

Zmanda Backup to Amazon S3: Harnessing Cloud Computing to Improve Data Protection for the SME

January 2009



The advent of cloud computing is offering new options to businesses of all sizes. At a high level, the appeal of cloud-based computing is to obtain specific IT services required to run your business on a “pay as you go” basis simply by plugging into a network connection. Services can be quickly started up with no capital expenditure, and can be scaled up or back without having to purchase and manage the infrastructure required to support them yourself.

The availability of internet-based storage, a form of cloud-based computing sometimes referred to as “storage as a service”, provides some particularly compelling options to small and medium enterprise (SME) for improving data protection. Enterprise backup software products have been using disk and tape as media targets for years, but enhancements to some of these products are now also offering internet storage as a media target. Internet storage based around an enterprise-class infrastructure can offer real advantages to SME customers who could not otherwise afford access to sophisticated capabilities such as nearly limitless capacity, service-level agreements, and multi-site disaster recovery.

The concept of using internet-based backup targets with your own backup software may be a new one to you. In this Solution Profile, we explore the value propositions it can offer, and also address some of the common questions people have when first confronting these types of offerings. Then we’ll take a close look at how Zmanda backs up to Amazon S3, providing the detail on how these solution works, how they are priced, and how these offerings can be leveraged to best advantage by SME customers.

Is Business Ready for Cloud Computing?

Consumers have been using online services from the likes of Amazon and Google for several years, and have reached a certain comfort level with these offerings. Businesses, however, have additional concerns and requirements that must be met if they are to use such services. In the data protection space, cloud computing offerings will be evaluated against the same set of

criteria the SME uses to evaluate conventional offerings. The ease of use and low cost advantages are readily understandable, but can these services meet the enterprise’s performance, scalability, availability, data reliability, security, and support requirements?

Culture and trust issues can be a stumbling block for those confronted with the concept of cloud computing for the first time. Can I trust an outside provider to manage my data

S O L U T I O N P R O F I L E

to my specific requirements? The concept of cloud computing is clearly a viable one, and has been validated by successful companies such as Workday and Salesforce.com, who reported Q208 revenues of \$248M and is on track to be a billion dollar company by year's end. Large, established companies such as Nasdaq, Deutsche Bank and SunTrust in financial services are using cloud-computing based solutions and you'll find a number of respected companies doing the same across the business services, distribution and retail, insurance and healthcare, non-profit and education, communications, energy, manufacturing, pharmaceuticals and life sciences verticals. At Taneja Group, we expect that by 2012 some 75-80% of Fortune 1000 companies will be using cloud computing services in some way. Services need to be evaluated on a case by case basis to ensure that they meet specific requirements for performance, scalability, availability, etc. but the question of whether cloud computing is a viable business model for enterprises has, in our opinion, been answered with a resounding "yes".

Conventional Approaches

At their most basic level, backup software products allow customers to set up schedules for backing up servers automatically on a regular basis, track and manage versions of the data backed up in a backup catalog, manage the different types of media to which backups occur, and provide a number of different options for restoring data. Historically, backup media (or targets) were comprised of different types of tape subsystems, but over the last 10 years the use of disk has become much more prevalent. Disk offers

some significant advantages over tape to address issues such as backup window, recovery point objective (RPO), recovery time objective (RTO), and recovery reliability. Almost all backup software now supports the use of either disk or tape as backup targets.

The conventional approach has been to build and manage your own data protection infrastructure. This provides enterprises strong control over their data resources with the trade-off of having to pay for and manage the infrastructure. Storage infrastructure resources are often overbought to ensure that businesses will not be brought to a standstill by a lack of storage space, so there is generally at least a 20% - 30% overhead in terms of unused storage which, if you are managing even tens of terabytes of data, can represent a significant amount of expense. Next, you have to worry not only about how to manage that storage, but how to continue to meet power, cooling, and floorspace requirements as data resources continue to grow at rates of 50% - 60% a year. And finally, in the face of evolving business and regulatory mandates, more and more companies are implementing DR strategies and struggling with the costs and processes of distributing data to and maintaining data at remote locations.

