

ADVERTISEMENT

MANAGING INFORMATION

in the Life Sciences

“For life sciences organizations and bio-IT professionals, data is necessary, but not sufficient, for organizational success.”



There are approximately 30,000 human genes that create more than one million proteins in the human body. Only one in thousands of these New Chemical Entities (NCEs) is transformed through the drug discovery process into a commercial pharmaceutical. In fact, since 1985 the number of NCEs commercialized per year has dropped from more than 60 per year to below five, while R&D expenditures have increased from \$35 billion to \$45 billion per year.

The economic and market pressure from rising R&D costs and longer product development cycles means that the biotechnology and pharmaceutical industries are working harder and longer to discover new and effective pharmaceutical agents; creating more business partners and more complex data during product development; and uncovering many more molecules that must be evaluated and processed through more comprehensive clinical trials before finally being approved by the U.S. FDA.

The combined effects of discovery, approval and marketing pressure on bio-IT and the professionals who manage bio-IT organizations are profound. New technologies, including genomics, proteomics, pharmacogenomics, and chemical informatics are developing rapidly. Data storage and information management requirements due to new drug discovery, research and development, manufacturing, and marketing are doubling in bio-IT every year.

For life sciences organizations and bio-IT professionals, data is necessary, but not sufficient, for organizational success. They must generate *information*—meaningful, actionable, organized, and reusable data. Data must be stored, protected, secured, organized, distributed, and audited—all without interruption. Moreover, it must be expanded efficiently and continuously from the NCE to a New Drug Application (NDA) and into the marketplace.

The Life Sciences IT Infrastructure

Over the decades, information delivery architectures in the pharmaceutical industry have evolved from batch computing and timesharing to distributed computing, and now to networked workstations and servers, including clusters for cost-effective computationally-intensive tasks. The infrastructure in bio-IT has also evolved from mainframes, mini-computers, and microcomputers to Internet Protocol (IP)-based computer networks where storage management systems may be globally distributed.

The operating model for the life sciences enterprise has evolved from an application-focused one to a distributed transaction-processing environment that provides a wide range of services to researchers, internal employees (using the internal intranet), partners (using the extended extranet), and customers and regulators (using the public Internet).

These services demand diverse communications methods, including Web protocols, voice, video, and document and communications sharing protocols, and they can generate and require increasingly complex data and data structures.



HP Superdome supercomputer with HP Surestore xp512 “Data center in a box” with no single point of failure for life sciences organizations

Rapid changes in both technology and bioinformatics have had effects on bio-IT that can be characterized by a number of trends and needs:

- Further integration of the Internet to support biotech eR&D initiatives and online partnering initiatives; for example, Internet-based clinical trials developed with multiple academic medical centers;
- Industry and regulatory demands for an extremely reliable networked infrastructure with no potential for data or document loss;
- Convergence of voice, data, and video across a common global network infrastructure;
- Global reach across geographic regions to reach global partners and regulators;
- A gaping shortage of IT skills, particularly in systems architecture, planning, operations, and Internet-related applications;
- A growing need to outsource the management of IT resources, allowing the life sciences enterprise to refocus resources on their core businesses—discovering and commercializing NDEs.

Creating Information Within the Life Sciences

For decades, Hewlett Packard has been providing IT technology and services to healthcare organizations, pharmaceutical companies, and life sciences organizations. HP has developed a storage systems strategy, Federated Storage Area Management (FSAM), that meets the information management requirements for bio-IT.

FSAM merges the concept of Federated (many-as-one) computing and Storage Area Management (SAM) to address two key storage challenges: infrastructure optimization and effective management of heterogeneous storage assets. SAM is the centralized management of resources and data across storage domains, providing shared storage services that can scale as needed. This is done in the FSAM architecture by adding new processors, storage devices, and technologies to current multi-vendor storage to achieve a greater total return on investment from the entire storage environment.

FSAM is an umbrella strategy that guides HP storage product development to meet key needs: standards-based, open components that are easy to manage, scale, and secure. FSAM encompasses three key focus areas as well as the architecture necessary for optimizing and managing storage: (1) Business Continuity, (2) Storage Management, and (3) Storage Consolidation and Virtualization:

- **Business Continuity**—Solutions that provide maximum performance and ensure that data is secure, protected, and available. These allow multiple copies of files and databases to be processed simultaneously, ensuring that applications required to be disaster-tolerant meet that requirement.

