

### KEY BENEFITS

- **Evaluate independently** or within DataSphere
- **Estimate the Total Cost of Ownership** for your data
- **Calculate ROI** based on actual data from your target infrastructure
- **Visualize or architect** tiered capacity and usage


Today's IT leaders and organizations face a number of conflicting pressures, including guaranteeing performance levels, providing non-disruptive data migrations, meeting service expectations with existing budgets, and seamlessly leveraging low cost cloud storage. Primary Data's new Data Profiler tool helps enterprises manage these conflicting demands.

The Data Profiler can operate independently as a standalone application or within Primary Data's DataSphere platform. It analyzes how data is used across the storage infrastructure to evaluate data management strategies and project the savings possible by improving utilization across resources including NAS, cloud, or on-premises object storage resources. When deployed with DataSphere, enterprises can quickly put their customized data management strategies into action with automated data migration, tiering storage across their enterprise infrastructure and adding low cost cloud storage for cold data management.

The Data Profiler uses statistics from a target share(s) and gives the IT user the ability to model the existing environment to then build a new tiered architecture by optimizing the use of existing storage, and replacing or adding new storage to meet budgetary and planning goals. In addition, the Data Profiler highlights:

- Total cost of data ownership within the target storage environment
- Cost Savings vs. previous or reference configurations
- Benefits of a tiered storage configuration
- Cost, amount and the number of files per storage tier
- The advantage of low cost object storage for cold less used data

The Data Profiler can collect statistics directly from a store independently, or storage not directly managed by DataSphere, as well as stores under DataSphere's control. The Data Profiler has no impact on the data being analyzed.



**Figure 1.** The Data Profiler shows how three storage tiers break down by cost, creating a 50% savings on future costs.

## HOW DOES IT WORK?

The Data Profiler can passively read metadata statistics from the target storage arrays or stores and use this information to model and present the state of the data's environment. In addition, the Data Profiler tool is also available within DataSphere, and can profile both stores managed within or outside DataSphere's global namespace. This means stores can be analyzed before determining where data can be best placed to meet performance and price objectives. The Data Profiler does not affect either the data being analyzed, nor clients accessing the data.



- 
- Assimilate existing shares
  - Does not impact existing data
  - Can be used on shares that are in use however it is recommended to use on a clone of a share

- 
- Assimilate existing shares
  - Does not impact existing data
  - Can be used on shares that are in use however it is recommended to use on a clone of a share

- 
- Assimilate existing shares
  - Does not impact existing data
  - Can be used on shares that are in use however it is recommended to use on a clone of a share

## HOW DO I GET IT?

While the Data Profiler is a module that comes integrated with every license of DataSphere, it can also be run as a stand-alone tool. This module can be run either as a part of a full production environment or as part of a proof-of-concept phase before DataSphere is deployed, building a plan for how to get the most out of DataSphere's intelligent data management.

## ABOUT PRIMARY DATA

Primary Data develops intelligence and automation software for enterprise data management across on-premises IT infrastructure and into the cloud. Its DataSphere platform combines metadata management and machine learning to move the right data to the right place at the right time across a global namespace, automatically and without application disruption. DataSphere makes heterogeneous data stores simultaneously available to all applications, enabling enterprises operating at petabyte scale to easily manage billions of files, automate data migration, integrate the cloud, and scale out NAS performance while getting the most value out of infrastructure investments on a per-client, per-file basis. To learn more, visit us at: [PrimaryData.com](http://PrimaryData.com), follow us on [Facebook.com/PdDataSphere](https://Facebook.com/PdDataSphere), or Twitter at [@Primary\\_Data](https://twitter.com/Primary_Data).