
Market Roundup

May 25, 2007

EMC Avamar Extends to NAS and Virtual
Machines

Microsoft Tackles Identity Management:
Peanuts for the Elephant in the Living
Room?

New Storage Offerings from Big Blue
Cisco and RSA: The Enemy of My Enemy Is
My Friend

HP Announces StorageWorks XP24000 Disk
Array



EMC Avamar Extends to NAS and Virtual Machines

By *Clay Ryder*

EMC Corporation has announced new data deduplication capabilities in EMC Avamar version 3.7, which now also supports VMware Consolidated Backup (VCB) for the protection and reduction of backup times within and across virtual machines. Avamar backup and recovery software features unique global data deduplication technology that eliminates the transmission of redundant backup data across the network that is saved on secondary storage. By deduplicating at the source organizations can dramatically shrink the amount of time required for backups, reduce network utilization, and ease the growth in demand for secondary storage. With Avamar's new support of VCB, which simplifies data protection by offloading backup to a centralized server, VMware customers have an easy way of deduplicating backup data stored in virtual machines. This can reduce the amount of data backed up which in turn minimizes the impact on host servers as well as reducing backup windows and storage requirements. Customers can now leverage Avamar deduplication capabilities with Celerra NAS systems, through Network Data Management Protocol backups. In addition, EMC Backup Advisor now supports Avamar in providing visibility to Avamar backup job completion, as well as diagnostics to analyze failed backup jobs. With this new release, Avamar now supports HP-UX and Mac OS operating systems in addition to existing support for Windows, Solaris, AIX, and Linux environments, as well as adding support for Oracle and IBM DB2 databases. EMC Avamar version 3.7 will be available in June 2007 from EMC and authorized resellers.

As demand for storage continues its unprecedented growth, IT personnel are under increasing pressure to deliver more storage capacity, but at the same time stay within often tight budgetary guidelines. Simply buying more capacity is not the most cost-effective approach to addressing storage needs; rather, maximizing the utilization and efficiency of the existing storage infrastructure is more often than not the sound approach. Judicious use of archiving and ILM can certainly help; however, one of the best ways to reduce secondary storage demand is to eliminate duplication of files resident on disks as well as eliminating duplicative backup traffic across the network. This is where we see EMC Avamar having a great deal of value to offer.

File systems are a wonderful hierarchical storage paradigm, but the many branches of the file tree can mask often significant duplication in storage. For example, a given file that was distributed to several users as an email attachment may be placed in several different directories on a central file store. Not only does this duplicate or waste storage capacity, but when backups are applied to the file store, the duplication is exacerbated as each duplicate copy of the identical file is backed up, sent across the network, and then stored independently on the backup volumes. In virtualized server environments, the duplication can become even more extreme. Consider the system boot and configuration files for a virtual server: each virtual machine accesses the same set of files, yet operationally each acts as if it were its own physical server. Traditional backup schemes would make copies of

each virtual server's environment, and send them off to the backup store. Hence, the greater the success in consolidating the servers the exact opposite takes place with respect to backing up said virtual machines.

Fortunately, Avamar's new support for VMware VCB addresses this irony. By eliminating duplication as an inherent part of the backup process, organizations should be able to significantly reduce the amount of storage needed for backups as well as reduce the overall impact on the network infrastructure in the delivery of backup data. EMC stated that this solution can reduce virtual machines backup windows by up to 90%, a substantial result. As the value proposition of virtualized server environments continues to be embraced by organizations of all sizes, ensuring that efficiency gains in servers are not offset by losses with respect to storage is a paramount concern for end users, the industry overall, and virtualization vendors in particular.

This announcement once again illustrates the thought leadership of EMC in its software acquisitions of the past several years, first with VMware, and more recently Avamar. While organizations could weave together the raw technologies needed to address some of their storage challenges, the integration hurdle often limits deployment and ultimately the value received by organizations for their storage investment. The symbiotic relationship of VMware, Avamar, and other EMC software properties when combined with EMC's hardware platforms illustrate the holistic approach the company is taking with respect to storage solutions, a position that we see continuing to strengthen the company's products and marketplace performance, and most importantly, its customers' daily operations.

Microsoft Tackles Identity Management: Peanuts for the Elephant in the Living Room?

By Lawrence D. Dietz

Microsoft this week announced a series of offerings that foster improved interoperability for online identity management. Microsoft has been a recurring voice in the creation of an identity metasytem, an ecosystem designed to enable the exchange of personal identity information on the Internet so all parties may understand whom they are working with online. Three core elements make up the identity metasytem: the people who are presenting their identities, the website or online service requesting proof of identity, and the identity providers who assert some information about those people. The projects announced today improve interoperability for each of the three metasytem components and represent the next step in Microsoft's commitment to deliver interoperability by design.

