vizioncore™
A QUEST SOFTWARE COMPANY

backup 2.0
Simply Better Data Protection

Backup 2.0: Simply Better Data Protection

Gain Net Savings of $15 for Every $1 Invested on B2.0 Technologies

Executive Summary

Traditional backup methods are reaching their technology end-of-life. Designed more than two decades before current datacenter infrastructure, these systems use inefficient methods for data collection and transmission, and offer limited recovery options. They are also dependent on their own set of complex infrastructure that includes backup agents, servers, networks, and storage systems. This complexity adds cost and administrative burden to already over-burdened environments. Ultimately, traditional backup methods – termed Backup 1.0 – are proving to be too costly to operate and maintain, with insufficient recovery options.

Beyond cost problems and recovery limitations, Backup 1.0 systems are also out of step with modern datacenter priorities. Modern priorities include initiatives that result in infrastructure consolidation, cost-cutting, and simplification of operations. Far from helping to further these initiatives, Backup 1.0 systems run counter to them: B1.0 systems add complexity, consume datacenter floor space, and increase energy consumption. Data protection systems must evolve to support current trends by taking advantage of new server capabilities.

Backup 2.0 solutions are the next generation in data protection technology. The break-through that enables "simply better data protection" is the use of system images to make backup copies rather than individual files. Virtual servers create image files which encapsulate Virtual Machine systems. Protecting these images turns out to be far faster and easier than scanning for the thousands of individual files that images represent. With simple conversion tools, equivalent image files can also be created for each physical system. By applying the advantages of image-based data protection learned in virtual environments to physical systems, simply better data protection is made available with Backup 2.0 solutions for all environments.

Using images reinvents how data is collected, transmitted and recovered. Image-based backup collects and protects more types of data, transmits and stores it more efficiently, and offers faster recovery at more frequent points in time of more types of data. The resulting Return on Investment on image-based data protection is significant; conservative estimates show at least $15 returned for every $1 invested in Backup 2.0 solutions.

Unlike migration between backup software brands, adoption of Backup 2.0 solutions is fast and easy in most environments. Backup 2.0 solutions can be added to Backup 1.0 deployments. Configuration options enable Backup 2.0 jobs to be merged into existing Backup 1.0 backup cycles. Console systems already in place can be used for environment-wide scheduling. Proxy Servers and attached SAN storage can be leveraged for secondary sweep-to-tape of all protected data in a single job. These Consoles and systems are already familiar to backup administrators, which should help to ease the adoption process. Operating a single Console system that manages all data protection for the environment is a priority for most organizations, for which the task of managing a variety of tools and technologies for data protection is already familiar.

Backup 2.0 is expected to receive incremental adoption into existing environments. Over time, as the benefits of B2.0 methods are experienced by adopting organizations, this adoption is expected to accelerate.

