

## IBM Spectrum Discover

Modern metadata management software for exabyte-scale file and object storage

With digital transformation comes rapid growth of unstructured data. According to a recent Forrester Analytics survey, the number of enterprises with a petabyte or more of unstructured data grew from 13 percent in 2016 to 38 percent in 2017.<sup>1</sup> Remarkably, this means, the percentage of companies with a petabyte or more of unstructured data nearly tripled in one year. So, it is no surprise that storage administrators are struggling to keep pace with rapid unstructured data growth—and that in response, they often find themselves simply throwing more storage at the problem.

---

### Highlights

- Event-notifications and policy-based workflows
  - Fine-grained views of storage consumption
  - Fast, efficient search through petabytes of data
  - Quickly differentiate mission-critical business data
  - Policy-based custom tagging
  - Extensible via Software Developers Kit (SDK)
- 

While storage volume is a challenge, limited visibility into stored data poses an even greater challenge for both storage administrators and users of large-scale unstructured data. Storage administrators often find that system metadata alone doesn't provide a fine-grained view of storage consumption and data quality, which is needed for effective storage optimization. Basic system-level metadata is also inadequate for data scientists, business analysts and knowledge workers who spend a significant amount of their time searching for data necessary to do their work. Data stewards also struggle to identify files and objects (records) that contain confidential or sensitive data.

To overcome these unstructured data challenges, large enterprises are turning to metadata management solutions that offer exceptional data visibility. Once organizations have a clear understanding of their unstructured data, they can optimize storage systems, mitigate risk and harness the value of unstructured data for competitive advantage and critical data insights.

## Improve unstructured data economics, governance, analytics and discovery

IBM Spectrum Discover is modern metadata management software that provides data insight for exabyte-scale unstructured storage. IBM Spectrum Discover easily connects to IBM Cloud Object Storage System and IBM Spectrum Scale to rapidly ingest, consolidate and index metadata for billions of files and objects, providing a rich metadata layer on top of these storage sources. This metadata enables data scientists, storage administrators, and data stewards to efficiently manage, classify and gain insights from massive amounts of unstructured data. The insights gained accelerate large-scale analytics, improve storage economics, and help with risk mitigation to create competitive advantage and speed critical research.

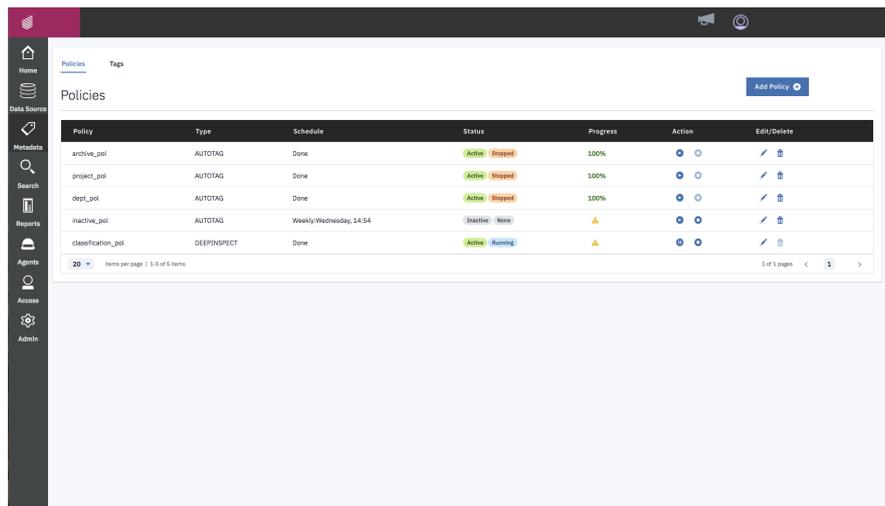
IBM Spectrum Discover highlights include:

- Event-notifications and policy-based workflows to automate metadata ingestion and metadata indexing at exabyte-scale
- Fine-grained views of storage consumption based on a wide range of system and custom metadata.
- Fast, efficient search through exabytes of data, resulting in highly relevant results for large-scale analytics.
- Ability to quickly differentiate mission-critical business data from data that can either be deleted or moved to a cheaper, colder tier.
- Policy-based custom tagging that enables organizations to classify and categorize data and align this data with the needs of the business.
- A Software Developers Kit (SDK) to build Action Agents that extract metadata from file headers and content, automate data movement and provide integration to open source software, such as Apache Spark, Apache Tika, PyTorch, Caffe and TensorFlow, which facilitates data identification and speeds large-scale data processing.