Let's take a closer look at this last point. Regardless of whether companies are backing up directly to disk or tape, they were generally staging the data to tape at some point and then contracting with outside vendors to ship tapes to remote locations where they could be stored for extended periods of time. Restore times for this data were very long since the data had to first be

S O L U T I O N P R O F I L E

found and shipped back to the requester before the tapes were mounted and the particular data sought accessed. Larger enterprises often invested in replication infrastructure that allowed them to replicate the data from source disks at their primary site across a network to target disks at one or more remote sites. This required that the enterprise have access to a second site, an expensive prospect in and of itself, but it did significantly improve both RPOs and RTOs for any restores from this data. To physically pack and ship a tape to these locations generally required several days, but data could be replicated in hours, and sometimes within minutes, after backups completed at the primary sites. Distributing the data electronically also increased the reliability and security of backup transfers, since it completely obviated the risk of boxes of tapes disappearing from trucks during transit. In the new regulatory environment, fines and other costs associated with public relations damage control make this a costly mistake, and there have been numerous high profile cases of this type where hundreds of thousands of customers' privileged information has been exposed in this manner. Distributing backup data electronically is not only faster and easier, but it is more secure. Infrastructure costs, however, have been prohibitive for many SME customers that just cannot afford to set up and maintain a secondary site.

Internet Storage For Backup

The concept behind this latest development is that internet storage is just another target type that can be selected as backup software is being configured. Don't think about this

as the single backup target that is used for all your corporate data. You likely already have some backup infrastructure in place, and unless you're already thinking about replacing it, think about cloud-based storage more tactically. Is a new application being added whose data protection requirements will require you to expand your existing backup infrastructure? Would you like to place backup data for particular applications off-site for DR purposes but can't afford to do so? Maybe you're already using services to transport and store physical tapes off-site, but you would like to move to a faster, more secure way to address this same requirement for particular applications. Or maybe you need to provide data protection for a short term project and don't have the excess capacity available in your existing backup infrastructure. These are the types of situations where cloud-based storage excels in the near term.

Cloud-based backup targets can offer some significant benefits:

Immediate, cost-effective access to storage. Cloud-based computing supports a "storage on demand" model that gives you access to what you need (storage capacity) without having to incur any of the management, maintenance, and upgrade costs associated with the supporting infrastructure. Particularly for smaller projects where you may need to allocate small storage capacities for short periods of time, there is no more cost-effective way to meet the requirement. By contracting for a specific level of service, you can meet your need for high availability storage without having to set up and manage logically

S O L U T I O N P R O F I L E

redundant configurations – it’s fast, simple, and gives you flexibility you may not have when you manage your own infrastructure.

Very inexpensive way to store data in an off-site location. One of the basic tenets of disaster recovery (DR) strategy is that the backup and DR data must be stored in locations far enough apart that they will not both be affected by an event that may require a disaster recovery such as a flood or a hurricane. Generally this requires a second site that has sufficient infrastructure to meet the requirement. Historically, the two options have been to leverage a second corporate site, an approach that can be very expensive, or to store tapes offsite with a company like Iron Mountain, an approach that has its own issues. By backing up directly to the internet, the geographical issue is resolved, and customers pay pennies per GB for this service, a cost structure that is in an entirely different league than that offered by conventional approaches.

No infrastructure management costs. Service providers own and manage the sites where internet storage is hosted. Typically they have a large infrastructure already in place, and are renting out a portion of that infrastructure to you for a monthly fee based on capacity. They have access to sophisticated administrative resources, and often can offer a number of advanced storage management features no small company could afford to build and manage themselves, such as 24x7 availability, geographical distribution, and archiving. No capital expenditures need to be made by you to start taking advantage of the space, or to increase the amount of capacity you need. You simply

request it from the service provider and begin paying an adjusted rate. Likewise, if you decide to shut down a particular application and no longer need the storage capacity for backup purposes, then your rate is adjusted downwards. You don’t pay for space you don’t need, and cloud-based storage allows you to grow or shrink the size of your capacity without having to deal with any of the infrastructure issues.

Hits the OPEX instead of the CAPEX budget. Usage fees are generally paid on a monthly basis based on the services (capacity) used, and as such hit an operational expenditure (OPEX) budget. This gives you a storage capacity solution that is easy to ramp up as necessary without having to incur any capital expenditures (CAPEX) or deal with the accounting ramifications of infrastructure purchases.

Identifying the Concerns

Other than the issue of whether or not cloud-based storage is a viable business model – and we think it is clear even at this early stage that it is – you should evaluate any solutions leveraging internet storage using the same set of criteria you use to evaluate any IT implementation. Will it meet my performance requirements? Does it scale and how easy is it to do so? Can it meet my requirements for high availability and data reliability? What about security? Can it meet any business or regulatory mandates specific to my industry for maintaining privacy? What is the support model? How easy is it to set up and manage? What are the costs associated with such a choice? And finally, am I paying more than I am now to

S O L U T I O N P R O F I L E

get additional capabilities, or does it allow me to continue to provide the same services at lower costs?

