



Questionnaire

ENTER
CATEGORY

Dear Data/Information Lifecycle Management Solution Provider,

The following is a questionnaire for our use in understanding and evaluating the numerous solutions that are collectively called data or information lifecycle management solutions by vendors. Our working definition of these types of solutions is "a software and/or hardware solution for migrating data through the storage infrastructure using automated policies that match data content and/or access characteristics to storage platform capabilities and costs." If you would like to suggest amendments or revisions to this definition, please make your suggestions as part of this questionnaire.

Please respond as per the cover letter attached to this questionnaire. You can edit this document directly to provide your information.

Usage Scenarios

Since the capabilities of these solutions may vary, we are going to assess their effectiveness in the context of different usage scenarios.

Capacity Utilization Efficiency. These users mostly use a D/ILM solution to place data on platforms or media that are least costly, most capable and best suited to their access characteristics and "inherited" requirements in terms of retention, security, and criticality. Automated data migration may also be seen as a strategy for reducing TCO by reducing or capping administrative staff requirements.

Capacity Allocation Efficiency. These users are primarily interested in D/ILM as a mechanism for sharing capacity in the most efficient way and for eliminating unnecessary replication or junk data – often to defer additional hardware purchases. Automated capacity allocation may also be seen as a strategy for reducing TCO by reducing or capping administrative staff requirements.

Disaster Recovery and Business Continuity. These users seek to use a D/ILM solution to segregate data by its restoration priority in the wake of a disaster and to route data to the appropriate protection process (mirroring or continuous data protection, snapshot or disk-to-disk replication processes or tape backup). These users may also seek to leverage D/ILM to provide input to change management processes designed to keep plans up to date with burgeoning data.

Questions

Development Status and Objectives. Describe your vision of D/ILM and identify what features or functions your platform currently provides and what features or functions are still in development.

Vision

Since Arkivio Inc. was formed in December 2000, our founders understood that the key to effectively managing storage resources is to combine data intelligence with automated management. We recognize that data and information have a value and that this value can vary over time to a business. Likewise, storage resources vary in terms of their cost, performance, and the level of data protection they provide. Our approach to Information Lifecycle Management (ILM) is to provide our customers with the solutions and methodologies

to understand the changing value of their data in order to automate placement and migration of information to the most appropriate storage resource, while optimizing the cost, performance, and availability within the storage infrastructure.

Key Features

There are six key elements or functions of an ILM solution: data/storage discovery, logical organization of resources, classification/valuation schemes, creation of data management policies, simulation of policy results, and evaluation (monitoring and reporting) of the entire ILM process. Our flagship product, ARKIVIO® auto-stor, is the first storage management product of its kind to integrate all these elements in one comprehensive ILM solution. Below is a more detailed description of the key features of the ARKIVIO® auto-stor solution.

- Discover:** scan and collect detailed metadata on data files and storage resources, as well as statistics on storage utilization and data usage patterns across heterogeneous file systems without installing server agents
- Organize:** logically group data and storage resources based on administrator-defined criteria to simplify reporting and centralize policy deployment across multiple systems (i.e. group files based on common attributes such as type, age, etc.)
- Classify:** prioritize data and storage resource groups according to their business value. For instance, rank a file group as high priority when the data is extremely important to the business, such as financial data summarizing the current fiscal quarter results. In contrast, a low priority file group might include data that is inconsequential to current business operations, such as temporary or log files. Storage (volume) groups are classified based on the characteristics of the device, such as device type, cost, make/model or performance characteristics.
- Create:** Create and schedule the data management policies that leverage the resource groups and classifications. Automate data management actions such as migrate, move, copy, and delete based on changing data retention and availability requirements and administrator-defined policies.
- Simulate:** Simulate the different policies and test their expected results before actually executing them. Various what-if scenarios can be run to determine the outcome and impact of the planned policies. For instance, the administrator could identify the capacity freed by a data migration policy.
- Evaluate:** Continuously monitor and alert on storage utilization levels, data usage patterns, and policy effectiveness to proactively resolve capacity management issues and monitor ability to meet department SLAs.

All of this capability for ARKIVIO® auto-stor is available today across heterogeneous platforms and storage topologies.

Development Partners. List the vendors with whom you are working to deliver your D/ILM solution functionality.

Arkivio has formal development partner agreements with EMC, IBM, Microsoft, Nexsan, NetApp, Sun, iLuman and Veritas. In addition, Arkivio is working on a variety of customer-driven product integration projects with leading solution providers.

Data Naming Scheme. What is your method for identifying data storage requirements or characteristics and for using those requirements or characteristics in building policies?

