THE DAFS BUSINESS CASE

Jon William Toigo, Independent Consultant and Author of “The Holy Grail of Data Storage Management”

Deconstruction: The 21st Century Business Model

If there is one recurrent theme as business enters the 21st Century, it is deconstruction. At one level, deconstruction refers to the initiatives increasingly seen within companies to change their business processes to take advantage of new opportunities enabled by web technology and the Internet. Following the path charted by Cisco Systems, Dell Computer Corporation, Ford Motor Company and others, more and more companies are deconstructing key internal processes - logistics, product design, product assembly, and so forth - and shifting responsibilities (and costs) for component tasks to their supply chain partners. The motivation driving this business process deconstruction is to create improved processes that will enable the company to operate more efficiently and profitably and to respond with greater flexibility and speed to rapidly changing market requirements.

Looking deeper at business process deconstruction, one can see its technology enablers. The 21st Century company is capitalizing on the improving technologies for Business-to-Consumer (B2C) interaction to facilitate customer order taking via wired and wireless network connections to the Internet. Similarly, these companies are depending upon improving Business-to-Business (B2B) technologies to integrate their supply chain processes and to fulfill customer orders efficiently and profitably. In essence, they are using technology to create "virtual companies" comprised of numerous functional building blocks, including service providers and supply chain partners. The modern business enterprise is no longer defined by the physical boundaries of a corporate premise, but by a virtual extra-enterprise network of seamlessly connected partners, suppliers and customers.

Supporting this virtual company is an integrated technology platform that is itself a virtual, deconstructed, system comprised of functional building blocks, such as the Internet, the universal web browser client, web technology-enabled software, next generation servers, and new enabling protocols.

In the realm of application software, monolithic data processing application architectures have given way to modularized application architectures that scale more efficiently and without performance degradation. The deconstructed application accomplishes what traditional client-server technology could not: using ubiquitous web technology and modular design, these applications are easier to grow and change in response to changing business requirements.

Deconstruction is also the order of the day in server host architectures. To improve efficiency, scalability and performance, vendors are providing specialized appliance engines that can be joined together in unique

combinations, much like technological LEGO* blocks, to support the innovative strategies of business planners. More and more the traditional data center, with its Big Iron systems and general purpose servers, is giving way to a model frequently seen in the web hosting environment: specialized appliances dedicated to specific discrete tasks and mounted in 19 inch racks. The latter provides an integrated platform that is at the same time more powerful and less costly to deploy and maintain than its predecessor. Isolating functions via a deconstructed server architecture also limits downtime, familiar in general purpose servers, but completely unacceptable in the 24x7x365 world of the Internet.

Looking more closely at next generation appliance servers, one can again see the impact of deconstruction. Intel Corporation's Virtual Interface (VI) architecture is a significant application of the concept. VI, the next generation architecture for all Intel servers, is fundamentally a distributed, networked (clustered) computing architecture implemented "inside" the computer itself. From 50,000 feet, the most noticeable attribute of a VI system is the proliferation of CPU chips. Instead of older system architectures, in which a single central processor polls numerous devices on the motherboard (or on daughter boards installed in bus slots), VI systems are comprised of multiple channels, each equipped with its own CPU. The main CPU acts like a cluster controller. Dedicated channel processors handle specific functions -- such as storage I/O, network stack processing, etc. -- locally and without requiring the services or cycles of the central processor on the motherboard. This enables servers utilizing VI architecture to be custom configured to meet very specific needs.

VI also provides the means to cluster numerous servers together in fault-tolerant configurations. So, the individual, specialized server appliances can work and act together as one virtual system.

Thus, from the business process to the technology infrastructure to the application and, finally, to the server, deconstruction is the philosophy that is guiding sea change in the way that we do business. Successful deconstruction of processes and their enabling technologies will determine the profitability - and ultimately the survival - of companies in the 21st Century. It is against this backdrop that the business value of the Direct Access File System (DAFS) can be most clearly perceived.

Introducing DAFS

The Direct Access File System protocol, under development by a consortium of more than 150 companies, is the first of what will likely become a generation of protocols designed to capitalize on the VI computing architecture. It answers the needs of business by providing a means for applications and end users to store and retrieve data efficiently within a deconstructed computing environment.

To provide a context for understanding DAFS and its importance, it may be useful to look at present day data storage trends and directions. The first trend, well understood by most businesses, is the explosion of data growth.

The exponential growth of data is now an established fact. Depending on the analyst one reads or the business experience one consults, the volume of data that is being generated within corporations is growing at a rate of between 80 to 100 percent annually.

