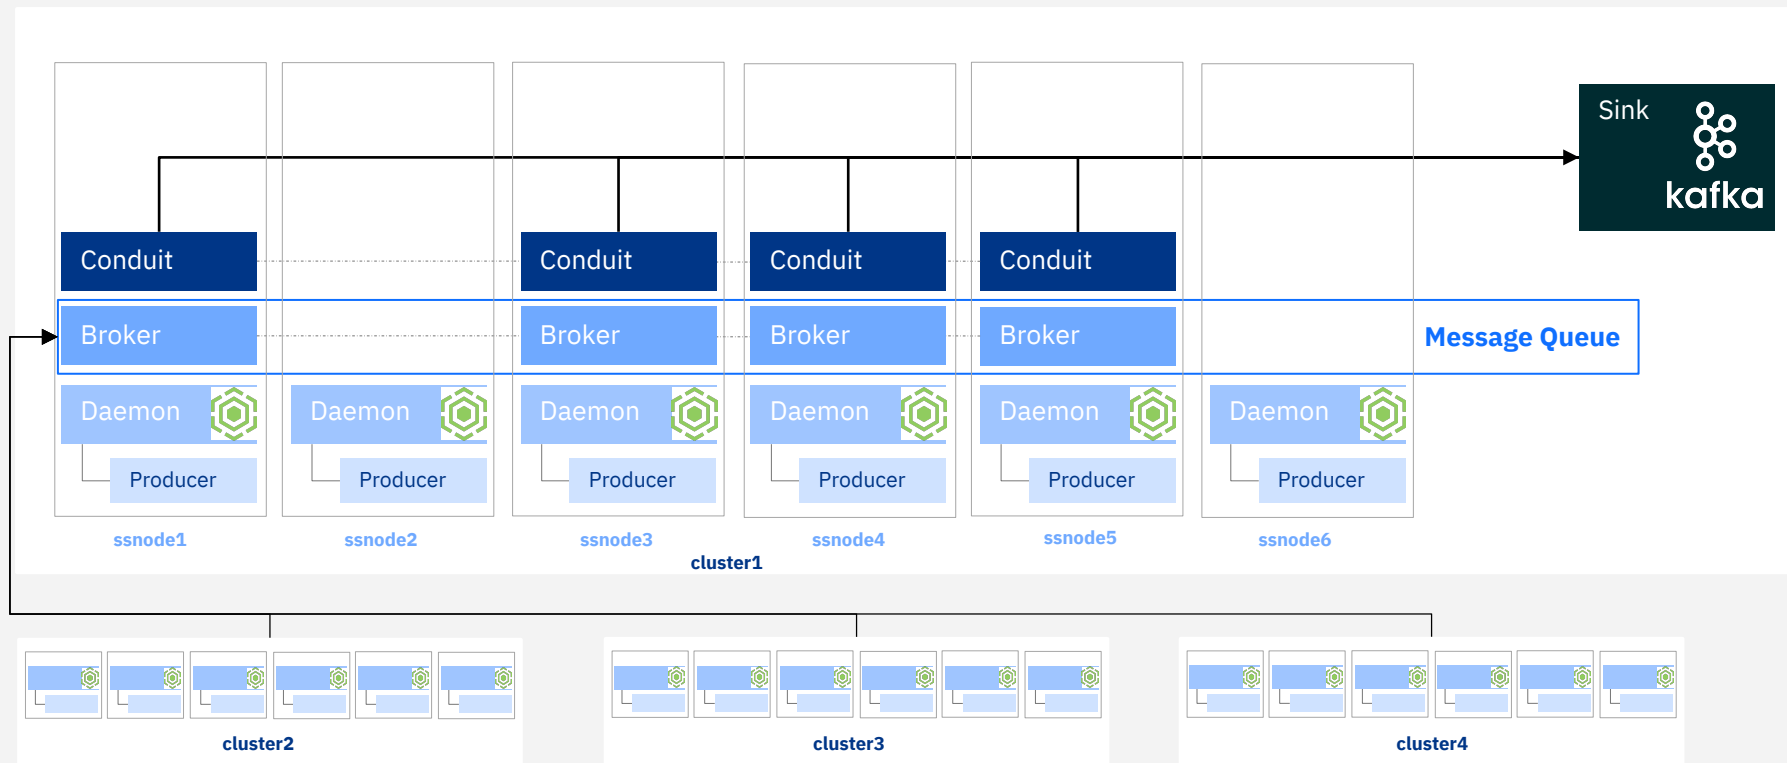
Rich Event Metadata

Take action, Respond to notifications:

- Automate Data Workflows
- Curate, Index, Discover
- Track Data Provenance and Lineage



Architecture



Spectrum Scale policies used to configure the daemons as to which accesses/files need to be watched.

Events you can Watch

- IN_OPEN *
- IN_ACCESS
- IN_MODIFY
- IN_ATTRIB *
- IN_CLOSE_WRITE *
- IN_CLOSE_NOWRITE
- IN_CREATE *
- IN_DELETE *
- IN_MOVED_FROM
- IN_MOVED_TO
- IN_DELETE_SELF
- IN_MOVE_SELF *
- IN_IGNORED **
- IN_UNMOUNT **

* IN_ISDIR also when operation concerns directory

** Cannot configure, but will be reported

<http://man7.org/linux/man-pages/man7/inotify.7.html>

```
{ "WF_JSON": "0.0.1",  
  "wd": "1",  
  "cookie": "0",  
  "event": "IN_CREATE",  
  "path": "/gpfs/gpfs1/DSTtest",  
  "clusterName": "solar.gpfs.ibm.net",  
  "nodeName": "node1.ibm.com",  
  "nfsClientIp": "",  
  "fsName": "gpfs1",  
  "inode": "134333440",  
  "fileSetID": "0",  
  "linkCount": "1",  
  "openFlags": "0",  
  "poolName": "system",  
  "fileSize": "0",  
  "ownerUserId": "0",  
  "ownerGroupId": "0",  
  "atime": "2019-03-21_18:14:05-0500",  
