



actionable intelligence
www.ideasinternational.com

Where is open source storage today? Can commercial vendors overcome stigma and create a profitable business model?
See page 3

TECH TRENDS MONTHLY

June 2010 Newsletter

REGULARS

Ask IDEAS

What Is the Best Way to Prepare for Cloud Computing in the Future? 2

IDEAS Insights

Blog Bites 5

FEATURE

The State of Open Source in Storage

The real challenge in the commercial storage world is finding a means to build-out a profitable business model that involves strong community-based development, while simultaneously overcoming the existing stigma of open source technology.

Tech Trends Monthly is free!

To subscribe click below link or visit:
www.ideasinternational.com/misc/msc060.aspx

[subscribe](#)

Ask IDEAS

What is the best way to prepare for cloud computing in the future?

Is it necessary to deploy some form of a private cloud before public clouds can be integrated, or can users start drawing on external cloud computing services without major changes to existing IT infrastructure?

Answer: The best approach to cloud computing will depend on which benefits an organization primarily hopes to achieve from the cloud. Many organizations are interested in cloud computing mainly because they want to optimize their IT investments so that they spend more on computing activities with the most strategic value, and less on maintenance and operations. A number of organizations, however, are less concerned with managing costs and more interested in other benefits of cloud computing. For example, some are interested in the potential for cloud computing to increase productivity through collaboration. Others see cloud computing as a way to gain processing capacity that they would not otherwise be able to deploy on premises. Such additional processing capacity would enable these organizations to develop analytical capabilities that may produce meaningful business differentiation.

Which of these benefits are most important to an organization will determine what the best approach is for that organization to take in pursuing cloud computing. If the primary goal is to lower costs, there are several paths to take. The simplest is perhaps using Software-as-a-Service (SaaS) for selected computing functions that are not considered strategic differentiations, such as e-mail and collaboration. SaaS can be introduced without significant changes to on-premise IT infrastructures. Organizations simply make a business decision to stop hosting the function internally, and instead draw on the functions from a third-party service provider.

If an organization wants to pursue costs savings more generally across the IT environment, a variety of optimizations can be applied to existing IT infrastructures through cloud-like approaches that will help to improve operational efficiencies. These optimizations may be based on enhanced management tools and practices, as well as emerging technologies. For example, power and cooling costs can be lowered by introducing newer and more efficient server platforms; improved management tools can increase the ratio of servers to administrators; and techniques such as virtualization can lower the costs of maintaining uptime and provisioning new services. When sufficiently coordinated and refined, these optimizations can produce an infrastructure that yields many of the benefits of cloud computing, but without any use of public cloud services. These so-called private clouds may eventually be supplemented with public cloud computing resources, becoming "hybrid" clouds if the public resources can somehow be acquired more cost-effectively than on-premise resources.

Organizations may be able to improve collaboration by refining the way that users access applications and data in existing computing infrastructures. For example, Virtual Desktop Infrastructure (VDI), in which virtualization is used to host desktop applications on centralized servers that are accessed remotely, may make it easier for users to share computing sessions than traditional client-server infrastructures, in which each user has a private computing system that has to be linked with others. Again, this benefit can be achieved in private clouds, without the use of public cloud services. However, the use of third-party services may facilitate more global collaboration (i.e., crossing organization boundaries).

A sufficiently parallelized application can be deployed easily and quickly on a public cloud, and does not necessarily require deployment of a private cloud . . .

The ability for cloud computing to increase processing capacity will depend on the workload being considered, and how well the applications required for that workload have been designed to scale on flexible infrastructure. As discussed in a previous column, applications that "scale out" are generally the most suitable for deployment on cloud infrastructures, since it is easier to assign computing resources to these applications on an incremental basis. A sufficiently parallelized application can be deployed easily and quickly on a public cloud, and does not necessarily require deployment of a private cloud (except for testing and development, or if there are concerns about data latency and/or security). However, the deployment of a private cloud for hosting these classes of applications on premises will give users the greatest flexibility for opportunistically switching between private and public infrastructure depending on cost, service level, or security requirements. ■

**Do you have a question for IDEAS?
AskIDEAS@ideasinternational.com**

The State of Open Source in Storage

Scott Kellogg, Research Analyst

An open source strategy does not often present a clear path to profitability, even when backed by commercial vendors. After all, open source is notable for releasing structured code that is freely available for public utilization. By its very definition, open source – and the community-based development it encourages – is not designed to facilitate a stout profit structure. However, many users are attracted to the benefits of open source, which often necessitates some type of open source strategy from commercial vendors. At the same time, other users are still questioning the robustness of open source solutions versus proprietary ones. Hence, the real challenge in the commercial storage world is finding a means to build-out a profitable business model that involves strong community-based development, while simultaneously overcoming the existing stigma of open source technology.

