Questionnaire

ENTER CATEGORY
Comparative Review

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.

Questions

Development status and objectives. Describe your vision of Data/Information Lifecycle Management (D/ILM) and identify what features or functions your platform currently provides and what features or functions are still in development.

Overview
The term Information Lifecycle Management (ILM) has a broad range of definitions. While its definition focuses mainly on storage, which is also where the term originated, the rest of the industry is just now beginning to include content management as part of the definition well.

This questionnaire will outline IBM's solutions for ILM as defined in this questionnaire, but will take a much broader approach to the concept of aligning the value of data with business priorities. This extends to the ability to integrate business policies and processes with the larger IT infrastructure that supports the business.

IBM's History of Delivering the Foundation for ILM
Before the term information lifecycle management was invented, IBM was delivering solutions for integrated information management and storage to manage the lifecycle of data from creation to disposal. For example, IBM introduced hierarchical storage management software for mainframes (DFSMS) in 1974. Over the years, IBM has worked with its vast customer install base to refine the original HSM concepts. One particularly important addition to HSM was data migration, the automated movement of data from older media types to newer ones. Frequently, a piece of data outlives the usefulness of the media it is stored on. Today, the transition from older media to newer media can be automated, allowing organizations to quickly gain new cost benefits.

Additional highlights of IBM's deep history and leadership in ILM include the advent of hierarchical storage management software for open systems in 1993; the delivery of the IBM Content Management solution in 1988; and the full integration of content management with storage management technologies in 1996.

More recently, IBM acquired Tarian Software, Inc. in 2002 to extend the content retention capabilities of IBM DB2 Content Manager and IBM Tivoli Storage Manager to address the need for certified records management. IBM Records Manager can be embedded into any business application and adds true records retention capabilities such as declare and classify, it also passes destruction control over to the records administrator of the organization. Unique to IBM is the fact that each of these solutions for information management are tightly integrated with each other and with a common repository for data management for a seamless solution to manage the lifecycle of the data.

Managing the Value of Information
The management and protection of information is one of the most important tasks facing IT organizations today. At the same time it is growing in complexity. Information's value to the business can fluctuate based on the type of data and the data's point in its lifespan. The value is not only in optimizing the cost and performance of storing data, but correctly managing information based on its value to the business and leveraging that data to help meet business objectives. It can even be appropriately leveraged to provide a competitive advantage.

Effectively managing the value of information requires a solution that:

- Aligns data content and access needs with cost for both active and inactive data
- Supports key business processes through events based retention (including but not limited to regulatory compliance and litigation)

- Requires robust business policy management integrated from the business application through the storage layer
- Uses robust content management software to manage information as an asset from creation, through active and inactive use, to disposition
- Supports heterogeneous environments

IBM DB2 Content Manager provides a foundation for managing, accessing and integrating critical business information. It helps integrate all forms of content (email, document, web, image and rich media) across diverse business processes and applications, including Siebel, PeopleSoft and SAP. Content Manager also helps deliver powerful information and services by providing a single repository for all content which can also be managed.

As compliance requirements and regulations become more complex, customers require more control for content lifecycle. IBM Records Manager applies formal records management capabilities to electronic as well as physical records. It is integrated with IBM DB2 Content Manager to provide one repository for records across the organization, but is also available as an engine that can be embedded into any business application and leverage that application's repository. This is especially valuable for companies that want to keep their critical records separate from other systems. IBM Records Manager offers several advantages to recordkeeping such as a building blocks approach to designing and managing corporate file plans, time and event-based retention of records, and legal holds (suspensions) of records for pending lawsuits or potential audits.

IBM's Vision
The IBM vision for managing the value of information extends beyond the lifecycle of the data to the integration of business processes with IT. It further extends to new advances in technology to help customers arrive at an operating environment that if fully integrated across their organization with the resiliency to respond to changing business needs.