  "ctime": "2019-03-21_18:14:05-0500",  
  "mtime": "2019-03-21_18:14:05-0500",  
  "eventTime": "2019-03-21_18:14:05-0500",  
  "clientUserId": "0",  
  "clientGroupId": "0",  
  "processId": "30432",  
  "permissions": "200100644",  
  "acls": null,  
  "xattrs": null,  
  "subEvent": "NONE"  
}
```

mmwatch

mmwatch *Device* **enable** { **-F** *ConfigFilePath* |

[**--fileset** *fsetname*] [**--watch-id** *WatchID*]

[**--events** {*Event*[,*Event*...] | **ALL**}]

--event-handler *handlertype*

--sink-brokers *BrokerIP:Port*[,*BrokerIP:Port*...]

--sink-topic *Topic*

[**--sink-auth-config** *Path*]

[**--degraded**] }

mmwatch *Device* **list** [**--events**] [**--watch-id** *WatchID*] [-Y]

mmwatch *Device* **list** **--watch-id** *WatchID* **--config** [-Y]

mmwatch **all** **list** [**--events**] [-Y]

mmwatch **all** **status**

mmwatch *Device* **status** [**--watch-id** *WatchID* [-v]]

Sink Authentication

- SELinux in enforcing mode is not supported.
- Root authority is required to run **mmmsgqueue** and **mmwatch**.
- The following TCP ports must be open on all nodes in the cluster:
 - 2181, 9092, and 9093 along with the range 2888:3888