## Policy-based metadata tagging for granular data classification

IBM Spectrum Discover automatically captures system metadata from source storage systems, creates custom metadata from search results and enables extraction of keyword metadata from file headers and content using IBM Spectrum Discover’s Action Agent API. The result is a rich layer of file and object metadata that is managed using one centralized solution. Here’s how policy-based metadata tagging works:

With IBM Spectrum Discover, policies are used to automate actions that enrich record metadata. Users can apply policies to any set of records and can configure actions. For example, storage administrators can easily coordinate with departments to archive aging data. To do this, they use the solution’s Policy Engine that leverages the search function to find records owned by a department (for example, marketing) and that have not been accessed for a specified period of time (for example, more than one year). Then, they select a pre-defined “archive” tag from a drop-down list and the archive tag is automatically applied to the relevant group of files. Policies can be executed as one-time events, or they can be scheduled to run periodically. Any new tags created with policies are indexed and become searchable.



IBM Spectrum Discover provides an intuitive user interface for policy-based metadata tagging.

## Fast searching through billions of metadata tags enables rapid discovery of data assets

IBM Spectrum Discover provides both a search bar and a more advanced search pane to help users quickly find subsets of records that have been indexed. Search results are displayed in a columnar table that contains information correlated to search criteria. What a user can see or not see is determined using role-based access controls.

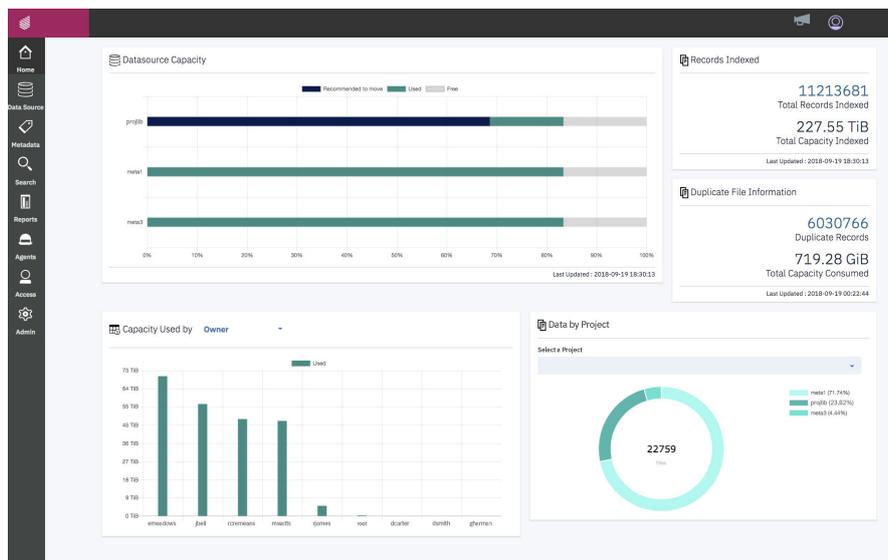
| Name                    | Datasource | Owner | Fileset | Location       | Size (Bytes) |
|-------------------------|------------|-------|---------|----------------|--------------|
| other_generic.14.tar.gz | ha         | root  | root    | /noR&B&H&L_C&V | 3429021928   |
| other_generic.15.tar.gz | ha         | root  | root    | /noR&B&H&L_C&V | 735109817    |
| other_generic.16.tar.gz | ha         | root  | root    | /noR&B&H&L_C&V | 58           |
| other_generic.17.tar.gz | ha         | root  | root    | /noR&B&H&L_C&V | 47           |
| other_generic.18.tar.gz | ha         | root  | root    | /noR&B&H&L_C&V | 58           |
| other_generic.19.tar.gz | ha         | root  | root    | /noR&B&H&L_C&V | 58           |
| other_generic.20.tar.gz | ha         | root  | root    | /noR&B&H&L_C&V | 47           |
| other_generic.21.tar.gz | ha         | root  | root    | /noR&B&H&L_C&V | 47           |
| other_generic.22.tar.gz | ha         | root  | root    | /noR&B&H&L_C&V | 47           |
| other_generic.23.tar.gz | ha         | root  | root    | /noR&B&H&L_C&V | 47           |
| other_generic.24.tar.gz | ha         | root  | root    | /noR&B&H&L_C&V | 47           |
| other_generic.25.tar.gz | ha         | root  | root    | /noR&B&H&L_C&V | 58           |
| other_generic.26.tar.gz | ha         | root  | root    | /noR&B&H&L_C&V | 58           |
| other_generic.27.tar.gz | ha         | root  | root    | /noR&B&H&L_C&V | 58           |
| other_generic.28.tar.gz | ha         | root  | root    | /noR&B&H&L_C&V | 58           |
| other_generic.29.tar.gz | ha         | root  | root    | /noR&B&H&L_C&V | 47           |
| other_generic.30.tar.gz | ha         | root  | root    | /noR&B&H&L_C&V | 47           |
| other_generic.31.tar.gz | ha         | root  | root    | /noR&B&H&L_C&V | 47           |
| other_generic.32.tar.gz | ha         | root  | root    | /noR&B&H&L_C&V | 47           |
| other_generic.33.tar.gz | ha         | root  | root    | /noR&B&H&L_C&V | 13           |
| other_generic.34.tar.gz | ha         | root  | root    | /noR&B&H&L_C&V | 109901271    |
| other_generic.35.tar.gz | ha         | root  | root    | /noR&B&H&L_C&V | 83210688     |
| other_generic.36.tar.gz | ha         | root  | root    | /noR&B&H&L_C&V | 91213755     |

IBM Spectrum Discover search results are displayed in a columnar table that contains information correlated to search criteria.