ARKIVIO® auto-stor scans and collects detailed metadata on data files and storage resources, as well as statistics on storage utilization and data usage patterns across heterogeneous file systems without the need to install server agents. The administrator can

create storage (volume) groups based on common characteristics such as cost, manufacturer/model, available capacity, or storage type. Additionally, file groups can be created based on attributes such as last access/modified, size, owner, or type and a priority can be assigned to each group based on its value to the business. These resource groups and their prioritization are easily incorporated into policies using ARKIVIO® auto-stor's wizard-based interface. Over time any new files created by users or new storage purchased by IT are automatically updated into the appropriate global resource groups to ease overall administration, and any existing data management policies are dynamically updated.

Access Frequency. Identify the mechanism by which your solution identifies how frequently specific files or datasets are being accessed and whether and how access frequency is used in migration policies.

ARKIVIO® auto-stor enables administrators to utilize multiple file attributes when logically organizing files into resource groups. Last accessed and last modified times are attributes that can be used. When included in policies these attributes are used to select which files should be migrated. In addition, the ARKIVIO policy engine uses these attributes in computing Data Value Scores (DVS) to determine the priority order for data migration. (see page 6 for more information on DVS)

While number of accesses to a file could also provide useful information, it would still need to be examined as a rate over time. For example, File A may have 1,000 accesses and file B may have 30 accesses, but file A may not have been accessed for 6 months and File B may have been accessed once per week for the last 30 weeks. Currently, file systems don't keep track and export file access counts, however, we have requested this capability from several vendors. If and when they do offer this, ARKIVIO® auto-stor will be able to readily take advantage of this information in file grouping and policy actions.

Storage Platform Characterization. Explain the mechanism that your solution provides for characterizing the performance capabilities and costs of specific hardware platforms for data storage so that this information can be used to target the appropriate storage platforms as destinations for automatically-migrated data.

ARKIVIO® auto-stor enables the administrator to classify one or more storage volumes based on criteria such as size, cost, performance, manufacturer, model, or device type (i.e. SAN, NAS, DASD, JBOD etc.). As the administrator selects the specific volumes to include in a volume group, ARKIVIO® auto-stor automatically suggests a cost for a volume based on the manufacturer model. The administrator is also provided with the option of over-writing the default costs with whatever cost/GB is applicable to their environment. The policy engine uses these attributes in computing Storage Value Scores (SVS) to determine the most appropriate target volume for data migration.

End of Useful Life. Explain how your platform facilitates the automated removal and clean-up of data that has outlived its useful life and restoration of freed capacity for use by applications.

Automated removal and clean-up can be achieved by data movement (i.e. moving data from primary storage to archival storage) or permanent deletion of files. The automated clean-up policies are specified based on criteria such as last access/modified time, file type, size, owners, or even directory (paths) locations. The administrator has the option to delete these files automatically on a scheduled basis or when the capacity of a volume reaches a certain threshold. The delete policy can be defined to automatically perform the clean-up or to do so only after the administrator has verified the list of files that the ARKIVIO® auto-stor Policy Automation Engine (PAE) intends to remove.

As an example, if the customer is concerned about copy right infringements in regards to MP3 files, they could implement a policy that automatically seeks and deletes this type of data. The file deletion policy is specified to delete all files matching a file group's criteria or

only enough files matching the criteria to relieve an over-capacity condition. Other examples of delete policies that may be applicable include removal all temp and log files or to remove files that have past their required retention timeframe according to government regulations.

Policy Articulation. Describe how policies are created and how they are applied to existing data.

The process of creating a policy is as simple as identifying the action desired (i.e. migrate, move, copy, or delete), the data and storage group the policy is to take action on and when the action should be taken. Using an intuitive, wizard-based management interface administrators follow a simple five-step process outlined below.

Step 1: Select source volumes or volume groups

Step 2: Select file groups on which the policy should apply

Step 3: Select target volume groups or a combination of target volume and directory

Step 4: Specify a trigger event (i.e. maximum capacity threshold exceeded)

Step 5: Schedule the policy (Policies can be scheduled to run only once or on a recurring basis. Policies may also be set to run automatically based on the specified condition in Step 4 only).

Administrators can deploy multiple policies on the same volume group and file group and then assign a priority to each policy so they are run in the proper order. Data movement across tiered-storage can be achieved by creating multiple interacting policies on the same resource or data set to manage the data as its classification changes over time. For example, six month old data can be migrated from primary storage to secondary storage. A year later the data can be completely removed from the online archive and saved on tape media. To improve data availability, a copy policy can be applied to the same data set. At the time new files are created they can be copied to a backup server, secondary storage, DR site, or tape.