Enable a watch

```
[root@cluster1 ~]# mmwatch fs0 enable -F myconf.conf
[I] Beginning enablement of Clustered Watch with newly created watch ID: CLW1553290256
[I] Verifying MsgQueue nodes meet minimum local space requirements for Clustered Watch to be enabled for watch: CLW1553290256
    Depending on cluster size, this may take some time
[I] Successfully verified all configured MsgQueue nodes meet minimum local space requirements for Clustered Watch to be enabled for watch:
CLW1553290256
[I] Verified the watch type is FSYS for filesystem fs0
[I] Successfully created Clustered Watch topic on the MsgQueue for watch: CLW1553290256
[I] Successfully added Clustered Watch configuration file into CCR for watch: CLW1553290256
[I] Successfully enabled Clustered Watch consumers for watch: CLW1553290256
[I] Successfully added Clustered Watch policy rules for watch: CLW1553290256
[I] Successfully enabled Clustered Watch: CLW1553290256
```

```
# cat myconf.conf
EVENT_HANDLER:kafkasink
SINK_BROKERS:nodel.ibm.com:9092,node2.ibm.com:9092,node3.ibm.com:9092
SINK_TOPIC:myExternalSinkTopic
SINK_AUTH_CONFIG:/home/secrets.txt
DEGRADED:false

# cat /home/secrets.txt
SINK_AUTH_TYPE:CERT
CA_CERT_LOCATION:/opt/kafka/SSL/ca-cert
CLIENT_PEM_CERT_LOCATION:/opt/kafka/SSL/clients/client.pem
CLIENT_KEY_FILE_LOCATION:/opt/kafka/SSL/clients/client.key
CLIENT_KEY_FILE_PASSWORD:DST4lyfE
```

Sink Authentication: NONE, PLAIN, SCRAM512, CERT

Use Cases


- File System Catalog for curating data – Spectrum Discover
- Workflow Automation
- File Sharing Portal
- Transparent Provenance for Analytical Pipelines