Aligning Storage Cost to Information Access and Value
IBM continues to deliver solutions that address the need to manage the lifecycle of data while addressing information's value to the business. IBM also offers a broad range of storage server capabilities with cost variability. The IBM TotalStorage hardware infrastructure solutions provide specialized storage servers to accommodate the changing value of information. These range from the IBM TotalStorage Enterprise Storage Server to the mid-range FAStT Storage Server, with a choice of high-performance fibre channel or low-cost Serial ATA (S-ATA) disk expansion. They also range from the flexible IBM TotalStorage Virtual Tape Server to the enterprise-class IBM TotalStorage Enterprise Tape Library with its 3592 Enterprise Tape Drive and to the IBM TotalStorage Ultrascalable Tape Library with its industry-standard Linear Tape Open (LTO) tape drives**.** Additionally, IBM's Write Once Read Many (WORM) media technology for the IBM TotalStorage 3592 helps secure tape drives from overwriting data.

Upon creation, data should be classified and placed on storage with the appropriate cost and performance. The lifecycle for data (a file, data set or any data object) begins the moment the data is created. To effectively categorize and place data as it is created, the allocation process must be "intelligent" and able to act on a set of customizable policies that place information into an appropriately priced pool of storage. As previously mentioned, IBM introduced this pooling capability and the policies that govern data placement on mainframe systems in 1988 and subsequently introduced this on open system platforms as well.

As data becomes less active or less valuable, it should automatically be moved. As data becomes less active or less valuable to an organization, or for TCO or SLA reasons, managing the data effectively requires software that can move the data to less expensive storage. As data goes through lifecycle changes, it may become appropriate to move it to

progressively lower-cost storage. The concept of *hierarchical storage management (HSM)*, invented by IBM, was introduced on mainframe systems 30 years ago and on open-system platforms nearly 16 years ago. Today, HSM techniques are available from IBM to manage the storage and movement of all kinds of data based upon management policies. This includes moving individual files on mainframe, moving open system file systems and moving historical e-mail and ERP data. This concept extends further to any kind of data object which embodies compliance data archives and backup copies of data that form the foundation of a business-continuance plan. Because the vast majority of all data maintained and managed by an IT organization varies in value and activity during its life, these HSM capabilities are critical to effective ILM.

Managing the value of information requires integration with business processes. Line-of-business (LOB) managers are increasingly implementing comprehensive content management solutions that can help them to:

- Search for and find information
- Reduce business process cycle times
- Improve worker collaboration
- Manage content on the Web

Corporate officers, driven by the need to show business control are also looking to content management systems to effectively manage the proper retention and disposition of business information and to comply with government regulations. To the CIO, these business imperatives represent additional data to be stored and managed. As an industry leader in content management, IBM provides the building blocks that help LOB managers meet their business imperatives. IBM has also fully integrated its content management offerings with its storage infrastructure and HSM offerings so that IT managers can cost effectively store this new information without the need to retrain administrators or implement a separate storage infrastructure.

Only IBM has the comprehensive solutions and services to help customers build an IT infrastructure that is flexible enough to change as quickly as business requirements. The ultimate vision for this type of flexible infrastructure can be described as an enterprise whose business processes — integrated end-to-end across the company and with key partners, suppliers and customers — can respond with speed to any customer demand, market opportunity or external threat.

Key to managing the value of information in this environment is the integration of business processes and information. This can mean leveraging formerly fragmented data and making it fully accessible across an organization to enable more accurate and timely business decisions. For example, content management solutions can be used to integrate content for suppliers, customers and employees while records management solutions can be used to meet compliance and regulatory demands by linking retention based business processes with application content management and storage management.

Second to this is aligning the management of the IT infrastructure with business needs. This extends to optimizing storage management to protect data, enhance utilization, and reduce costs. This can be accomplished across heterogeneous environments to reduce the complexity in IT infrastructure.

Finally, the IBM vision for leveraging the value of information helps organizations innovate the business to differentiate themselves and deliver new value while making better use of the resources they have and becoming more productive.

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

IBM offers a complete set of building-block technologies that help users to effectively implement basic, entry-level or highly scalable ILM. IBM also works with a wide variety of Business Partners to expand the breadth of customer solutions for leveraging the IBM infrastructure building blocks.

The IBM Content Manager application development toolkit provides a foundation for business partners and customers to efficiently integrate content into existing applications and business processes. This includes object-oriented APIs in C++ and Java, both visual and non-visual, a component toolkit for rapid application development, and portlets for easy access via IBM WebSphere Portal.