- **Storage Management**—Central management of distributed storage resources ensures that growth of data can be efficiently and effectively managed with the same resources in place today. This management architecture is essential for optimizing storage asset utilization.

- **Consolidation & Virtualization**—Multi-

ple vendor storage systems, server attached storage, and storage area networks are merged into a virtual enterprise storage solution that consolidates distributed storage into an easily managed common pool of storage.

- **Architecture**—An architecture for the life sciences that ensures that life sciences organizations can successfully continue their exponential growth of data and information, leading to more rapid time-to-market of NCEs and pharmaceutical products.

Implementing state-of-the-art storage hardware and software from HP enables these capabilities.

Business Continuity

With each NCE costing more than \$800 million to reach the marketplace, every day counts. HP provides business continuity products and services that are tuned to the needs of the life sciences enterprise. HP is the only vendor that provides complete high availability backup with no backup window, no recovery window, and no application disruption. HP products include:

- **HP Surestore Continuous Access XP and XP Extension** are high-availability data and disaster recovery solutions that deliver host- and operating system-independent real-time remote data mirroring between HP XP disk arrays. HP has unique asynchronous mirroring advantages that do not require the expense of extra arrays and data centers traditionally required in multi-hop solutions.

- **HP Surestore Business Copy** creates multiple, real-time online copies for zero-downtime split-mirror backup, disaster recovery mirroring, decision support, data warehousing, and application development and testing.

- **Oracle operating on HP servers.** HP and Oracle have developed a unique partnership to ensure that life science databases are more easily designed, created, and available. HP's

Precise/Crosspoint, for example, manages the Oracle infrastructure, and HP Oracle9i Real Application Clusters ensure business continuity. HP supports world-class database performance supported by comprehensive database administration tools for DBMS backup/recovery and failover using inexpensive server clusters. Oracle 9i RAC was developed in partnership with HP over a seven-year period and the 9i RAC was stress-tested on the HP Surestore xp512.

Storage Management

Genetic, proteomic and combinatorial chemistry applications, as well as clinical trials, can generate hundreds of gigabytes of data per week. Managing this constant storage growth efficiently without increasing your storage personnel at the same rate will become critical for the rapidly growing life sciences entity. HP provides integrated software and services for infrastructure management, such as:

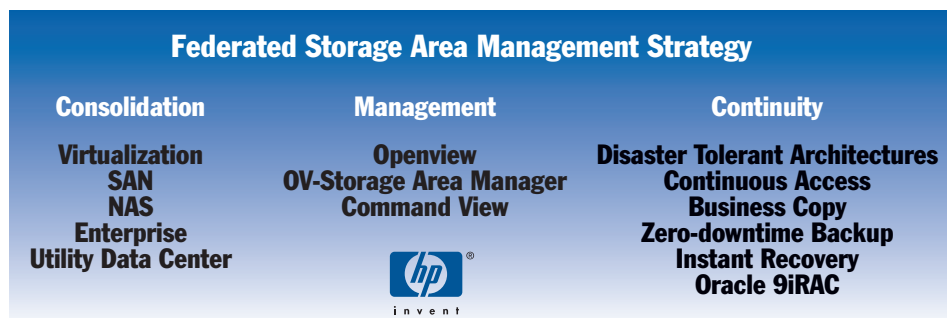
- **HP OpenView Services and Software.** HP OpenView provides comprehensive, multi-vendor services and software to manage the complex computing infrastructure that will be created within a life sciences organization. HP OpenView can manage your end-to-end IT environment—storage networks, servers, mainframes, and supercomputers.

- **HP OpenView Storage Area Manager.** HP OpenView Storage Area Manager software enables you to centrally manage multi-vendor storage devices across distributed, networked storage environments, spanning direct attached storage, network attached storage (NAS), and storage area networks (SANs). Thus, it provides a uniform management console and a common reporting structure for highly efficient and cost-effective storage area management. OpenView Storage Area Manager can discover, manage, allocate, optimize, and account for storage systems across multiple platforms and processing nodes.

Consolidation & Virtualization

Computational power and information management are growing at startling rates in the life sciences. From Apple® computers to the largest massively parallel processors, to multi-platform computational and database management servers, the life sciences organization needs integrated platform consolidation and management tools. HP provides integrated hardware, software and services to consolidate and manage the IT infrastructure, particularly storage systems which house your most important asset—your data—and enable efficient information flow. Solutions include:

- **Storage Virtualization.** Storage Virtualization creates and manages storage devices and resources as one heterogeneous storage pool, shared by all servers, enabling them to do more with less.