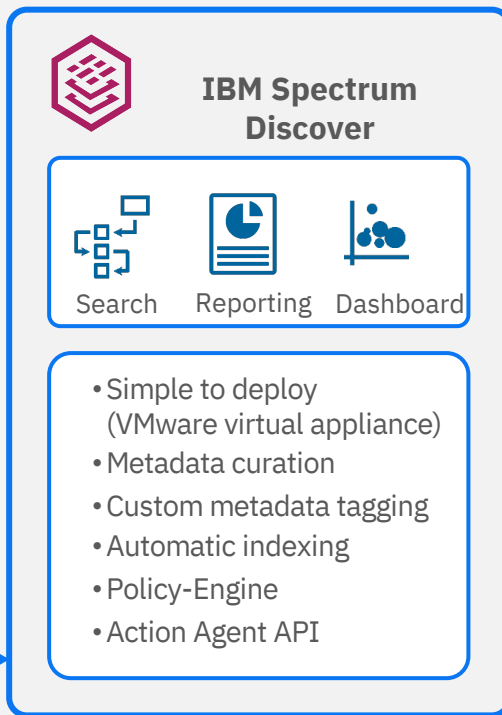
IBM Spectrum Discover



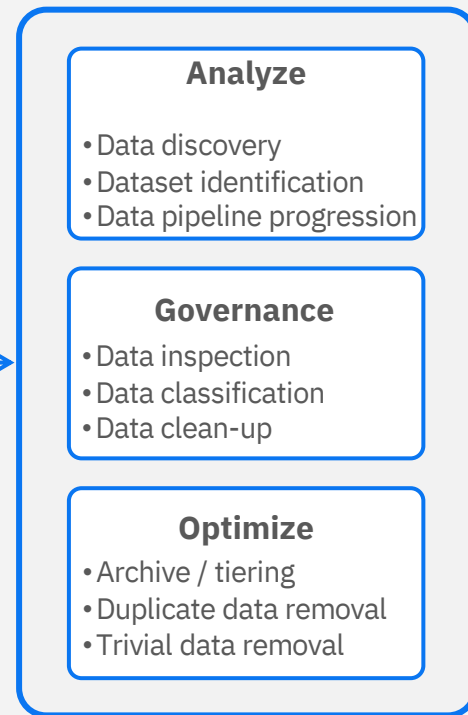
File and Object Storage



Data Insight



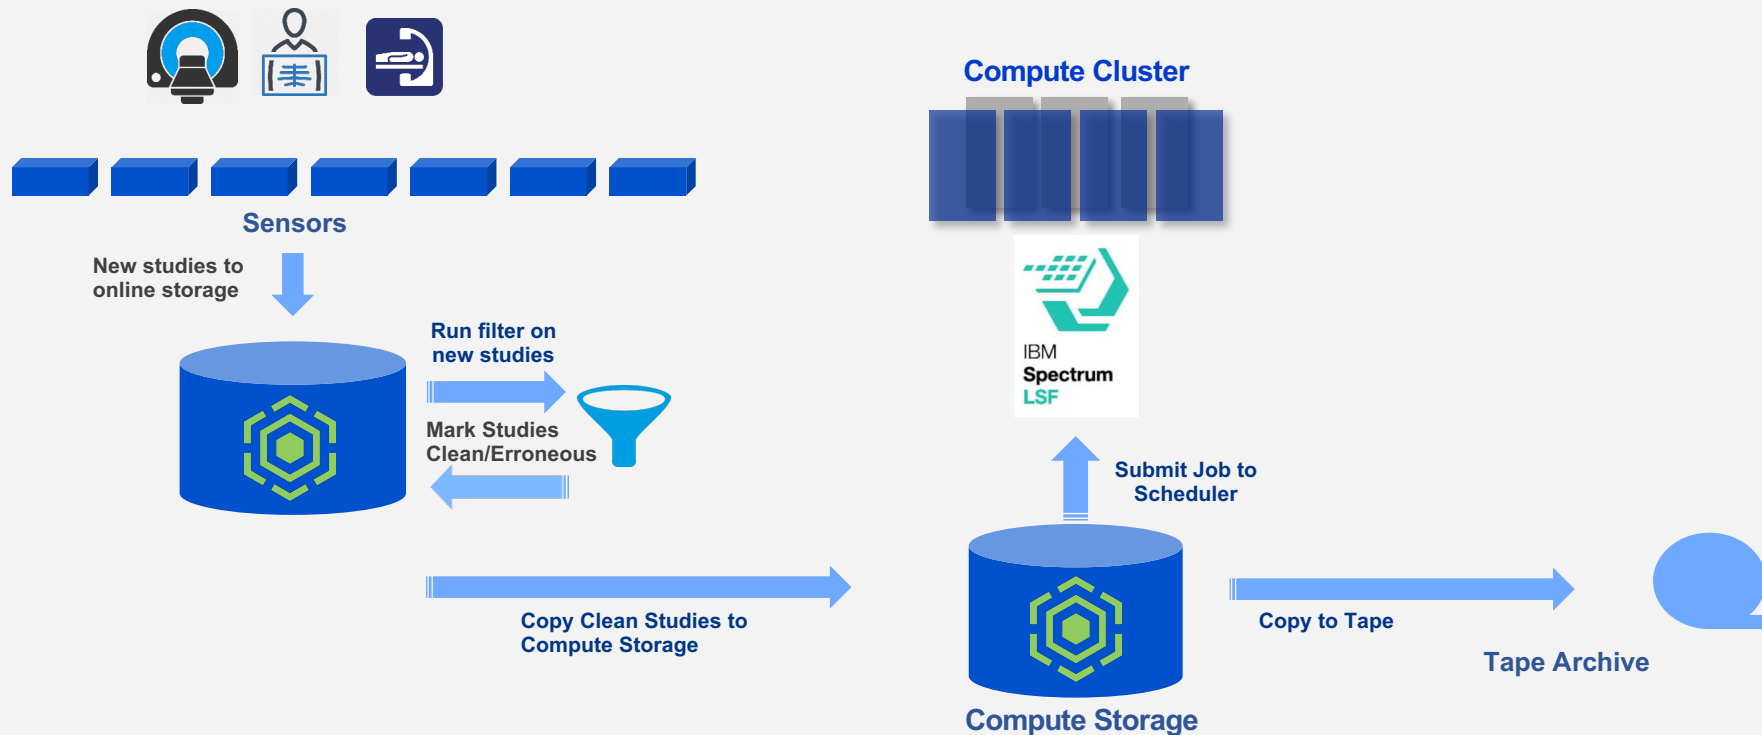