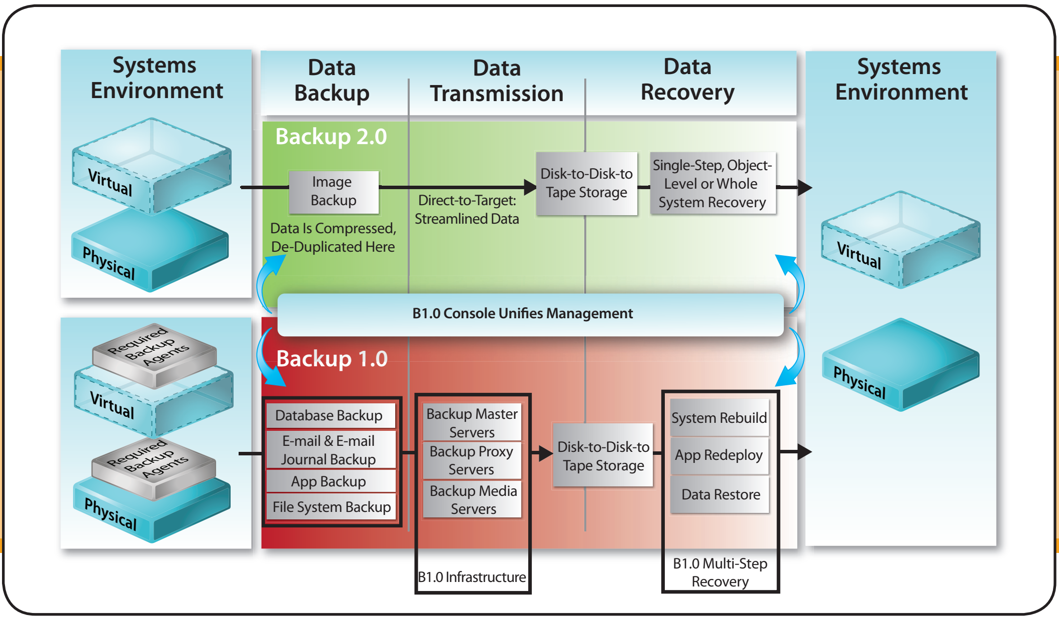

Figure 1. Backup 2.0 Solutions Offer Simply Better Data Protection by Reinventing How Data is Collected in Backup, Transmitted, and Recovered

Estimating Backup 2.0 Return on Investment

Table 1 provides a break-down of how Backup 2.0 methods offer significant savings and Return on Investment (ROI) in all of the key steps required to protect data. ROI calculations are based on conservative cost estimates for traditional backup agents, servers, network and storage resources, and productivity impacts.

Model Assumptions:

- This model assumes that at least one backup agent can be removed from each system in a deployment; in practice, most systems require more than one backup agent. Depending on the number of backup agents deployed, ROI savings opportunities can be substantially higher.

- This model assumes that the time required to manually rebuild a server impacts the business for the time of the rebuild. On average, it takes about six hours to rebuild a server with about $1K per hour of business impact during that time. However, different types of servers have different business value. Depending on the nature of the server andthe time required to rebuild it, the business costs from lost productivity can be much higher. In worst-case scenarios, prominent studies show that the loss of some types of servers for more than a few days can threaten to bankrupt a business.

- This model assumes complete adoption of B2.0 technology onto all systems in the deployment. In reality, the adoption of B2.0 technology is expected to happen more gradually within a given environment. The exact ROI realized in an environment will depend on many factors, including the pace of adoption.

Table 1. Backup 2.0 ROI Offers Net Savings of at Least $15 for Every $1 Invested

Backup 2.0 ROI Summary	Data Backup	Data Transmission	Data Recovery
Efficiency Gains	• No Agents • Less Time - No Scan • More Types of Data Collected • Less Burden to Manage	• No Backup Servers • One Backup Copy-Backup-Once-Restore-Many • Less Network • Less Storage	• Productivity Gains • Faster Recovery • Recovery of All Types of Data • Less Burden to Mangae
ROI Estimates	• No Agents • ~$400 per Agent Provisioning Cost • ~$400 per Agent On-Going Mgmt Cost	• ~$3K per Backup Server • 2:1 Savings on Networks • 2:1 Savings on Storage	• ~6 Hours to Manually Rebuild a Server • ~$1K per Hour in Productivity Cost • ~$6K per Server Crash
NET Savings Opportunity per 20 Systems (physical or virtual machine)	• NET $14K for 1 ESX Server running 20 VMs • NET $10K for 20 physical systems	• NET ~$9K for 3 Removed Backup Servers • NET ~$6K for 2:1 Data Compression on Networks and Storage	• NET $18K for 3 System Crash Recoveries

Why have Backup 1.0 Methods Failed?

Back when Backup 1.0 methods were invented, file systems were smaller and the overall size of data in an environment was orders of magnitudes less than what is found in modern datacenters. Scanning file systems to collect data for backup was logical: it was easy, and it didn't take very long. Sending data over networks and retaining multiple copies of the same data on storage was likewise straightforward and cost-effective.