The answers to many of these questions will be specific to a particular solution. Instead of discussing them generically, we'll introduce a new joint offering from Zmanda and Amazon, and answer these questions accordingly.

Zmanda's Amanda Enterprise

In 2006, Zmanda took a venerable open source data protection offering, Amanda, and commercialized it. Amanda was originally developed at the University of Maryland in 1991, and has been deployed in hundreds of thousands of installations since then. Zmanda's core product, called Amanda Enterprise, retains all the advantages of open source software – simplicity, support for industry standards, and low cost – and adds usability and enhanced application coverage. In commercializing the offering, Zmanda has provided 24x7 enterprise-class support, maintains updated documentation, releases regular patches, and has provided a managed development direction for the product with a dedicated engineering group.

Amanda Enterprise is a file-based network backup solution built around a client-server architecture. Software agents reside on backup clients, sending the data back to a centrally managed backup server. A single backup server can support up to hundreds of heterogeneous clients, including Windows, Unix, Linux, and Mac. To maintain high backup performance, Amanda Enterprise backs data up to a “holding disk” which acts

as the primary backup target. Once the data is on the holding disk, it can be staged to other backup targets such as NAS, optical, or tape, among others. Amanda Enterprise supports a number of advanced features, including disk-based backup, simultaneous backups to dissimilar backup targets, hot backup using snapshots such as Solaris ZFS and Windows VSS, compression, encryption, and alternative client and server independent restores. Amanda Enterprise uses non-proprietary media formats so that restores can be performed even when Amanda Enterprise is not available, and leverages operating system-resident device drivers (meaning that new OS releases will not break the backup software, a problem experienced by many commercial backup software products that use proprietary device drivers). Amanda Enterprise's most unique feature is probably its Intelligent Scheduler, which allows an administrator to specify parameters within which the software will calculate the backup schedule to optimally smooth resources requirements across the days in each week, dynamically re-adjusting to accommodate disconnected clients. Amanda Enterprise is targeted at the SME, and offers a very compelling alternative to commercial backup software products at 20%-30% of the cost.

Amanda Enterprise Amazon S3 Storage Option

Amanda Enterprise allows customers to stage data to a large number of media types, including internet storage. In January 2008, Zmanda and Amazon jointly announced an integrated offering that allows customers to define Amazon's Simple Storage Service (S3)

S O L U T I O N P R O F I L E

as a staging target through the Zmanda Management Console (ZMC). Customers work with Zmanda as the single point of contact in setting up the service to make start-up easy.

Amazon made its reputation in e-commerce as an online bookstore. Having a scalable, dynamic, and highly reliable IT infrastructure was a competitive weapon for Amazon against its competitors. In 2006, Amazon entered the cloud-computing business, providing access across the internet to compute and storage services which leveraged their sophisticated IT infrastructure. Amazon's first cloud service was S3, and it provided unlimited storage which was initially targeted at consumers. In 2008, Amazon added 24x7 support and a 99.9% service level agreement (SLA) which guaranteed less than 9 hours of downtime per year, making the service much more attractive to businesses who were considering the service for the same reasons that consumers did: low up-front costs, scalability up and down, and IT resource flexibility.

The availability of S3 as just another backup target for Amanda Enterprise offers a very low cost secondary storage solution to SME that has significant flexibility. To choose S3 as a target, simply select it on the ZMC GUI, define a "bucket" name and certification path, and you're ready to back up. Zmanda customers can set up policies to define when data is stored in the cloud (e.g. the daily backup, Sunday's full backup, etc.), and data is then copied to the specified location(s). Companies may decide to back up to local media 6 days a week, distributing a single

full backup weekly to a remote site for DR purposes by using the S3 option. S3 can be the only media target or it can be one of several targets. Customers can specify multiple copies of their data if so desired, and store them in geographically separate regions within the Amazon "cloud" to set up a multi tier DR solution. This gives SME customers instant access to a widely distributed, highly available storage infrastructure for pennies per GB.