Multi-vendor storage systems can be pooled into a continuous storage "fabric" for optimization, tuning, and management. Virtualization enables customers to increase management efficiency, and increase the utilization of their storage resources, including hardware and administrators. This allows bio-IT professionals to manage more than ten times more storage than they do today, because virtualization minimizes the complexity of disk management and capacity expansion.

• **Storage Area Networks and Network Attached Storage Integration.** Consolidation implies building an enterprise storage system based on a SAN architecture and growing it to facilitate the sharing of SAN and NAS storage in a network-centric storage model. Through the capabilities of Fibre Channel, NAS heads, and storage security software, a single architecture can dynamically scale and support virtually unlimited NAS and SAN growth.

Architecture

With more than 60 years experience in computing and information management, HP is a global leader in the development of innovative systems architectures. HP is developing IT "blueprints" for life sciences bio-IT professionals. HP architectural solutions will support discovery, document management, eR&D, NDA automation, and post-market support for newly-introduced pharmaceuticals. HP understands that for the life sciences IT systems, services, and solutions must focus on speed-to-market and bring world-class reliability, security, and business continuity solutions to the life sciences organization. HP architectural components for the life sciences include:

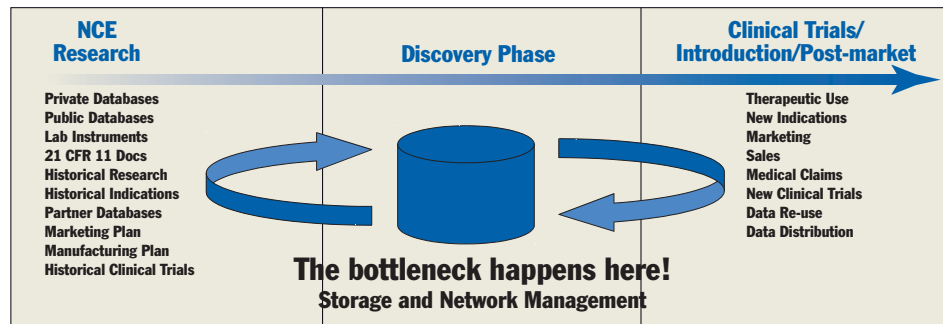
• **A Computational and Storage Architecture.** Using NAS and particularly SANs, a foundation storage architecture can be developed using HP software and services. This architecture can evolve to create a life sciences computational utility, providing computational power and storage on demand using a utility pricing model.

• **The HP Superdome.** The HP Superdome is a 16, 32, or 64 processor supercomputer using hardware and software partitioning to create and manage "virtual" computers and processing spaces. Coupled with a utility pricing model and high performance virtualized XP storage systems, the Superdome becomes a "data center in a box." For the life sciences entity, the Superdome allows consolidation of applications, from eR&D and NDE discovery to document management.

• **The HP Utility Data Center (UDC) architecture.** An innovative data center and information "utility" IT strategy and architecture designed to ensure high performance and high availability—all modeled on the "pay as you go" pricing model of modern utilities. The UDC is the future of Internet-based life sciences applications.

Breaking Through the Life Sciences Storage Bottleneck

The Life Sciences Challenge: "Continuous Data Re-use"



The CIO of the life sciences organization is charged with managing complexity and growth within bio-IT. This means focusing on architecture, planning, and cost control. Rapid expansion of platforms, applications, and storage subsystems can lead to a proliferation of "islands" of underutilized storage and rapidly escalating management and data protection costs. Managing for success in the life sciences in IT means focusing on consolidation, control and visibility, and continuous availability.

The life sciences have a tradition of scientific and technical planning unlike most other industries, for it is rare in any industry for product development to require 10 years and nearly a billion dollars of investment. Breaking through the information storage management problems in the life sciences requires planning, architecture, and world-class tools and techniques. HP provides integrated services, computing and information systems hardware and software, coupled with more than 60 years of engineering excellence. These solutions are available to life sciences organizations as required. To arrive at a successful strategy and architecture for enterprise storage management, HP recommends using a federated storage area management strategy and architecture.