The Content Manager Object Oriented (OO) framework is an advanced object model based on Object Management Group (OMG) standards utilizing Persistent Object Services (POS) and Object Query Services (OQS). This framework of APIs or 'connectors' provides access to one or multiple heterogeneous content repositories in a single API call. With the complete OO toolkit, powerful end user applications, system administration applications, and workflow applications can be easily and efficiently developed to fully exploit the power of IBM Content Management capabilities.

To learn more about IBM DB2 Content Manager partners, see
ibm.com/software/data/cm/solutions.html

IBM Tivoli Storage Manager Partners
IBM Tivoli Storage Manager application programming interface (API) is the foundation of the open X/Open Backup Services API (XBSA) standard. It allows customers to choose any storage hardware and device vendor to use with content manager applications. Because the Tivoli Storage Manager API has been open and free for use by developers for nearly 10 years, many independent software vendors, including IBM customers, have developed to this API. Additionally, IBM has relationships with partners for Tivoli Storage Manager support for Oracle, SAP, Microsoft® Exchange, Microsoft SQL as well as the IBM DB2, IBM Lotus® Domino® and IBM Lotus Notes® applications.

There are many ways to integrate with the Tivoli Storage Manager API and we are currently finalizing agreements with a number of development partners for use of the API in conjunction with TSM for data retention. Many prominent content management application vendors are in the process of writing to this API.

To find out more about Tivoli Storage Manager partners, see
ibm.com/software/tivoli/partners/public.jsp?tab=connect&content=rft-validation-tsm

Data naming scheme. How do you identify data storage requirements or characteristics, and how do you use them in building policies

Storage management policies are required to manage the growing volume of business content such as electronic documents, facsimiles, scanned images and reports. IBM DB2 Content Manager is integrated with IBM Tivoli Storage Manager to provide hierarchical storage management by which information can be stored at central and/or distributed locations, automatically transitioned to lower cost storage over its lifecycle, and retained or dispositioned based on relevant records management policies.

File naming (including extensions and directory location), ownership, and create and last-reference dates can all be considered.

In open system environments, the IBM Tivoli Storage Resource Manager help can perform storage classification (see **ibm.com**/software/tivoli/products/storage-resource-mgr)

After the storage has been classified, capacity utilization, growth and access metrics can be calculated on the different classes of data. This gives IT administrators supporting information for designing the ILM policies that best fit their environments.

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.

The policies that drive file, data-set and data object HSM (migration) consider the classification of the file (the business policies governing the file) along with the file size, create and last-access dates, and the overall utilization of the current storage media.

Classification is important. Data classified as “business critical,” for example, may migrate more slowly and stay on disk media types longer. By contrast, data classified as “archive” may migrate quite quickly and go directly to lower-cost removable media such as tape. Data classified as “backup copies” may migrate first to low-cost disks so that it is available for rapid restore within the first several days — and then migrates to a lower-cost tape for longer-term storage and remote vaulting.

IBM software that implements these HSM techniques is designed to provide flexible policy-based automation for the wide variety of

data types that IT managers encounter. This software can rapidly adapt to changing hardware infrastructures. In open system environments, IBM Tivoli Storage Manager provides the software infrastructure for HSM. Building on this base, a heterogeneous variety of data types (files, records objects, etc.) are sent to and managed by this central hierarchy — allowing IT managers to leverage existing investments, rather than building a unique set of skill and infrastructure to manage the different types of inactive data.

For more information on managing file-system data and Tivoli Storage Manager for Space Management, see **ibm.com**/storage/software

To learn more about backup copies of data using Tivoli Storage Manager application protection, see **ibm.com**/software/tivoli/solutions/storage/application and **ibm.com**/software/tivoli/solutions/storage/backup

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

To place a given class of data onto the appropriate storage (based on price, performance media type, etc.) — and to facilitate the movement of that data to more- or less-expensive storage media during its life cycle — IT administrators use IBM software to create pools of the different media types in the storage infrastructure. IT administrators can customize these storage pools to represent the varying cost of the media, or to represent a more business-driven grouping, such as a particular project or LOB.

After administrators have defined these logical storage pools, policies direct the allocation and movement of files among them. By separating the logical definition of the storage pool from the physical media used to populate the pool, IBM enables the IT administrator to change physical media without having to rework policy definitions.

In open system environments, the IBM TotalStorage SAN File System can implement storage pools for many different classes of data (see **ibm.com**/storage/software).

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.