In September 2006 Microsoft announced the availability of thirty-eight Web services specifications under the Open Specification Promise (OSP). A subset of those specifications, such as WS-Trust and WS-SecureConversation, addressed identity metasytem scenarios and have led to interoperable identity solutions such as Novell's Bandit project and the Eclipse Foundation's Higgins Trust Framework Project. Microsoft is now making the Identity Selector Interoperability Profile available under the OSP to enhance interoperability in the identity metasytem for client computers using any platform. An individual open source software developer or a commercial software developer can build its identity selector software and pay no licensing fees to Microsoft, nor will it need to worry about future patent concerns related to the covered specifications for that technology. Microsoft is also starting four open source projects that will help Web developers support information cards, the primary mechanism for representing user identities in the identity metasytem. These projects will implement software for specifying the website's security policy and accepting information cards in Java for Sun Java System Web Servers, Apache Tomcat, IBM WebSphere Application Server, Ruby on Rails, and PHP for Apache Web servers. An additional project will implement a C Library that may be used generically for any website or service. These implementations will complement the existing ability to support information cards on the Microsoft Windows platform using the Microsoft Visual Studio development environment.

The growth of identity theft and general public distrust of doing business on the Internet is a deterrent to the overall growth of Internet commerce. Identity management is in some ways the Augean Stables of security. Governments, businesses, and organizations recognize the need to validate the identities of those who access their IT infrastructure but have been loath to put big bucks behind cleaning up the mess. Though many approaches have been offered ranging from Public Key encryption schemes to two-factor authentication featuring tokens with time-limited identity numbers, none have achieved general acceptance. While vendors are jockeying for relative market position, none has emerged as a standard.

We believe the widespread penetration of Active Directory, OpenLDAP Directory, and the capabilities of Microsoft Identity Lifecycle Manager 2007 put these products in the eye of the storm. Over the past several months we have noticed an increasing number of smaller, specialized vendors who are leaning on Active Directory in particular as an anchor to their implementation. We will be watching to see if there are large and powerful early adopters of this technology such as health care providers, payers, and insurers. Once a critical mass has been established, we believe a standard will emerge mostly because everyone wants one, and Microsoft could just be the appropriate party to catalyze a solution to this problem.

New Storage Offerings from Big Blue

By *Clay Ryder*

IBM has announced new unified storage offerings and enhancements across a range of storage products in its portfolio, including virtualization, storage resource management and enterprise disk. The new IBM System Storage N5300 features scalability up to 252 drives and 126TB of capacity; storage tiering with support for both FC and SATA drives; 64-bit controller architecture with high-bandwidth I/O design; and integrated Gigabit Ethernet and 4GB Fibre Channel ports. The N5300 also offers Snapshot copies, a comprehensive set of storage resiliency features including RAID-DP (RAID 6), and disaster recovery capabilities. IBM also announced the IBM System Storage N5300 Gateway and N5600 Gateways that leverage the dynamic provisioning capabilities of Data ONTAP software across existing Fibre Channel SAN infrastructures to enable connectivity of IP applications and user communities to Fibre Channel SAN storage resources. In addition, a new feature, known as Advanced-Single Instance Storage (A-SIS) deduplication, will now be available in the System Storage N series. Under A-SIS deduplication, data is automatically scanned and redundant copies are eliminated, resulting in immediate space savings with minimal performance overhead or operational impact. The IBM System Storage N5300, N5300 Gateway, N5600 Gateway, and A-SIS deduplication for all N5000 and N7000 appliance models will be available on June 8, 2007.

May is shaping up to be the Month of Storage. With three of the big players—EMC, HP, and IBM—all having substantial announcements for the marketplace, it is hard to remember a recent time when there has been so much emphasis on the seemingly mundane world of storage. However, the bevy of product and service-related activity in the market actually illustrates that storage is anything but mundane, and the opportunity for its vendors is as good as it has been for quite some time. The unprecedented growth in data/information being stored in organizations combined with regulatory and compliance pressures have transformed the task of simply storing bytes for later retrieval to an exacting process, if not science, of cost-effectively managing and securing a growing pool of information within the context of business operations. These latest offerings for the System Storage N series help to fill out the already substantial N series portfolio, but we are drawn towards a couple aspects, namely the expansion of the media supported, and most importantly the new deduplication offering.