As datacenter deployments have grown and became more complex, however, Backup 1.0 methods have struggled to keep pace. Complex topologies for handling data volumes, like Network Attached Storage (NAS) and Storage Area Networks (SAN) architectures, were invented to alleviate the burden that backup data put on business networks. Shared storage architectures also make large storage systems easier to share among backup systems. Disk-based backup alleviated the performance challenge associated with streaming data to tape. Deduplication technologies have more recently attempted to further streamline backup, by making notoriously inefficient backup copies less wasteful by removing duplicate blocks and files from stored backup copies.

Despite all of this effort, however, backup systems remain more of a problem than a solution in most environments. With all of their agents, workloads, jobs, and servers, Backup 1.0 systems are simply too costlyto purchase and maintain. Provisioning the infrastructure required to operate Backup 1.0 and keep jobs running at sufficient performance levels has consumed a larger and larger percentage of IT budgets every year. Likewise, the burden required to operate, troubleshoot, and maintain these complex backup systems is too high for already over-burdened administration teams and requires dedicated specialists just to complete the basics.

The expense of traditional data protection systems comes in large part from their complexity. Disk scanning methods require expensive, specialized backup agents integrated to work with file systems and applications. These agents must be installed and maintained on every system that organizations need to protect. The agents must be maintained with versions that match the applications and file systems. Licensing must be tracked and maintained. Traditional methods also require a large number of backup servers including Master Servers, Media Servers, and Proxy Server systems. The cost of provisioning all of these agents and servers is large. The on-going cost and burden of keeping this software operational has proven to be too heavy for most IT budgets.

With all of this cost and effort, it's perhaps a surprise that Backup 1.0 methods fail to protect environments sufficiently. In most environments, getting all of the data that needs to be protected copied in the available backup window is still an enormous daily challenge. The result is that, often, critical data is not protected at sufficient intervals and an organization's Recovery Point Objectives cannot be met. Even when data is protected, meeting Recovery Time Objectives is still a challenge because recovery options are limited. Recovery options do not include whole environments and systems, and do not easily support restore of data across systems and onto dissimilar systems. This means that recovery often starts with a complex, time-consuming system rebuild which must be completed before restore of application data can even begin. Protecting whole systems requires add-on of expensive bare metal recovery tools, the creation of extra backup copies, and the management of extra steps in restore.

What's Different with Virtual Servers?

Virtual servers operate differently than physical servers, and create data files at a faster pace than equivalent physical systems. Each time a new Virtual Machine is created, which is often, a new Virtual Machine file is written. Protecting these VM files, called images, is critical for enabling fast and comprehensive recovery of Virtual Machines. Protecting an image also protects all of the individual data files in the Virtual Machine. In this way, image-based backup enables the recovery of individual files and application objects along with complete systems.

The challenge in protecting system images is that these files are very large and complex to capture coherently. Traditional backup methods simply take too long to work, and saturate server, network, and storage resources in the process. They are also hard to install and manage for Virtual Machines, because most Backup 1.0 systems were not designed to run at the hyper-visor level. As a result, most organizations using traditional backup methods do nothing to protect Virtual Machine images. When they do, the cost in time, wasted resources, and productivity delays is unnecessarily high. Traditional data protection methods are not only insufficient for protecting all types of data; they are also too burdensome to operate. Due to their heavy load, traditional backup jobs must be scheduled to run during windows in which end-users are not trying to work. The load impact is worse on virtual servers, where the slowdown of one Virtual Machine affects every VM on the server. The biggest challenge for backup administrators working with these tools is getting backup jobs to complete successfully before the backup window expires. As a result, most data simply cannot be protected frequently enough.

What's Better with Backup 2.0?

Backup 2.0 solutions use image files to reinvent how data is collected, transmitted, and stored for protection purposes.

Image files encapsulate an entire system, including the system state, OS and application configuration data, and the many thousands of individual system and application files. Virtual Servers automatically create image files for every Virtual Machine. On physical systems, simple conversion tools can be used to create equivalent images.