Zmanda Internet Backup for Windows

Subsequent to the release of the initial Zmanda Internet Backup product, Zmanda also released a copy specifically targeted for use in Windows-only environments. Like the initial product, Zmanda Internet Backup for Windows offers S3 as an optional backup target (primary or secondary), but it goes beyond the original product's capabilities by offering a comprehensive ability to back up all types of Windows data, including files, applications, database, networked devices, and Windows system state, in a single product. With this product, all data is stored in native Windows formats, allowing easy recovery whether Zmanda Internet Backup is available or not. With its flexible support for cloud-based backup targets, Zmanda Internet Backup for Windows now offers complete protection for any type of Windows server (e.g. Exchange, SQL Server, etc.) in a single, integrated product with all the advantages of Zmanda's original commercialized "open source" data protection solution.

S O L U T I O N P R O F I L E



Figure 1. Choosing S3 as a backup target is as simple as selecting and naming an S3 “bucket” name.

Performance. If you’re already a Zmanda customer, then you are currently backing data up over your networks to the holding disks. With the Amanda Enterprise S3 Storage Option as well as ZIB, that performance does not have to change; holding disks can continue to be primary backup targets. S3 can be defined as a media target, allowing you to still define how many backups you want to keep locally on your holding disk and when they are deleted. By leaving the most recent backup on your holding disk (a common best practice with disk-based backup), you can offer very high performance for most restore requests. The bottom line: while the use of S3 adds resiliency to your environment, it does not affect initial backup performance. It does, however, offer you the ability to improve the both the RPO and RTO performance of your DR solution. Since backup data is catalogued to enable point-in-time restore and distributed to remote locations much faster than can be achieved using physical transport

methods, you may actually be able to improve your RPO and RTO by several days.

Scalability. There is no limit to the amount of data that you can store using S3. Billing is tied to usage, and you can keep data in S3 for as long as you want provided you keep paying the monthly fee. While it is extremely convenient as a way to provide quick access to small amounts of storage, particularly to meet storage requirements that may be short-lived, it can just as well be used to quickly add several terabytes of storage or more to your “virtual” backup infrastructure.

Availability and reliability. High availability was a requirement for running Amazon’s own e-commerce business, and their IT infrastructure was built to provide that. The infrastructure includes technologies, such as multi-pathing and RAID, that support high availability, transparent recovery, online maintenance and repair, and supports Amazon’s 99.9% SLA. The

S O L U T I O N P R O F I L E

infrastructure's ability to support multiple concurrent failures transparently ensures that data is reliably available, while the cyclic redundancy checking performed by the back end storage infrastructure combines with the Zmanda backup software to ensure reliable data transfers and storage. Customers have added flexibility to increase the resiliency of their stored data even further by requesting that it be distributed to multiple geographies by the S3 infrastructure. The cost for this additional resiliency is just the cost of the added utilization (in terms of capacity consumed and data transferred).

Security. Zmanda supports 256-bit AES encryption. Keys for this encryption are retained on the customer premises, ensuring that even Amazon personnel cannot access the data they are storing for any of their clients. They can manage and distribute the data, but they cannot actually read it. Both Zmanda products also support password-based access controls, giving customers full authentication capabilities with internet storage as a backup target.

The use of internet-based storage also provides a faster, more secure data distribution mechanism than the conventional approach of shipping tapes to off-site locations. The risk of losing customer or other important data because physical tapes were lost in transit goes to zero.

Support. These joint offerings provide a single point of support contact, managed through the Zmanda Network Account. Zmanda retains all responsibility for tracking cases to closure, working with Amazon on the back end as necessary to resolve issues.

A single level of service (24x7) is available for the Amanda Enterprise Amazon S3 Storage Option. Support for ZIB, during its limited release, is provided through an actively monitored web-based forum. Both products are supported by Amazon S3's 99.9% SLA, a guarantee of less than 9 hours of downtime a year. For any month where Amazon does not meet this uptime guarantee, customers are entitled to service credits. Amazon is unique in the cloud-based computing space today in standing behind its offering with an SLA.

Ease of use. There are several aspects to bring out that contribute to Zmanda's overall ease of use. First, customers can sign up for the service directly through their Zmanda Network Account. S3-specific registration is facilitated by a re-direct to the Amazon Web Services site. Once this storage starts to be used, Amazon billing starts up with no additional effort on the part of the users. Backup administrators do not need to deal with any of the conventional storage provisioning tasks when setting up S3-based backup targets – this is all handled automatically on the back end by Amazon's infrastructure. The S3 targets inherently know they are being configured as backup targets for Zmanda, so this is easier than trying to configure backup targets yourself with raw internet storage (which thinks of itself as a "copy" destination rather than a backup target). Once the data has been backed up to S3, it can be restored to any location worldwide that has web access, and there are no limits to the amount of data that can be restored provided the administrator has proper authorization. Recovery is as simple as finding the files or directories you want to restore using the Amanda Enterprise

S O L U T I O N P R O F I L E

or ZIB management console and then clicking “restore”.