Key Attributes of the HP Federated Storage Area Management Strategy

FSAM combines physically separate storage domains (subsystems), types of storage media, and various vendor storage subsystems operating as networked storage devices into a logical storage resource. Federated storage breaks down "islands" of data into a uniform enterprise storage network that frees underutilized storage devices for use by the entire enterprise.

HP has identified the following attributes that are characteristic of a successful FSAM strategy:

• **Scalability.** Storage resources expand and contract on demand or by policy without application or database interruption.

• **Performance.** Multi-vendor storage systems are optimized for access, maximizing the performance of the storage devices.

• **Cost Effectiveness.** The total cost of ownership of physical storage decreases as the scale increases by utilizing fewer personnel.

• **Optimization.** Data is easy to move and reorganize, and can be copied on demand or by policy among multiple logical storage subsystems and physical locations.

• **Virtualization.** Storage devices from multiple vendors are integrated into an enterprise pool of storage resources that can be allocated on demand.

• **Openness.** The expansion and maintenance of multi-vendor storage is non-disruptive to the computing environment and easily accomplished both physically and logically.

• **Self-Management.** Storage systems are self-managing, self-tuning, and self-healing resulting in the highest levels of data access.

• **Secure.** Data is protected from access by an unauthorized application or user across a distributed network of storage with the ability to change access rules on demand.

• **Easy to Maintain.** Storage appears as a uniform network infrastructure, a resource that is configured once and maintained centrally.

These attributes describe a high performance storage architecture and environment—an environment that is consistent with HP's goal to develop and deliver a network storage infrastructure that meets the demands and challenges expected within the life sciences.

A Recipe for Success

HP's experience with mission-critical applications in multiple industries indicates that a number of features are critical for the success of any complex computing infrastructure built on top of a storage management architecture. For the life sciences organization, this means increasing the speed-to-market of NCEs and rapidly getting through the clinical trials and NDA approval processes. For the bio-IT executive, this means focusing on:

- **Vision & Strategy.** The overall bio-IT architecture must be conceived as a whole, executed through a comprehensive plan and not built in a piecemeal manner.
- **Security.** The architecture must emphasize security yet be efficiently operated.
- **Auditability.** 21 CFR 11 requires that all documents, for example, be stored, versioned as modified, and continuously audited throughout their life cycle.
- **Flexibility.** The operational architecture must be flexible in terms of both the delivery of application solutions and in managing data content. Both partners and employees will inevitably want to personalize it. The form and content of data will also change over time.
- **Complexity.** A storage architecture must be introduced through a manageable plan. The growth of data will create complex demands on the infrastructure and the operational systems.
- **Availability.** The system must operate on a 24X7 basis, without interruption.
- **Control.** Storage systems must allow for ease of maintenance and uninterrupted local and remote control. They should allow easy expansion and reallocation.
- **Re-use.** Data must be easily re-used and reformatted to support discovery, clinical trials, marketing, and post-market analyses.
- **Ease of use.** Storage systems and applications must be user-friendly. Research and development partners may be spread throughout the world and use multiple languages and possess local data management conventions.
- **Manageability.** Storage and application systems must be centrally managed to ensure optimal efficiency and resource utilization.

Tips to Improve Traditional Processes

- Have an architecture.
- Involve partners.
- Integrate all discovery data in databases.
- Merge clinical, marketing, and post-market data.
- Submit NDAs electronically.

Conclusion

Storage systems are the foundation of integrated life sciences data management and the NCE discovery process. Storage systems exist to maximize the operational flow of information and facilitate data sharing, re-use, and analysis.

The use of a federated storage management strategy within a virtual storage architecture will encourage the life sciences organization to optimize and manage its information quickly and effectively. Given that the fully amortized costs of research and development for an NCE can average \$6 million

per month, effective management of data and information within the life sciences organization is not optional—it's imperative.

HP stands ready to provide the life sciences industry with fully integrated storage management, computational and information management strategies, services, and products to ensure that its life sciences customers maximize their technical excellence, while simultaneously improving their speed-to-market of NCEs and new pharmaceutical products.

**To help you add more efficiency and manageability to your storage environment, visit HP online to receive your free design guides for SAN and Disaster Recovery:
www.hp.com/go/storage_blueprint**



Hewlett-Packard
3000 Hanover Street
Palo Alto, CA 94304-1184
www.hp.com