Where Is Open Source Storage Today?

Storage vendors today are not strangers to the concept of open source; however, the industry is largely still in the growth phase with regard to shared software and hardware blueprint projects for storage products. Some open source vendors like greenBytes, Gluster, Zmanda, Open Source Storage, and Nexenta are trying to build some traction in the storage industry by providing fairly robust operating platforms, management software, security software, file systems, and backup solutions for NAS and SAN environments. Even cloud storage software is available in the open source community via Parascala and Twisted Storage. Zmanda also has a unique Windows-based cloud backup offering that uses Amazon's S3 Storage Cloud as a target for backup and disaster recovery solutions. These companies and many others have done a good

job appealing to the development community by focusing efforts on commercially disruptive storage features, transparent project goals, and well-documented support. Additionally, they are all pushing cost-effective alternatives to simplify the customer acquisition process and enhance storage flexibility without compromising on key enterprise features, such as deduplication, massive scaling file systems, data-at-rest encryption, and content addressable storage (CAS).

Can Commercial Vendors Profit from Open Source?

The economic benefits of free software (with a possible small license subscription fee) appeal in particular to small and medium businesses (SMBs) and universities, but also to some users in enterprise environments. The current recession has helped create a more attractive marketing pitch for open source products, but there are still the same old hurdles that stand in the way of both vendors and customers. As evidenced by some of Oracle's posturing with open source technology (brought on by its acquisition of Sun Microsystems), there is still room for large vendor investment in open source technology, if the technology can stand on its own two feet. Despite concerns about the risks for vendors to have a blended proprietary and open product strategy, an effective product positioning model can bear fruit in the financial books and in customer mind share. If the product has a strong adoption rate, install base, and contrived technology, there will be room for project funding, regardless of the source code nature. Of course, there may also be some concerns around cannibalizing proprietary products; however the ability to provide options to customers will better help them perform a cost-benefit analysis. Meanwhile, the vendors can still keep their chips on the table regardless of the customer's strategic inclination toward the type of technology it wants to implement.

Vendors can expand their perspective on open source by segmenting their offerings. For example, a storage developer could provide a simple no-cost solution with an option for support and maintenance at extra cost, or it could charge incremental fees for additional features that enhance storage software capabilities. The vendor can also bundle open source programs with additional product offerings as a means to grow usage for the open source code, while also creating an incentive to adopt other open or proprietary offerings. The flexibility in designing a business model is as broad as the company's interests and marketing resources will endure.

How Can Vendors Overcome Open Source Stigma?

A number of customers still feel that open source solutions are not as robust as proprietary ones. Many IT organizations are typically so worried about issues like data protection, regulatory compliance, and scalability that they will focus only on the proven storage solutions from well-known vendors. In the storage marketplace, Sun Microsystems, now absorbed into Oracle, has been able to take advantage of this mind set by leveraging its enterprise presence and turning storage discussions around to open source solutions with the Oracle Sun Storage 7000 series arrays running its open microcode OpenSolaris and ZFS file system. Through rigorous testing and certification, Oracle's commercial distributions of its open source software

(Continued on page 4)

Open source continues to evolve through community cooperation and as storage establishes more significance in data management capabilities, open source technology will be a key contributor in driving new production designs.

have eased some of the customer worries. For example, ZFS is claimed to offer "19 nines" of certainty around data integrity, and OpenSolaris provides unique functionality with end-to-end check summing for all data.

Among the emerging vendors in the open source market, Gluster has been one of the more visible companies. Gluster is in the business of offering a scalable storage clustering solution for unstructured file data within a global namespace. The Gluster Storage Platform combines an open stackable file system (GlusterFS), OS, and management capability that run on standard hardware. The Gluster Storage Platform appeals to customers with its simple cluster node install process, petabyte-level scalability, self-healing capabilities, and many other benefits. The software is free for download with the option to buy various subscription and professional service offerings.

In their positioning, many open source vendors are trying to emphasize that high performance, reliability, and scalability are paramount in the design of their products, and that these benefits can be realized by customers without the need for premium priced, high-end solutions. Granted, the use of open source solutions may come at a cost of not having the "single throat to choke" support capabilities that Oracle, IBM, and HP like to tout, but the products certainly appeal to companies that appreciate a low cost of acquisition, ease of install, integration, and management with enterprise features that lack operational complexity.