Users familiar with SQL syntax can enter a search string in the search bar. Or, IBM Spectrum Discover provides an easy-to-use search pane to filter records using predefined selection boxes. For example, the “File System” selection box allows users to select one or more source storage systems. The “Time” selector allows users to specify a range of time based on when records were last accessed. The “Size” selector allows users to identify records based on minimum and/or maximum file sizes. These and other search capabilities allow users to employ any combination of search boxes that best suits their needs.

## Dashboard and customizable reporting for record visualization

The IBM Spectrum Discover dashboard represents a user’s environment at a glance. What a user can see or not see is determined using role-based access controls. The dashboard contains widgets that graphically present information about records indexed in the IBM Spectrum Discover cluster allowing users to visualize their data environment. For example, the dashboard can show usage vs. capacity of their registered storage systems, information about potential duplicate files, and breakdowns of how capacity is being used by projects or departments.



IBM Spectrum Discover Dashboard includes widgets that show a user’s environment at a glance.

For users who want additional record detail, IBM Spectrum Discover provides customizable reports. Both summary and detailed reports can be generated. Summary reports aggregate and group information, such as record count or record capacity by different criteria, for example: object vault, file system or user. Detailed reports provide detailed information for each record in the system that matches a report’s filtering criteria.

1 “Global Business Technographics Data And Analytics Survey, 2017,” Forrester, June 2017.  
<https://www.forrester.com/Global+Business+Technographics+Data+And+Analytics+Survey+2017/-/E-sus3671#>

## IBM Spectrum Discover at a glance

| Single node            |   |
|------------------------|---|
| Memory                 | 228 GB  |
| CPU                    | 16 cores  |
| Storage                |   |
| Primary pool           | 3 TB local SSD or 3 TB SAN-attached SSD/Flash storage   |
| Backup pool            | 3 TB  |
| Network                | Single gigabit Ethernet or 10Gb Ethernet                |
| Multi-node (3 nodes)   |   |
| Memory                 | 256 GB  |
| CPU                    | 24 cores  |
| Primary pool           | 15 TB local SSD or 15 TB SAN-attached SSD/Flash storage |
| Backup pool            | 15 TB   |
| Network                | 10 Gb Ethernet  |
| Software prerequisites |   |
| VMware ESXi 6.5        |   |

## IBM Spectrum Discover capabilities

|                                    |  |
|------------------------------------|--|
| Continuous metadata ingestion      | <ul style="list-style-type: none"> <li>• Built-in connectors provide tight integration with IBM Cloud Object Storage and IBM Spectrum Scale</li> <li>• Event-notifications automate continuous metadata ingestion</li> <li>• Metadata indexing enables rapid data queries</li> </ul> |
| Systematic metadata curation       | <ul style="list-style-type: none"> <li>• Policy-driven workflows automate custom labeling</li> <li>• Custom data labels help pinpoint data for large-scale analytics</li> <li>• Ability to link system and custom data labels accelerates storage optimization</li> </ul>            |
| Real-time data insight             | <ul style="list-style-type: none"> <li>• Fast search locates highly relevant files and objects in seconds</li> <li>• Dashboards with drill-down chart elements simplify storage management</li> <li>• Customizable reports expedite audits and communication</li> </ul>              |
| Secure and extensible architecture | <ul style="list-style-type: none"> <li>• Role-based access control ensures only authorized access to data</li> <li>• Action Agent API supports integration with customer-developed and/or third-party software</li> <li>• Policy engine hooks enable automated workflows</li> </ul>  |

## Why IBM?

As an industry-leading provider of data storage products, IBM is investing in data management solutions that improve storage economics, data quality, data governance and data identification for large-scale analytics. IBM Spectrum Discover is a key aspect of the overall IBM data management advantage, and provides powerful metadata management software that brings visibility and classification to IBM Cloud Object Storage and IBM Spectrum Scale customers.

## For more information

To learn more about IBM Spectrum Discover, please contact your IBM representative or IBM Business Partner, or visit: [ibm.com/us-en/marketplace/spectrum-discover](https://ibm.com/us-en/marketplace/spectrum-discover)



---

© Copyright IBM Corporation 2018.

IBM, the IBM logo, and [ibm.com](https://ibm.com) are trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at <https://www.ibm.com/legal/us/en/copytrade.shtml>, and select third party trademarks that might be referenced in this document is available at [https://www.ibm.com/legal/us/en/copytrade.shtml#section\\_4](https://www.ibm.com/legal/us/en/copytrade.shtml#section_4).

This document contains information pertaining to the following IBM products which are registered trademarks of IBM Corporation:  
IBM®, IBM Cloud™, IBM Spectrum Scale™, Apache Spark™, Cloud Object Storage System™

All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

---