Opinions vary about the source of data growth. Some analysts attribute the trend to the Internet, others to the proliferation of new, resource-hungry applications. Still others cite as a key growth factor the increasing tendency of organizations to store all data on-line, rather than migrating older or less frequently accessed data to secondary storage media such as optical and tape.

Not to be ignored is the amount of data growth that can be attributed to data replication. Data replication can account for up to 50 percent of increased capacity demand. Replication may be necessitated by applications themselves, as in the case of data warehousing, but it is more often the result of inefficient data sharing methods.

Whatever the key drivers of data growth, one fact is certain: the data that is being generated by organizations needs to be stored in a manner that makes it readily accessible to users and applications authorized to use it. Put simply, data must be stored in an architecture that is cost-effective, scalable, manageable, and secure. And which is consistent with the distributed character of the great majority of the ultimate users of data. Hence, it is key that organizations deploy architectures that support distributed data access while centralizing data management. In addition, the architecture needs to support distributed central pools of shared data since quality of service normally falls below acceptable levels if all data is held exclusively in physically central repositories.

Most analysts agree that meeting mounting storage requirements will necessitate a shift away from internal, or "server-captive," storage approaches and toward external, or "network-based," storage architectures. To meet the capacity, scalability and manageability requirements of IT organizations, in which terabyte-sized storage platforms have become the rule rather than the exception, server-captive storage has become too costly, prone to downtime and unmanageable. The market has awakened to the need for external storage architectures that scale beyond the server or array box.

The onset of an architectural shift within data storage is confirmed by several industry studies. A recent one, prepared by International Data Corporation, places the compound annual growth rate (CAGR) of revenues realized by vendors of server-captive storage technologies -- including server cabinet-mounted disk drives and server-attached storage arrays -- at negative 3 percent for the period under study (from 1999 through 2003). By contrast, IDC anticipates a positive 67 percent CAGR in revenues resulting from the sale of external "networked" storage in the same period.

The industry has responded to the need for networked data storage with two architectures: network-attached storage (NAS) and storage area networks (SANs).

NAS embodies an appliance-based concept for storage networking: with NAS, a dedicated, storage optimized operating system kernel controls a storage array that connects directly to the corporate LAN. NAS appliances capitalize on TCP/IP protocols such as Dynamic Host Configuration Protocol to install in a plug-and-play fashion onto the company LAN and advertise their disk volumes for use via network file system protocols, such as the Network File System (NFS) and Common Internet File System (CIFS).

SANs, by contrast, provide unlimited scalability by delivering a mechanism for interconnecting servers and storage devices in a back-end network, rather than to the LAN directly. To grow SAN storage, in theory, one needs only to add more storage devices to the SAN. Challenges confronting the SAN, however, are numerous and current-generation products are difficult and costly to implement and scale.

DAFS provides a deconstructionist approach for scaleable, highly accessible, high-performance storage that captures the best of both NAS and SAN. As a network file system protocol, DAFS enables the creation of storage appliances along much the same lines as current generation NAS appliances. However, a DAFS-compliant NAS appliance can be used to provide efficient shared access with high performance to any back-end storage platform - whether an integral appliance array or a storage area network. Thus, DAFS can be viewed as an enabler of SAN-NAS hybrid solutions aimed at delivering the plug-and-play simplicity of NAS installation together with the scalability and high speed access of SAN or direct-attached arrays.

Benefits of DAFS-enabled Storage

The key feature to understand about DAFS, and what sets it apart from present day network file system or SAN protocols, is that it takes advantage of VI architecture directly. The protocol leverages VI's direct memory mapping to support high-speed storage access and retrieval. The result is that DAFS-enabled storage is as responsive as dedicated block-level storage arrays in high performance, transaction-processing applications.

Other extremely important features of DAFS-enabled storage include the following:

- DAFS is designed to support VI's distributed bus architecture. DAFS enabled storage appliances communicate directly with the dedicated I/O processors in VI application servers, without requiring additional central processor unit cycles. Thus, to the application server, the presence of DAFS-enabled storage on the network is handled in much the same way as a locally attached disk drive.

- The DAFS protocol will utilize the same, well-understood and tenured IP networking technology as other network file system protocols. There is no new technology service to learn or provision at the network or system level in order to make DAFS storage work.

- Like all other server appliances, DAFS appliances can be configured for redundancy and fail-over in order to achieve fault tolerance and zero-downtime objectives.

- DAFS, as an IP-based, network file system protocol, is compatible with a broad range of physical and network transports. Whether storage accessibility is required within a corporate Ethernet or across a metropolitan or wide area network, or over a wireless or satellite-based net, if the network supports IP, it also supports DAFS.