Data disposition is driven by the same policy system that directs HSM. The policy definitions vary based on the nature of the data. Some data, such as backup copies, is governed by a defined number for backup versions. Archive data is more readily governed by a fixed life span, expiration date or business event-driven expiration. In regulated environments that implement a full content management system like IBM DB2 Content Manger, business processes integrated with the content management system can control data expiration.

Regardless of the expiration method, the IBM software that manages the centralized hierarchy — Tivoli Storage Manager — then automates the process of deleting the data from the media it was stored on, removing the inventory records and reclaiming the media for future use.

Further to this, policies can be set using IBM Tivoli Storage Manager for Data Retention to store data indefinitely or based on a retention event or predetermined expiration date. In addition, enforcement retention may be applied to data using deletion hold and release which holds data for an indefinite period of time while subject to audits or investigations. The software also verifies data is written correctly and that no modifications or deletions are made after it is stored.

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

When articulating a policy for retaining records existing, four key questions are pertinent:

1. What records or data must be retained?
2. How long should it be retained?
3. Why is it being retained?
4. What happens to the records or data when they are no longer needed?

These key questions are the basis for developing record retention policies. The data must then be classified to decide how each of the different data types must be retained in order to adhere to these policies. Finally, these policies must be entered into IBM DB2 Content Manger so the software can provide automatic classification and indexing of data as described in the record management policy.

Device support. What storage devices does your product support? Are their any proprietary devices (controllers, arrays, host bus adapters [HBAs], switches, SAN topology, virtualization products, and so on) that are required for your solution to work?

The IBM ILM building blocks are designed to operate on the heterogeneous hardware infrastructures that businesses run on.

For a complete list of heterogeneous storage supported see
ibm.com/storage/software/virtualization/sfs/datasheet.html

For inactive data, see
ibm.com/software/sysmgmt/products/support/IBM_TSM_Supported_Devices_for_AIXHPSUNWIN.html

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.).

ILM is a set of building block technologies that are designed to reduce the overall cost of storing and managing vast quantities of data. IBM has a long history of delivering these solutions to its customers and has found that they deliver a positive ROI. The actual physical overhead needed to achieve this ROI varies by implementation. IBM Global Services and IBM Business Partners can provide individual customers with implementation planning services that help them to more specifically understand the complete picture in their environment.

Resource consumption. How much bandwidth and server processor “overhead” is introduced by your solution (for example, to support polling processes, migration processes, agent processing, access frequency counting and so on).

ILM is a set of building block technologies that are designed to reduce the overall cost of storing and managing vast quantities of data. IBM has a long history of delivering these solutions to its customers and has found that they are able to deliver a positive return on investment (ROI). The actual physical overhead needed to achieve this ROI varies by implementation. IBM Global Services and IBM Business Partners can provide individual customers with implementation-planning services that help them to more specifically understand the complete picture in their environment.

Flexibility. Does your solution have multivendor support? Can data, after it has been 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?

IBM DB2 Content Manager runs on a variety of platforms, including Windows, AIX, Sun Solaris, HP-UX, Linux, iSeries and z/OS. It is also supported on IBM DB2 Universal Database and Oracle, with select support for Microsoft SQL Server. (see **ibm.com**/software/data/cm/cmgr/mp/architecture_open.html).

The IBM TotalStorage SAN File System adheres to Portable Operating System Interface (POSIX) standards for file systems. IT managers can choose to move data into the SAN File System and use its policy-based data-pooling features and, just as easily, move data out of the SAN File System back to a traditional file system without ILM features.

The IBM Tivoli Storage Manager interacts with a variety of open interfaces to capture inactive data. These include the Data Management API (DMAPI) — see www.opengroup.org/products/publications/catalog/c429.htm), the XBSA – see www.opengroup.org/publications/catalog/c425.htm and the Network Data Management Protocol (NDMP) — see www.ndmp.org).

Building on this open-standards approach to creating solutions, there is a wide range of existing ISV applications that are interoperable with IBM ILM building blocks (see “Development partners.”)

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?

A key to effective ILM is data classification — understanding which classes of data require what level of quality of service (QOS). A good QOS definition, described in generic terms, allows an IT manager to take advantage of rapidly improving hardware capabilities without having to alter the basic policy structure. For example, a flexible policy for the placement of a database data file might talk about best I/O performance, point-in-time replication capability and continuous replication capability. A policy like this allows the IT manager to transition from the storage hardware device that provided the best I/O performance two years ago to the storage device that provides the best I/O performance today without having to alter the policy system.