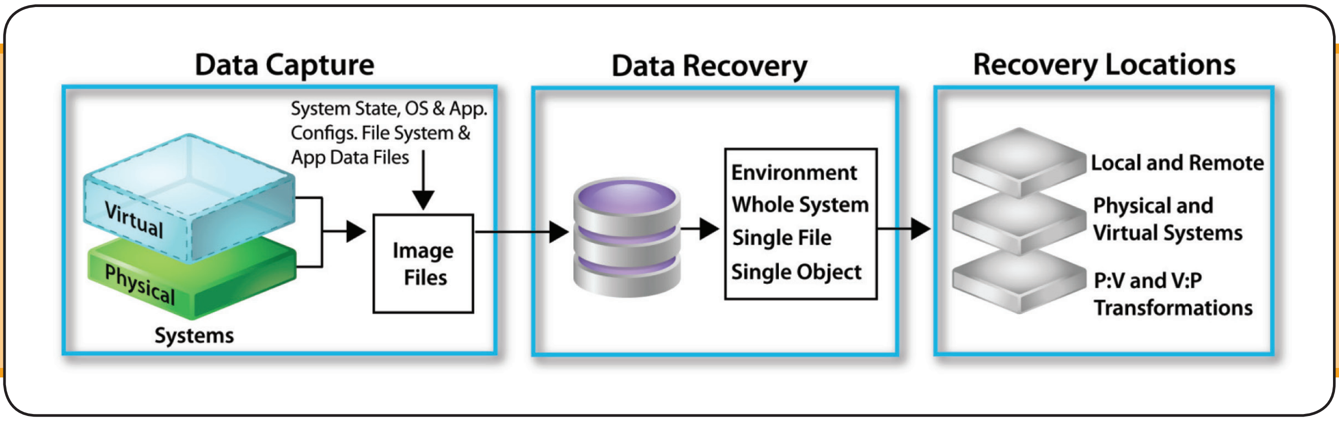

Figure 2. Image-Based Data Protection is Faster, Protects More Types of Data, and Enables Better Recovery Options

Because Backup 2.0 methods do not have to process the thousands of discrete files that images actually contain, they work very quickly. The information captured in each VM is better than what is possible with traditional backup, because it includes not only the application data but also the system state and configuration. Recovery options are dramatically expanded, to include the rapid reinstatement of entire systems complete with applications along with individual data files or application objects. Backup processes are simplified because individual application agents are no longer required. As a result, Backup 2.0 software costs less, and is easier to deploy and maintain, because there is simply a lot less of it.

To accommodate the extreme size of image files, Backup 2.0 solutions must integrate efficiency capabilities to make them effective. These efficiency capabilities include deduplication, zero elimination, and active block management to handle deletions. In combination, these efficiency capabilities dramatically reduce the size of the image data being protected.

Understanding Backup 2.0 Architecture: Two Methods

Backup 2.0 solutions can be built using one of two architectures: they can rely on Proxy Servers found in traditional data protection deployment or they can be built using Direct-to-Target methods.

Solutions that rely on Proxy Servers require server systems to manage the collection of data from clients, and to move that data from the clients through the Proxy to the attached storage devices. Part of the job that a Proxy Server does is to manage access to shared storage devices, devices which are typically attached to the Proxy on a Storage Area Network (SAN). In these deployments, in fact, SAN deployments are often required raising the overall cost of the deployment still higher.

Rather than sending data through a backup server, Direct-to-Target Backup 2.0 systems send collected images directly from source servers to target storage devices. Direct-to-Target implementations deliver the maximum performance available from network, I/O, and storage resources by eliminating the bottleneck of sending all collected data through a Proxy Server. There is no need to manage access to shared storage, because all storage on the network is available to every source system. This approach also avoids the added complexity and management overhead of having one or more Proxy Servers and SANs in the environment. In fact, this approach consolidates Backup 2.0 server infrastructure to the bare essentials which further reduces backup costs and associated burdens.