- DAFS is “wire” agnostic. That means that storage devices in a DAFS-enabled storage appliance can be connected via SCSI, Fibre Channel, Infiniband or virtually any other interconnect selected by the product developer. This fact prevents the business from becoming limited to a single storage product provider and enables decision makers to select best of breed products for inclusion in their scalable storage infrastructure.

- DAFS enables a flexible storage infrastructure strategy. DAFS enabled storage appliances can be deployed to serve the local file sharing requirements of a workgroup, or they can be clustered together to form large storage pools for use in server farms or server clustering environments. With the DAFS protocol, there are no "isolated islands of storage" - files can be shared immediately, based on requirements and policy, across campus, across town, across the country or around the world.

Making the Business case

These benefits provide a compelling technology business value proposition for the Direct Access File System within the context of the 21st Century company. Specifically:

- Cost Savings: DAFS-enabled storage appliances eliminate the cost to acquire and maintain expensive, dedicated, server-attached storage devices. DAFS storage can be deployed more cost effectively than current-generation SANs, and may include a SAN "back-end" that is managed by the DAFS appliance as a packaged solution. DAFS storage requires no staff retraining, network upgrades or other enabling infrastructure changes. Most importantly, DAFS-enabled storage participates in the cost-reducing, performance-enhancing infrastructure that is produced by VI appliance servers and clusters, deconstructed application architectures, and next-generation B2B and B2C platforms.

- Risk Reduction: DAFS-enabled storage appliances provide support for redundancy and fail-over configurations that can greatly reduce the likelihood of interruptions of access to data. The DAFS protocol leverages the established standard protocol suite for networking, TCP/IP, thereby ensuring its compatibility with existing and future networks and transports that support IP. DAFS itself is being developed as an open standard, enabling the development of a variety of storage products by a variety of vendors that can meet company needs. As a consequence,

the risk of becoming locked in to a particular vendor's product is eliminated.

- Business Enablement: DAFS-enabled storage is part of an evolving technology infrastructure designed to support business process deconstruction and next-generation extra-enterprise business processes. By delivering a storage infrastructure characterized by flexibility, high performance, accessibility, resilience and scalability, DAFS is a key component of a technology infrastructure that will position a company to achieve its goals.

Next Steps

The description and business case articulated above are visionary. The Direct Access File System is an evolving technology spearheaded by the collaboration of more than 150 members of the DAFS Collaborative. The Collaborative is not merely seeking to develop a new protocol for file system networking, but an integral part of a larger solution aimed at enabling the evolution of business itself.

The DAFS Collaborative needs you.

If you are an end user of technology, the DAFS Collaborative seeks your input and insights. Business requirements are driving the deconstruction process and dictating the role that technology will be asked to play. As a business decision maker, you are "in the trenches" and know first hand what problems need to be addressed. The DAFS Collaborative challenges you to become an active participant of the DAFS development effort by:

- Learning more about current development directions and objectives. Visit the DAFS Collaborative website at www.dafscollaborative.org.

- E-mailing the DAFS Collaborative to let us know what is missing with current data storage and access methods used by your company and what you are seeking from a next-generation solution. The email address is listed below.

- Supporting the DAFS Collaborative initiative by volunteering to serve as a beta test site for protocol-based products.

- Spreading the word about DAFS to your peers and to your vendors, encouraging the former to become knowledgeable about this important technology, and the latter to ensure that the protocol is understood and embraced within their technology solutions.

- If you are a technology vendor, the DAFS Collaborative invites you to join this important initiative. DAFS will be submitted for a formal standards review once it is finalized within the Collaborative. It is hoped that the agreement among the more than 150 leading companies who are already contributing their energies to protocol development (and, in some cases, implementing it in new products) will enable quick

ratification of what will almost certainly be a de facto industry standard. You can become a full participant in the DAFS Collaborative by:

- Visiting the DAFS Collaborative web site for information about the current protocol definition and the terms of collaboration by participating developers. Go to www.dafscollaborative.com.
- Determining the level of participation you wish to provide and signing up using on-line registration materials.
- Participating in DAFS Collaborative Working Groups and Developers Conferences.
- Participating in DAFS Collaborative promotional efforts.

The DAFS Collaborative is a vendor-neutral organization. The Collaborative does not endorse the products of specific vendors, but plays a key role in articulating the value of DAFS-enabled storage.

Thank you for your interest in DAFS.

For additional information, contact

Werner Glinka
Executive Director
DAFS Collaborative
650 851 5909
fax 851 5987
werner.glinka@dafscollaborative.org

or visit **www.dafscollaborative.org**