Device Support. What storage devices does your product support? Are there any proprietary devices (controllers, arrays, HBAs, switches, SAN topology, virtualization products, etc.) that are required for your solution to work?

ARKIVIO® auto-stor operates at the volume and file system level and is independent of hardware devices and storage topologies. No proprietary devices are required.

Resource Consumption. How much bandwidth and server CPU “overhead” is introduced by your solution (e.g., to support polling processes, migration processes, agent processing, access frequency counting, etc.).

ARKIVIO® auto-stor has a modular architecture that is designed to minimize resource requirements on the managed hosts, storage, and network. Discovery, data collection and policy execution can be done without the need to install server agents. These tasks can be scheduled to run as background processes or during off-peak hours. For example, volume capacity information can be collected every 15 minutes (configurable); whereas detailed data usage information can be collected on a daily basis. The customer can also choose to have a server agent reside locally on a managed device with minimal overhead (avg. 3 - 5% CPU load). In addition, any policy action (e.g. data migration) is done as a background process, does not take up significant CPU load, and is schedulable during off-peak hours to minimize any impact on the network.

Flexibility. Does your solution have multi-vendor support? Can data, once integrated into your D/ILM scheme be migrated readily between your solution and other solutions in this space? Is your solution interoperable with other solutions in this space?

ARKIVIO® auto-stor operates at the volume and file system level and is independent of hardware devices and storage topologies. The solution can be used to facilitate the

movement within a particular vendor's storage platforms or between multiple vendors' platforms. ARKIVIO® auto-stor can also be used to move data between storage topologies (DAS, NAS, SAN).

ARKIVIO® auto-stor interoperates seamlessly with other data management applications (i.e. backup, archiving, and virus protection) without disrupting existing processes. For instance, the periodic execution of virus protection and backup software will not cause unwanted recalls of migrated data. ARKIVIO® auto-stor can also migrate data to a media management server that, in turn, stores the data to other storage (disk, tape, or optical). In addition, ARKIVIO® auto-stor's alerting mechanism integrates with other monitoring tools and systems management frameworks.

Speeds and Feeds. How should a prospective customer compare the relative performance of competitive solutions in this space? What are the appropriate performance measures or metrics to use in evaluating competitive solutions?

There are several performance metrics that customers can use:

1. Access latency. What is the latency when the user tries to access a migrated file that is no longer stored in the primary storage?
2. System performance overhead. What are the changes to throughput as measured by industry benchmarks such as NetBench and SPEC when the ILM solution is installed?

Interconnect Support. List the network or fabric interconnects that you support. Is there an optimal interconnect for data movement for your solution? What about for management?

ARKIVIO® auto-stor supports the management of storage resources on any network interconnect, including IP, Fibre Channel, and iSCSI. Data movement and management control information are transferred on the IP network.

Protocol Support. List the storage networking technologies supported, including FCP, Ethernet, iSCSI, FCIP, iFCP, Parallel SCSI, SAS, and network file system protocols (NFS, CIFS, HTTP, DAFS).

ARKIVIO® auto-stor operates at the volume and file system level and therefore supports any storage networking technology. Our software supports industry standard file system protocols including CIFS and NFS (both v2 and v3 for UDP and TCP).

Data Type Support. List the data types supported by your solution, including file types and file systems, databases, and hybrids such as email. Comment on the granularity of your D/ILM solution: will it support the migration of database components or subsets, subsets of email files, etc.?

ARKIVIO® auto-stor supports any non-structured file type for both CIFS and NFS. The ARKIVIO data management policies can also act on pst files and e-mail attachments. We also integrate with leading e-mail archiving applications such as iLumin and IXOS. Entire database or log files can be managed but not tables or records within the database.

Cost. List the cost of your solution or provide some means for calculating cost for a specific environment.

The ARKIVIO pricing model is based on the philosophy that an ARKIVIO customer should only pay for the specific use and value they derive from the deployment of our technology. The ARKIVIO technology consists of reporting, classification, and monitoring (ARKIVIO® auto-view) and automated data management (ARKIVIO® auto-stor). Arkivio recognizes that many of our customers wish to deploy the reporting and monitoring capability across the entire environment, but only desire the automation management capability on specific devices

within their network; therefore, we have developed a pricing model that easily and cost effectively allows them to accomplish this objective.

The cost per TB for ARKIVIO® auto-view is exclusively based on the amount of TB that is being monitored regardless of number of servers. The cost for ARKIVIO® auto-stor is only applied to any volume of storage that is utilized as a repository for data that is migrated, copied, or moved. This pricing model has proven to produce a very strong ROI model for ARKIVIO customers and prospects. The MSRP for ARKIVIO® auto-view starts at \$2,500 and ARKIVIO® auto-stor starts at \$6,000 per TB. As the volume of storage under management increases, the price per TB decreases at a rate of about 15% for every 10 TB.