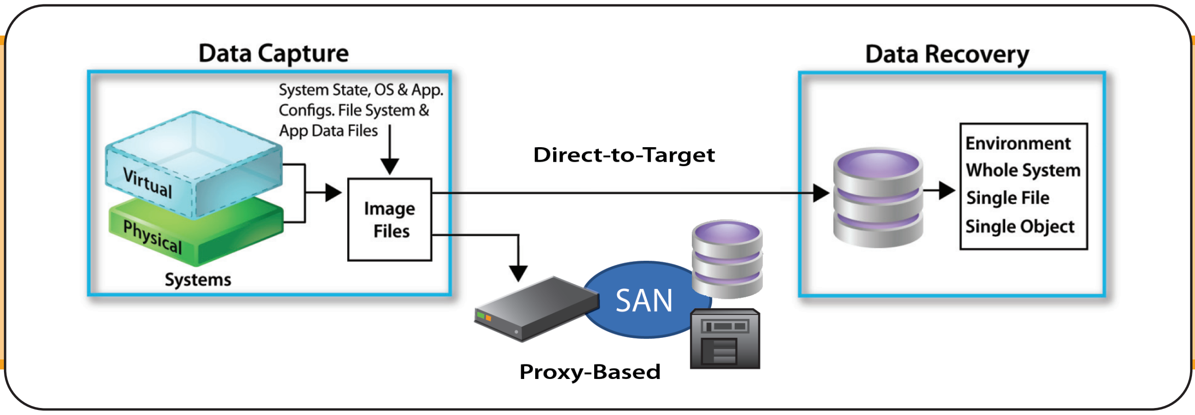

Figure 3. Direct-to-Target Backup 2.0 Methods Offer Shortest Possible Backup Windows and Fastest Recovery; Use of Proxy-Based Method Preserves SAN Investment

As with backup, Direct-to-Target architectures speed restore. Direct-to-Target systems do this by sending images directly from storage to client systems. This avoids the Proxy-Server bottleneck and enables many simultaneous recovery processes to occur.

Direct-to-Target implementations raise ROI still more by creating an architecture uniquely suited to adding efficiency capabilities. With Direct-to-Target, efficiencies can be added where they have the most benefit, which is right at the source of the data. At the source, data can be collected with better efficiency, and better compressed and de-duplicated before it is transmitted over networks, and before it is stored. Advanced Efficiency capabilities which are considered critical for handling image-based backups and which should be present in Direct-to-Target implementations include:

- **Zero Elimination** – removes empty blocks to provide image compression before transmission
- **Active Block Management** – removes application and file data that has been deleted from within the image before transmission
- **Global, Inline Deduplication** – eliminates duplicate blocks from within images globally from across the entire protected environment
- **Integrated Change Block Tracking** – uses system tracking of changed blocks to avoid the need to scan systems during incremental image-based backup

In combination, Direct-to-Target architecture combined with Advanced Efficiency capabilities reduces the amount of data that must be protected across an entire environment. This makes data collection faster and more efficient. This makes data transmission faster and less bandwidth intensive. This also reduces the amount of storage required to protect and retain backup images. Equally important, recovery times are improved because there is less data to restore. In short, Backup 2.0 solutions with Direct-to-Target architecture are necessary to achieve the smallest global backup windows, Recovery Point and Recovery Time Objectives.

Organizations may have significant investment in SAN-attached storage managed by Proxy Server systems already present in their environment. In this case, Backup 2.0 solutions can leverage these systems for LAN-free backup and restore. SAN resources can also be shared to store B1.0 and B2.0 protected data, including disk and tape systems.

Obtaining Backup 2.0 Benefits in Backup 1.0 Environments

In environments using Backup 1.0 systems, Backup 2.0 methods can be added without removing older systems. Most traditional backup deployments treat Virtual Machines as physical servers, and do nothing to protect the VM itself. In these environments, the first priority must be to get a VM protection strategy in place.