Not all that long ago, “serious” storage solutions would be typified by high-speed FC disks, with rotation speeds of 10k or faster. These solutions were definitely high performance, but with a commensurate price. It is rather remarkable that in a very short period of time we are now seeing very low-cost and high-capacity SATA and SAS drives now being included as options, often targeted at second-tier solutions, in enterprise-oriented NAS or even SAN offerings. This has changed the potential usage scenario for some classes of storage away from being monolithic in their capacities, e.g., Tier I only, to support multiple classes of storage (mixed Tier I and Tier II or more) within a single solution. This mix-and-match with respect to disk drives offers organizations greater flexibility in aligning service needs with capital expenditure while maintaining and thus managing storage within a single framework. The latest editions of N series gateways take this integration a step further as they bridge the disparate iSCSI, NAS, and FC access methods to unify the underlying storage in its presentation to applications. As we have maintained for some time, storage is logically “just” another resource on the network and as such applications should not be concerned with its physical instantiation. Consistent with this approach is the mixing of tiers within storage cabinets as well as bridging protocols and other elements that have historically created motes around the various storage castles within the enterprise.

Duplicative files are one of the more obvious, yet not always easy to tackle waste in enterprise storage. From an efficiency perspective, having multiple copies of the same content just doesn't make sense. Happily the industry is

beginning to address this in a more strategic nature. One approach to reduce file duplication is deduplication in the backup process, as addressed by EMC Avamar and other solutions; another is to eliminate duplication in the primary storage volumes. The newly announced A-SIS deduplication from IBM is positioned to do exactly this. Since the A-SIS operation takes place behind the scenes, there is an almost immediate improvement in efficiency, especially after the initial deduplication tasks complete. As the deduplication remains in effect when transferring or backing up files across supported tiers, the efficiency gains can be retained, which will reduce the overall impact of backups on the storage and network infrastructure. While deduplication will not end the need for further storage investments, for many organizations it may delay the need and thus smooth out capital spending while also offering the opportunity to gain more value from existing investments for an incremental investment. For many organizations, we expect that this capability will be well received.

Overall, it is clear from the last several days that the major storage players are continuing to up the competitive ante in their storage solutions. We are pleased to see attention paid to deduplication as well as unification, and management efforts surrounding enterprise storage. With the roadmaps that some vendors have released regarding blades and other storage technologies, we expect that 2007 will remain an interesting year for the storage marketplace and ultimately the organizations that choose to avail themselves of the interesting wares that vendors are plying to trade.

Cisco and RSA: The Enemy of My Enemy Is My Friend

By *Lawrence D. Dietz*

Extending their strategic alliance, Cisco and RSA (the Security Division of EMC) have announced they are working together to develop technology to help customers improve and simplify the encryption of confidential information such as medical records, social security data, and credit card numbers. Together, these technologies will help customers encrypt data-at-rest on tapes and other types of storage media and manage the associated encryption keys within the storage area network infrastructure, a process that is less invasive, more secure, and easier to manage.

The planned Cisco/RSA technology will include Cisco Storage Media Encryption (SME), which provides encryption of data-at-rest as a fabric service, and RSA Key Manager, a centralized solution for encryption key lifecycle management. The intent is to provide mutual customers with a highly secure, scalable technology capable of managing both encrypted media and their keys. This SME technology is designed to help enable customers meet various regulatory and privacy requirements by ensuring that confidential information is not compromised if a storage tape or disk is lost or stolen. The impending technology will also support an open API for key management, giving customers the flexibility to deploy stored, encrypted data solutions that best meet their operational needs. Cisco will integrate Storage Media Encryption into Cisco-based SAN fabrics to offer seamless management of data encryption across multiple types of storage devices, such as disks, tapes drives, and virtual tape libraries. By encrypting data in the network fabric, customers can secure data on media that lack native encryption capabilities, such as legacy tapes and disks. The Cisco SAN fabric promises to eliminate the need to manage separate stand-alone encryption appliances. Cisco SME technology promoted as non-intrusive, working seamlessly with existing backup applications.

When two large complementary vendors work together, it is often necessary to look behind the curtain to see if there is more to the relationship than appears at first glance. Cisco is all about the network and RSA is all about encryption. EMC, RSA's parent, is storage. Both Cisco and EMC would like to develop the ways and means of outflanking rivals, especially big ones. This announcement extends their mutual reach into new turf, and sends a shot across the bow of rival Symantec as it moves to expand the position in the storage space enjoyed by the former Veritas product line.