Data Activation/Optimization



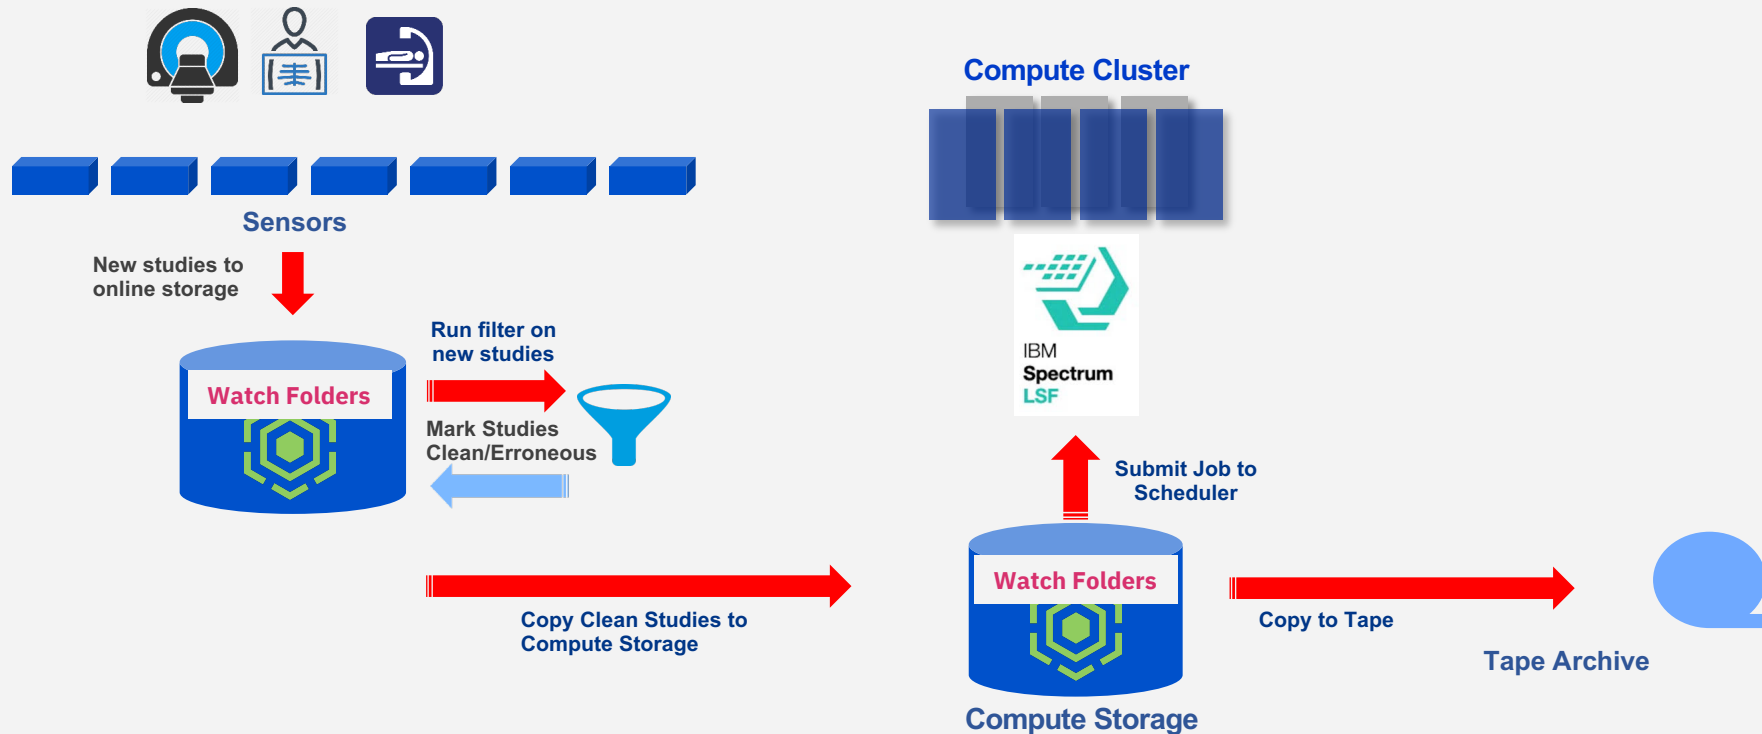
Event
Notifications

Use
Cases

Workflow Automation (POC)



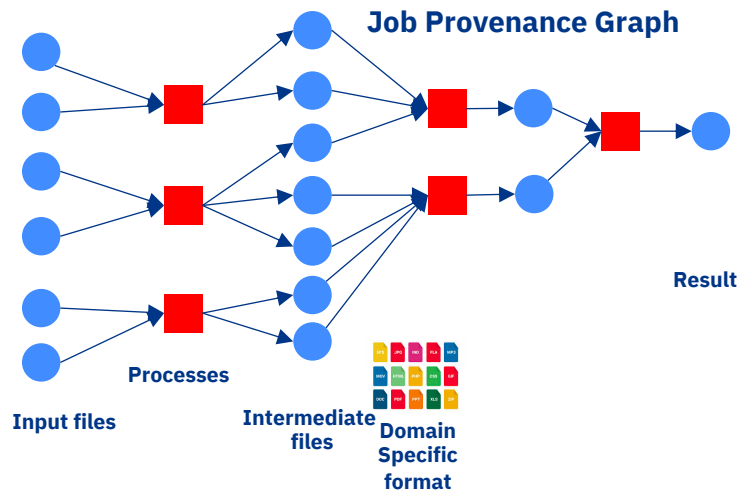