Tactically, as IT managers evaluate the hardware and software infrastructure that these policies will act on, they should use industry-standard benchmarks to compare competitive offerings from different vendors. IBM has completed benchmarking its disk systems with the Storage Performance Council, the industry's leading storage performance benchmarking group. The entire line of IBM TotalStorage Enterprise Storage Server models and IBM TotalStorage FAStT storage servers have been publicly recognized and audited through independent and accredited benchmarks.

To get a better view of the whole solution's overall performance, IT managers should try to test all the selected building blocks together. IBM also provides testing capabilities for customers through IBM TotalStorage Solution Centers. These unique centers provide customers with the opportunity to test and evaluate software and hardware solutions in a real-world environment.

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?

In ILM, data is usually moved at life-cycle changes. Depending on the life-cycle stage and the physical location of the data, IBM ILM building blocks can move data over either IP or fibre-channel-based networks. IBM Tivoli NetView® software can help users to manage these networks (see **ibm.com**/software/tivoli/products/netview/).

The management and use of fibre channel networks, although newer, is also quite mature. Users can manage these networks with the IBM Tivoli Storage Area Network Manager (see **ibm.com**/software/sysmgmt/products/support/IBM_TSANM_Device_Compatibility.html).

For enterprises that choose to consolidate the network management of both IP and fibre channel-based networks, the Tivoli Storage Area Network Manager and the Tivoli NetView user interfaces can be integrated.

Protocol support. List the storage networking technologies supported, including Fibre Channel Protocol (FCP), Ethernet, iSCSI, Fibre Channel over IP (FCIP), Internet Fibre Channel Protocol (iFCP), Parallel SCSI, Serial Attached SCSI (SAS) and network file system protocols — Network File System (NFS), Common Internet Filesystem (CIFS), HTTP, Direct Access File System (DAFS).

Data managed by the IBM TotalStorage SAN File System can be accessed through FCP or by file system protocols such as NFS, CIFS, HTTP and FTP.

Users can send data to IBM Tivoli Storage Manager using either FCP or network protocols such as TCP/IP or more infrequently used protocols such as advanced program-to-program communication (APPC), SPX/IPX, LU2 and NetBIOS.

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.?

The data that affects LOBs and that CIOs are responsible for managing is widely varied. As such, an effective ILM solution must have support for a wide variety of environments. As a group, the IBM ILM building blocks support files that reside on IBM OS/390®, IBM z/OS®, IBM AIX®, IBM OS/400®, Microsoft Windows®, HP/UX and Linux operating systems and the Sun Solaris operating environment on Intel®, IBM eServer(logo) pSeries® and IBM eServer(logo) zSeries® platforms. It also supports mailbox and attachment-level data from Microsoft Exchange and IBM Lotus Domino mail systems as well as all varieties of SAP document-level data stored in IBM DB2 Universal Database, IBM Informix®, Microsoft SQL Server or Oracle databases. The IBM ILM group also facilitates the integration of all kinds of structured data with unstructured content. In particular, IBM has looked at a variety of data in delivering its compliance and data retention solutions. In this space, it is not the data type that is important, it is the management of the data. In these kinds of solutions it is data that is kept (retained) and protected for a specific (or unspecified) period of time, usually years. At the end of the retention period, it is important to dispose of the data appropriately.

IBM Content Management runs on a variety of platforms, including Windows, AIX, Sun Solaris, HP-UX, Linux, iSeries and z/OS. It is also supported on IBM DB2 Universal Database and Oracle, with select support for Microsoft SQL Server. IBM Content Management supports storing and managing any file type, and supports viewing of more than 200 file formats. Most common file types include TIF, GIF, HTML, XML, PDF, MPEG, common office file formats such as the Microsoft and Lotus office suites, and a variety of print streams including, but not limited to, AFP, Postscript and Line Data. Any file types not viewable with IBM's native viewing technology can leverage other premier existing viewers of choice.

The IBM Content Management portfolio also provides e-mail archiving capability so that messages and document attachments may be later efficiently organized and retrieved directly from the Microsoft Outlook or IBM Lotus Notes desktops with ease.