Cisco has moved to outflank Symantec in an attack from another direction with its recent acquisition of Broadware. Not only does that acquisition expand Cisco's presence in the physical security world but it is another step down the road of fostering its customer's use of the IP network as a platform to converge applications from diverse portions of the organization. Video and safety/security systems are a step along the road to mainstreaming the often complex world of SCADA applications not to mention kiosks, special purpose workstations, etc. We believe this announcement is more noteworthy for the growing interrelationships being developed by Cisco than

for the specific products covered. We also believe that both EMC and Cisco will be able to use these new capabilities to further encroach on the market of their competition, notably Symantec.

HP Announces StorageWorks XP24000 Disk Array

By *Clay Ryder*

HP has introduced the HP StorageWorks XP24000 Disk Array, a high-end fault-tolerant storage solution that targets organizations seeking to consolidate workloads, reduce datacenter management costs, and provide 24x7 operational continuity. The XP24000's new processor and provisioning technologies offer increased performance and lower power consumption than previous XP arrays as well as hot-swappable components, non-disruptive online upgrades, dynamic partitioning, and thin-provisioning software. XP StorageWorks Thin Provisioning Software also lowers power consumption and heat generation by reducing the total number of disks required in typical configurations. The XP24000 supports 4GBps interconnections between disks and hosts, and supports up to 1,152 disk drives and 332TB of capacity. It features enhanced HP StorageWorks XP Array Manager software, which allows customers to easily configure, manage, partition, and secure their data through a Web-based interface. With the HP StorageWorks XP External Storage Software, the XP24000 can virtualize HP and third-party storage to consolidate and manage multi-tier storage solutions while supporting up to 247 petabytes of external storage. In addition, the XP External Storage software integrates storage resources while aiding in data migration, array repurposing, and tiered storage. The HP StorageWorks XP24000 is expected to be available in July 2007. The company will continue to sell and support previously announced XP disk arrays, including the XP12000 and XP10000.

This latest offering from the HP StorageWorks line features various incremental improvements as well as more significant upgrades or new capabilities. To us, the most notable aspects of this announcement is the new crossbar and associated 4 Bps end-to-end bandwidth, the increased scale to support up to 579 TB of internal plus external capacity, XP Array Manager, and thin provisioning. The upgrades in the processors, throughput, and scalability are all indicative of the continued investment in storage performance by HP, but they also enable organizations to undertake larger scale consolidation projects than they previously would have been able to. 579TB of storage is not usually in the scale of small organizations; however, it may well be for mid-sized firms or those who are involved in M&A activities, as well as life sciences, media, or other image-intensive verticals. In addition, the increased scale of externally managed resources allows for consolidation, at least from a management perspective of other storage silos, even if the data is not physically moved to internal XP24000 drives.

Although storage vendors of most every stripe continue to improve the performance and scope of their hardware, during the past few years the capability and importance of storage software has increased considerably. At the same time, the sheer volume of data being stored by organizations is growing dramatically, and so have the associated costs for storage software. We are pleased by HP's decision to simplify its licensing model and to provide the equivalent function of several disparate products within its XP Array Manager software at a lower overall price. In addition, the elimination of license fees associated with external storage for raw capacity priced products encourages organizations to bring more of their storage under the XP24000 umbrella. Organizations garner considerably more value and productivity from their storage infrastructure through deployment of operational management software. By moving to a base price plus model, organizations have a simpler cost model to manage as their storage needs grow, but also have a model that is more conducive to departmental chargebacks or other cost-sharing regimens. Since the cost is segregated into base-plus-storage, incremental expense can be more easily matched to the users of the storage. With expectations of increased transparency both inside and outside of organizations, this approach simply makes more sense and allows IT to manage expense more effectively.

Of course, the most effective way to manage expense is to not incur the expense in the first place, or at least to defer the incurrence of the expense until a later time. This is one aspect of thin provisioning that we find intriguing. Through a thin provisioning approach, planning and allocating of storage for long-term needs continue, but the deployment of the storage is staged or otherwise aligned with the actual present need for capacity. As a result, strategic planning continues, which is a good thing, but expenses are in better alignment with

productive value received, which is even better. Perhaps this is why an increasing number of vendors have started talking about thin provisioning.

Overall, we believe the target organizations would appreciate the improved price/performance and scalability that this solution offers. However, from a strategic perspective, the notion of thin provisioning is one that is well positioned to resonate in the marketplace. The combination of enhanced software with simplified pricing, increased performance and capacity, and thin provisioning illustrate HP's desire to improve the price/performance of its solutions while simultaneously bringing customers' capital investments more in line with the current need and value received.