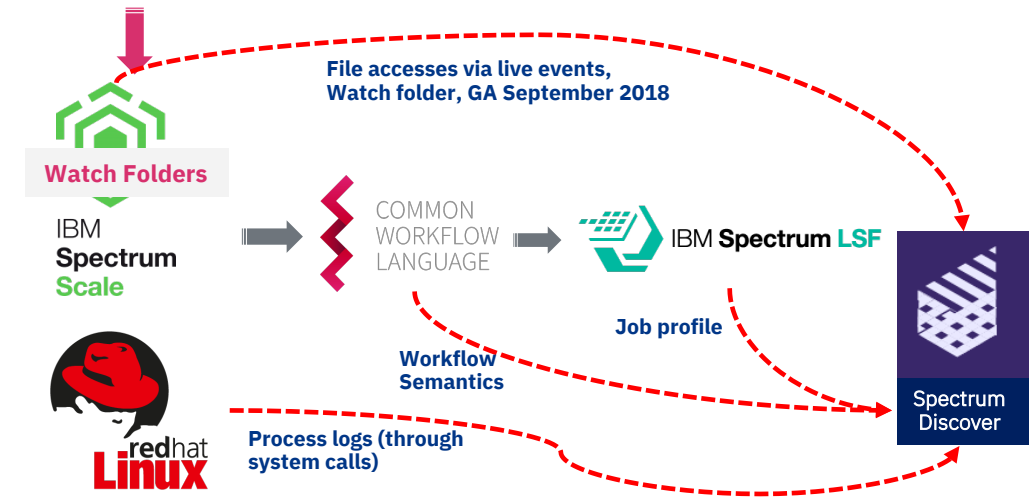
Workflow Automation




File Sharing Portal (POC)



Transparent Provenance for Analytical Pipelines (POC)



Live Demo (hopefully)



Thank you! Questions?