To get the complete picture of IBM's ILM software building blocks, see:

- **ibm.com**/storage/disk/dr
- **ibm.com/**storage/software
- **ibm.com**/software/data/commonstore
- **ibm.com**/software/data/cm/cmgr

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

To help IT managers streamline their implementations of ILM, IBM has created a solution that combines many of its ILM building blocks — the IBM TotalStorage Data Retention 450. It combines IBM Tivoli Storage Manager for Data Retention with IBM TotalStorage FAStT Storage Server with S-ATA disks for low-cost disk storage, plus a supporting IBM pSeries POWER™ processor and a lockable cabinet to deliver a complete ILM solution. One feature the Data Retention 450 offers is the ability to use tape and other media as storage tiers. In this manner, IBM provides a unique capability to significantly lower the total cost of ownership of data as it is retained and managed over long periods of time.

List price for an entry-level 3.5TB IBM TotalStorage Data Retention 450 system is US $141,600.

For pricing on individual IBM hardware components and software for managing and storing data, please contact an IBM business partner or an IBM representative.

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

See “Flexibility” above.

Other features. List other pertinent aspects of your solution.

IBM Content Management provides a reliable, scalable and robust infrastructure. It delivers robust, secure and scalable services that provide access to e-business content across the enterprise. This infrastructure supports the convergence of disparate content management technologies, such as email management, document management, digital asset management, and Web content management, with a single repository solution based on a powerful and flexible data model. This data model forms the basis of content sharing across both existing and newly developed applications within any industry or organization. In turn, this content sharing is the basis for major improvements in business process efficiency.

Miscellany: 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.
ILM Services piece with Business Consulting Services Set up services.

IBM offers a consulting organization that combines business insight with technology leadership and world-class delivery under one roof. IBM Business Consulting Services (BCS) is accountable for measurable results at every step of the solution, from strategy to process transformation, application development and delivery. BCS offers deep expertise in 17 industries, including banking and financial services, electronics, automotive, telecommunications, retail, pharmaceuticals and consumer packaged goods (CPG). The group also has a broad set of services capabilities: strategy and change, customer relationship management (CRM), supply chain management (SCM), financial management, human capital management and application innovation.

IBM BCS offers clients the full scope of IBM capabilities, including:
· IBM Global Services — a world leader in IT services, including integrated technology services, application management services, strategic outsourcing, e-business hosting and learning services.
· IBM Research —- IBM invests US$5 billion annually in technology and business research at eight labs with 3,000 employees around the world.
· On Demand Innovation Services — A new consulting practice, created in partnership with IBM Research, brings more than 200 researchers to solve clients’ most complex business problems. IBM will invest US$1 billion in this program over the next three years.
· IBM Global Financing — World's largest IT financing entity, with over $40B in assets.
IBM Software and Systems – Offering a broad portfolio of market-leading software and hardware products.

Usage scenarios
Capacity allocation and efficiency for effective information lifecycle management (ILM)
The ability to classify data — and to measure the amount of data by class in the storage environment — is critical to both capacity management and allocation. This is what enables management of the data based on the value of the information. IBM has a variety of solutions to help accomplish this at several levels such as IBM Tivoli® Storage Resource Manager and IBM DB2 Content Manager.

Capacity utilization efficiency: IBM Tivoli Storage Manager is a key tool that IBM offers for capacity utilization. For more than 10 years, Tivoli Storage Manager has incorporated automated date expiration, using policy-based retention, which automatically deletes data and then reclaims the space afterwards to ensure more effective storage usage. IBM Tivoli Storage Manager’s operational reporting feature also provides daily reporting, relaying

information about space utilization and backup performance and confirming successful backups on a daily basis.

In addition to the Tivoli Storage Manager, IBM Tivoli Storage Resource Manager (SRM) provides historical data for the entire network, detailing how much storage is in use to help determine where users can attain additional capacity. For example, although a customer's storage may grow at 50 percent per year, further analysis using SRM may indicate that 25 percent of the data growth is occurring from duplicate files or storing unnecessary data. SRM can also automatically set policies to delete certain types of files or send automatic notification to systems administrators or users based on the appearance of certain file types.