Even if the features and functionality are addressed by open source solutions, there is still the stigma that support for the technology will be costly due to lack of product maturity, and that open source is too simplistic to deploy in a critical application datacenter environment. In order to change these traditional viewpoints, the value proposition for open source storage

solutions needs to be redirected so that it reflects the economic benefits of standardized technology versus vendor lock-in. If open source is freedom for developers, open standards are freedom for customers. Unlike proprietary storage deployments, open source solutions that adhere to open standards can allow customers to freely mix, match, and maximize components of their computing infrastructures as business needs change. Customers can repurpose and reuse existing hardware simply by adding new software with their own customized or community-based features. Customers can also track community development projects that are in the works, but may not be commercially available. The most popular resource for viewing and downloading open applications is SourceForge, which provides a development portal for open source software products.

The IDEAS Bottom Line

Innovation is very much a driving element in the open source market, with the understanding that the adoption lifecycle slows as unfamiliar elements are introduced. By today's standards, successful open source initiatives are not just defined by features and functionality, but also by the ability to embed the technology in a total solution offering. Companies like Nexenta have cued in to these trends with integrated unified storage solutions, like NexentaStor, which is optimized for virtual environments running on standardized platforms. The intent is to make NexentaStor's commercially packaged solution more appealing through ease of install; dynamic block, file data, and virtual machine management; and enterprise-class features with ZFS at the heart of the software design. It is also worth noting that more mature companies like EMC are actively building open source strategies by developing software products like PowerPath to

run on Linux. EMC also acquired the assets of SourceLabs in 2009 as a possible statement around supporting open source development in the cloud space with SourceLabs' diagnostic and troubleshooting self-support service.

Open source is not about reinventing the wheel, but rather providing unfastened solutions based on customer requirements. Providing a shared technology environment harnesses the brainpower of entire development communities and encourages solutions to real-world customer needs, without the restrictive closed-source corporate agenda. Open source continues to evolve through community cooperation and as storage establishes more significance in data management capabilities, open source technology will be a key contributor in driving new production designs. ■





IDEAS Insights

Blog Bites (<http://www.ideasint.blogs.com>)

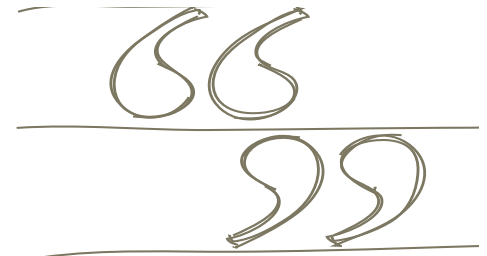
The message is clear. **EMC is ready to take a leadership position in the emerging cloud market.** Nevertheless, it might be EMC's only option. If cloud computing does take over the industry, and customers only accept integrated solution deliveries, remaining a storage-only vendor at its size would be impossible.

EMC Bets Its Future on Private Cloud

Joseph Zhou | May 13, 2010

[read more»](#)

... it is probably fair to say that growth of the web has created new computing workloads, which are suited to a horizontally scaled approach, such as the front-end web serving and related applications. So, with conditions ripe for more general horizontal workload applications and the modularization that gains pace, through the Cloud Computing movement, one would expect the commodity cluster architecture to continue to grow in popularity well beyond this supercomputing niche.



The Growth of Commodity-Based Clusters

Gary Burgess | June 1, 2010

[read more»](#)

About Ideas International

Ideas International (IDEAS) is the leading global supplier of comparative intelligence on enterprise IT infrastructure. Servicing both IT vendors and large-scale IT users, IDEAS products focus on: servers; storage; virtualization and consolidation; and software and infrastructure services. IDEAS online

products are backed up by expert industry analysts who can also deliver specialized advisory and consulting services.

For more information about Ideas International, please visit:

www.ideasinternational.com



Americas

Ideas International, Inc.
800 Westchester Avenue
Suite N337
Rye Brook, NY 10573-1354
USA

Tel +1 914 937 4302
Fax +1 914 937 2485

Asia/Pacific and Worldwide Headquarters

Ideas International Limited
Level 3
20 George Street
Hornsby, NSW 2077
Australia

Tel +61 2 9472 7777
Fax +61 2 9472 7788

Europe, Middle East, Africa

Ideas International Europe
Milton Park Innovation Centre
99 Milton Park
Abingdon, Oxon, OX14 4RY
United Kingdom

Tel +44 (0) 1235 841 510
Fax +44 (0) 1235 841 511