Another key component of the ARKIVIO solution is the ARKIVIO Central Server (ACS). The ACS is broken into three separate elements: the ACS for auto-view, ACS for CIFS, and ACS for NFS. Each of the ACS options is priced at \$5,000. Therefore, if the customer only requires ARKIVIO® auto-view, the total ACS price is \$5,000. If the customer would like to automate data management (ARKIVIO® auto-stor) for either CIFS or NFS environments, the total ACS charge is \$10,000. However, if the customer would like to automate data management across both CIFS and NFS simultaneously; the total ACS charge is \$15,000.

Arkivio utilizes industry standard rates of 18% of MSRP for standard support and 24% for advanced 24x7 support.

Standards. List any relevant open standards upon which your product is built.

File system protocols (CIFS, NFS), network protocols TCP/IP, and network management protocols SNMP.

Other features. List other pertinent aspects of your solution.

Data Value Scores / Storage Value Scores

The ARKIVIO Policy Automation Engine (PAE) utilizes patented algorithms to compute a Data Value Score (DVS) for the file groups and a Storage Value Score (SVS) for the volume groups to determine the most appropriate files and volumes to include in data management actions. DVS scores are calculated based on the detailed parameters of the data, such as age, size, last modified, last accessed, etc. and the business valuation that the administrator specified for the different file groups. Similarly, SVS scores are calculated based on volume attributes, such as cost or percent utilized. Based on these scores, the PAE intelligently selects the most appropriate source and target volumes and end user files, which should be migrated, moved, copied or deleted to achieve the desired storage management objectives. This ensures, for example, that high-value data is matched to high-performance, highly available storage, while data that is lower in value (i.e. less frequently accessed) is matched to more costs effective storage (i.e. ATA RAID)

ARKIVIO® auto-stor LSA / ARKIVIO® auto-stor RSA

Arkivio has developed two unique solutions to address a variety of data management requirements. ARKIVIO® auto-stor LSA which is an agent-based ILM solution and ARKIVIO® auto-stor RSA, which does not require installation of agents. Although these solutions offer identical features in the areas of discovery, monitoring and reporting, and data/storage classification, they differ in their specific approaches to automated management.

ARKIVIO® auto-stor LSA (Local Server Agent) relies on the deployment of server agents installed directly on managed hosts to accomplish management and monitoring tasks. Currently, Arkivio provides local agents for Windows NT®, Windows® 2000, Windows powered NAS, and Solaris™ 2.8. The LSA approach allows for the use of “tag” or “stub” files to be left behind after migration, which provides complete transparency for user and application access to migrated data. In addition, having the server agents reside locally on the managed devices allows for near-real-time performance for the reporting and management of data.

ARKIVIO® auto-stor RSA (Remote Server Agent) utilizes remote server agents to manage and monitor any NFS or CIFS compatible storage without requiring agents to reside locally on managed hosts. The RSA product allows for the monitoring and management of any platforms that support CIFS or NFS file systems, which includes all flavors of Windows® and UNIX™ operating systems as well as proprietary NAS solution such as NetApp's ONTAP and EMC Celerra®. Using standard file system features, such as CIFS shortcuts and NFS symbolic links, the RSA product is able to migrate data in a way that allows it to be accessed from the original source location, while preserving compatibility with existing server and client platforms. Because the RSA product performs all management from a remote platform, it is completely out of the data path. Access to data on any managed platform never depends on the availability of any component other than the resident storage.

Miscellaneous: Please note any additional information that you think would be worthwhile for prospective customers to consider about your solution or other solutions in this space.

Policy Simulation

Once data management policies have been defined, the ARKIVIO solution is able to simulate the different policies and test their expected results before actually executing them. Various what-if scenarios can be run to determine the outcome and impact of the planned policies. For instance, the administrator could identify the capacity freed by a data migration policy. Additionally, the same detailed set of reports which are available for implemented policies are available for simulated policies.

Resource reporting and capacity planning

The reporting and trending analysis provided by the ARKIVIO® auto-stor solution enables customers to evaluate current storage resources, implement intelligent policy management, and monitor policies for effectiveness. This can be integrated with the logical groups of data and storage resources to specifically tailor reports to customer's priorities.

In addition, resource reporting can be used as a tool for capacity planning, to ensure that future storage needs are met. Resource reporting provides information such as the following:

- „ Types of files requiring the most storage space
- „ User groups using the most storage
- „ Number of stale or junk files in current storage
- „ Under or over utilized storage