Capacity allocation efficiency: IBM Tivoli Storage Manager includes archive capabilities that set flexible policies by file type, user, volume and groups. This archival policy-making allows users to automatically allocate the data migration to tape or other background media.

To improve capacity, valid data — the data that is used to generate revenue — can be pooled and shared. Customers struggling with the constant fluctuation of storage hardware — and the application outages that result — can improve flexibility in the storage infrastructure through virtualization.

Additionally, the IBM TotalStorage SAN Volume Controller and SAN File System (based on IBM Storage Tank™ technology) solutions provide a virtualized storage environment that allows users to share capacity and reduce the cost of maintaining a physical storage environment. SAN File System provides innovative technology such as policy-based file provisioning, distributed file locking and a file-based IBM FlashCopy® function. IBM designed these features to help increase performance when sharing information and to improve productivity by automating routine data-management tasks. For example, the policy-based file provisioning function can automatically allocate space for files into storage pools, using customer-defined rules. This allows the customer to better target and address specific performance or availability requirements.

Disaster recovery and business continuity

The rationale for ILM is that information has variable business value, tending to change over time and through business events, and that to control costs, CIOs need a variable-cost hierarchy of data storage. This cost variability can include the cost of delivering highly available data to support a critical business application. The most-demanding availability requirements require seamless integration among server and storage resources, coupled with automation. The IBM Geographically Dispersed Parallel Sysplex® (IBM GDPS™) is a highly effective disaster recovery and continuous availability solution for multisite enterprises. This solution automatically mirrors critical data and efficiently balances workload between the sites. GDPS also uses automation and IBM Parallel Sysplex® technology to help manage multisite databases, processors, network resources and storage subsystem mirroring. (See **ibm.com**/servers/eserver/zseries/announce/april2002/gdps.html).

However, not all businesses, and not all classes of business data, require GDPS. CIOs and IT managers should evaluate the business value of the information, the recovery-time objective (for example, how quickly the data must be available following a disaster) and the recovery-point objective — that is, as soon as the data is available again, how far back from the failure point can the data be recovered? After the CIO or IT manager has relayed this information, IBM can use ILM building blocks to effectively classify the data (see **Data naming scheme** later in the text). IBM can also use design policies to pool the active data on media with the appropriate point-in-time or continuous replication services (see **Storage platform characterization** later in the text). Other design policies can help users store the backup copies on the appropriate media pool to support rapid recovery or off-site vaulting.

ILM allows users to move data to lower-cost (and frequently offline) media as it goes through life-cycle changes, leaving more active, tactically relevant data on the highest-cost storage. This remaining active data frequently requires the quickest recovery time and the most recent recovery point. Using ILM, CIOs can focus their investments on providing a highly available infrastructure for this critical data without having to buy extra high-cost capacity for inactive data.

In addition to the basic ILM capabilities of data classification, active and inactive data pooling, and hierarchical storage management **(**HSM), IBM offers a complete range of building blocks to help CIOs meet a wide range of recovery-time and recovery-point objectives, including:

- Capturing reliable recovery points for file systems, files, databases, application servers, mail systems and enterprise resource planning (ERP) systems with IBM Tivoli Storage Manager application recovery solutions (see **ibm.com**/software/tivoli/solutions/storage/application and **ibm.com**/software/tivoli/solutions/storage/backup)
- Providing point-in-time and continuous replication services for files and disk volumes across heterogeneous storage infrastructures with IBM TotalStorage SAN File System, IBM TotalStorage SAN Volume Controller and IBM TotalStorage SAN Volume Controller for Cisco MDS 9000 (see **ibm.com**/storage/software
- **ibm.com**/storage/software/virtualization/svc/index.html
- **ibm.com**/storage/software/virtualization/svc_cisco/index.html)
- Providing point-in-time and continuous replication services for IBM disk volumes, using the IBM TotalStorage Enterprise Storage Server® (ESS) and the IBM TotalStorage FAStT Storage Server (see **ibm.com**/storage/disk)
- Automated disaster recovery plan generation, maintenance and system-level restoration (see **ibm.com**/software/tivoli/solutions/storage/disaster)

IBM Business Continuity and Recovery Services — with a presence worldwide in 78 countries across four continents — providing consulting services and recovery support for large systems, midrange and distributed environment customers (see **ibm.com**/services/continuity/recover1.nsf/worldwide/countries)