Zmanda offers more flexibility than competing, cloud-based storage offerings. Many online backup services offer online storage as the only target for backup, whereas Zmanda offers the flexibility to still back data up locally to a disk for high performance restores and to stage data from the holding disk to one or more other targets. Zmanda supports backing up heterogeneous servers (including Windows servers) where many cloud-based competitors support Windows only. Zmanda also puts no limits on how long backups may be retained in the cloud – S3 will store them for as long as you like.

Costs. As would be expected based on Zmanda’s open-source background, the cost of this joint service is quite a bit less expensive than competing offerings based around commercial backup software. Amanda Enterprise Amazon S3 Storage Option is licensed on a subscription basis, typically at prices which are 20-30% of commercial offerings. ZIB, during its limited release, will charge a one-time fee of \$50 per Windows server protected. The Amazon S3 pricing model consists of two components: capacity and volume of I/O traffic. Users are charged \$.20/GB per month for data stored plus \$.20/GB per month for data transferred into the S3 infrastructure. Usage is tracked and made available to customers through a secure web portal. Recovering and transferring data out of S3 using Zmanda is free. S3 has data centers in the United States and Europe, and users will spend a little more to store data in European data centers.

Competitive cloud-based offerings charge per server fees (which Amazon does not) as well as fees for storing and transferring data, often at rates that are 2-3 times as high as S3’s prices.

When comparing the costs to more conventional methods, you’ll need to take into account the cost of the required infrastructure, administrative overhead to create, pack, and manage backup tapes, the cost of the service to ship tapes, and ongoing media replacement costs. When it comes to the “hassle” factor of dealing with off-site tape transport, Zmanda Internet Backup is the clear winner. It’s faster, easier, more secure, supports much better RPOs and RTOs, and is quite a bit less expensive.

Getting Started

If you’re new to using internet-based backup targets, there are some operational considerations to take into account as you think about how to integrate Zmanda backup to Amazon S3 into your environment.

Backup set selection. Customers can select one, many, or all backups, setting up schedules for staging them to S3. The choice will be driven by the particular data protection problem you are trying to solve. If you are trying to set up your company’s first DR solution, then determine the critical servers and/or applications which must be protected by storing data at a remote site. If you are already shipping physical tapes to a remote site, then you may already have a list of servers and/or applications that will be on this list. If, on the other hand, you are trying to distribute your backup data to multiple

S O L U T I O N P R O F I L E

geographical locations to provide faster access, you will select backup sets on the basis of the applications that require this improved distributed access.

In the case where you may have a transient project that needs backup protection, you can perform backups to the local holding disk and then immediately stage them to S3 to free up the local space. Backups can be scheduled to occur so as to make the most of existing holding disk capacity. This flexibility makes it quick and easy to add the capacity needed to provide data protection for a particular project without having to increase the capacity of your existing infrastructure.

Next, you need to determine the RPO and RTO you intend to meet. Electronic distribution generally allows much more stringent RPOs and RTOs to be met than the shipping of physical tapes can. Once you've defined your RPO and RTO, you know how frequently you want to schedule backup and stage upload to S3. If you are staging backups once a week or less frequently, you will want to stage full rather than incremental backups. Full backups will generally require more bandwidth, but tend to support faster, easier recoveries than incrementals or differentials (which require more steps to complete the recovery process). If you plan to stage daily backups to S3, then you will likely be staging incrementals and/or differentials as well as the periodic full backup.

Performance. Zmanda leverages a local holding disk as the primary backup target, and primary scheduling of backups can be

handled automatically by Amanda Enterprise or ZIB's Intelligent Scheduler or through specific schedules set up for certain backup sets that operate independently of the Intelligent Scheduler. Determining when backups are downloaded from the holding disk to one or more S3 targets can be scheduled to occur at times when WAN usage by other applications is relatively lower. As mentioned earlier, most users leverage S3 as a staging target to store selected backups, such as a weekly backup, or to implement cost-effective DR plans. Backups do not need to be immediately staged to S3 as they occur – they can be held in the holding disk and scheduled to occur at times when WAN bandwidth is normally quite low. For example, if completing Sunday's weekly backups takes up most of the backup window available on Sunday, those backups may not be staged to S3 until later in the week. Understanding the peaks and valleys in your network bandwidth usage is an important prerequisite in leveraging S3 in a minimally impactful manner.

Amazon has multiple data centers located in North America and Europe that support the S3 service. Additional data centers will likely be brought online in the Asia Pacific region in the future. For performance and/or DR reasons, you may want to identify specific S3 targets, located in certain geographical regions, as your primary staging targets. For example, if you are located in California, but also have operations in France, you may want to distribute data to Amazon data centers in the American Midwest and to a continental location in Europe outside of France.

S O L U T I O N P R O F I L E

Security. Best practices for maintaining the confidentiality of data prescribe that it be encrypted both in-flight and at rest. Amanda Enterprise offers options for encrypting data at either the backup client or the backup server. Encrypting at the client requires client resources, but ensures that the data is secure anytime it crosses a network, regardless of whether that is a LAN or a WAN. If you're more interested in the use of encryption when data traverses a WAN, then you can encrypt the data after it hits the local Amanda Server. By encrypting data locally on the Amanda Server, Amanda Enterprise protects it while it is at rest in the holding disk(s) and ensures that it is kept secure as it crosses the WAN. ZIB offers the option to encrypt data over the network or on your disk using strong 256-bit SSL and AES encryption. Data should be left in its encrypted form as it is staged to S3, with the encryption key(s) retained at the customer site in the local backup server.

Taneja Group Opinion

The introduction of internet storage as just another backup target type offers a lot of value to SME customers who would have no other way to implement DR plans or respond to requests for additional capacity in a timely manner. It leverages a business model that has been proven for business use not only by enterprise applications providers like Salesforce.com but also by online backup service providers that have been successfully selling into the enterprise for the last 4-5 years. Hybrid services that leverage a backup product and an internet storage backup target like Amanda Enterprise and ZIB hit a new price point which is roughly an order of

magnitude lower than what online backup service providers offer, and because it uses a customer's own backup software, the hybrid model gives end users more control over how their data is backed up. We think that, over the next 12-18 months, you'll see many other players enter this arena.

There are several best practices we can suggest for customers leveraging Zmanda backup to S3 solutions:

- Cost-effective DR for SMEs that could otherwise not afford it
- Fast, easy ramp-up of storage capacity needed to support data protection infrastructure for new or short term projects
- Simple, inexpensive way to provide secondary storage capacity to remote offices that may not have much IT infrastructure
- A more secure way to handle the geographic distribution of backup data than physical tape transport can offer
- Improve the overall availability of your backup infrastructure to 99.9% while at the same time lessening your dependence on tape-based backup

The combination of Zmanda and Amazon provide some unique differentiators that are meaningful to the SME. First, it is based around Amanda Enterprise, the commercialized version of Amanda, arguably the industry's most mature open source backup software. As such, it leverages open source advantages like use of industry standard media formats and device drivers, and provides a mature, feature-rich data

S O L U T I O N P R O F I L E

protection solution which fits well with the requirements of SME – all at price points that are 20%-30% of alternative commercial software offerings. Amazon offers a highly scalable, highly available storage infrastructure that is managed by seasoned storage professionals, makes it easy to get storage on demand, and backs this service by a 99.9% SLA – all at aggressive price points as well. While we do not think that internet storage is yet appropriate for most primary business data, we do think it is a smart way to add flexibility to existing IT infrastructures

for storing secondary data. Together, these two companies provide a joint offering with a single point of support that makes it easy for IT organizations with very limited budgets to implement very effective DR plans. If you are already a Zmanda customer, it would pay to consider how this offering can help you address some of your existing storage challenges. If you are not yet a Zmanda customer, it is yet another feature that makes Amanda Enterprise and ZIB compelling data protection solutions for the SME.

***NOTICE:** The information and product recommendations made by the TANEJA GROUP are based upon public information and sources and may also include personal opinions both of the TANEJA GROUP and others, all of which we believe to be accurate and reliable. However, as market conditions change and not within our control, the information and recommendations are made without warranty of any kind. All product names used and mentioned herein are the trademarks of their respective owners. The TANEJA GROUP, Inc. assumes no responsibility or liability for any damages whatsoever (including incidental, consequential or otherwise), caused by your use of, or reliance upon, the information and recommendations presented herein, nor for any inadvertent errors which may appear in this document.*