In this case, removing the legacy backup environment is not necessary. Instead, incrementally adding Backup 2.0 to work seamlessly alongside traditional Backup 1.0 environments is the recommended approach. Using the same systems already in place, including those from Symantec, CommVault, and Tivoli, Backup 2.0 image-based backup can be managed as part of the traditional protection process.

This process of adding Backup 2.0 capabilities to Backup 1.0 environments is simple working with vRanger, the best Direct-to-Target Backup 2.0 solution available in the market. Proxy-based Backup 2.0 implementations have performance bottlenecks and lack client-based efficiency capabilities required for large-scale image-based protection. vRanger's Direct-to-Target Backup 2.0 implementation is scalable and can be configured to work as part of the backup cycle of existing systems . In this way, the value of existing Proxy Servers is preserved because they are not overloaded with primary backup and can still be used for sweep-to-tape.

Moreover, no additional operational steps are ever required to manage backup. Likewise, once deployed no additional changes or adjustments in the vRanger backup configuration are required to find new images, as these are found automatically using a query process that is part of vRanger's implementation. The Consoles from existing backup deployments can be used for environment-wide backup scheduling, for all Backup 1.0 and 2.0 processes, which helps to unify the backup environment.

All of these features add up to fast and easy adoption of vRanger Backup 2.0 into existing environments, with no additional infrastructure. In fact, the value and effective life of existing Backup 1.0 deployments is preserved and maximized with vRanger Backup 2.0.

Conclusion: Adopting Backup 2.0 Technologies is the Next Step in Evolving Datacenter Data Protection Strategies

As organizations continue to consolidate datacenter infrastructure while attempting to accommodate exponential year-on-year data growth, data protection solutions must improve. IT budgets have been consumed just keeping up with data protection costs. Despite all of the spending, data protection remains unreliable as traditional file-based backup jobs do not have the time required to complete, fail often, and offer limited recovery. The result for business is high costs and low data availability, along with stressed IT administration teams struggling to make it all work.

The source of the problem is out-dated backup methods which were originally designed for data volumes and infrastructures present two decades in the past. Moreover, the advent of virtual server technology offers better access to all types of data. Using images which encapsulate Virtual Machine data, all data can be protected without having to scan disk storage systems. These techniques can also be applied to physical systems to reinvent how data is collected, transmitted, and recovered in all environments.

Backup 2.0 offers simply better data protection with significant ROI results. For every $1 spent on B2.0 technologies, the adopting organization can expect to reap at least $15 in savings within the first year. Also, adopting is not a rip-and-replace exercise within an environment. Instead, B2.0 technologies can be added to B1.0 deployments. In these deployments, the value of B1.0 components including Proxy Servers and Consoles is not only preserved but enhanced.

The business stakes are high. If enough IT budget can be freed from data protection, then more strategic initiatives can be funded. Business competitiveness will improve. IT teams will get out from under the heavy weight of data management.

vRanger: The Best B2.0 Solution

- First to market with image-based capabilities
- Most widely used with over 12,000 customers
- Recognized market leader in virtualization management by Gartner and IDC
- First and only product to feature image encapsulation for physical systems alongside virtual systems
- First product to incorporate Direct-to-Target architecture
- Easy and fast to adopt into existing backup deployments, working alongside console systems which can be used for:
    - environment-wide scheduling
    - sweep-to-tape
- True image handling delivers Backup-Once-Recover-Many, Any-to-Any recovery
- Backup of a few images is faster than backup of thousands of discrete files
- Recovery is faster and more reliable for full environments, individual servers, single files and application objects
- Direct-to-Target advanced efficiency capabilities speed performance and reduce network and storage consumption
    - Zero handling
    - Deleted data handling
    - Global, inline deduplication
    - Change block tracking
- Image-based backup is available for both physical and virtual systems
- Migration of physical systems into virtual servers is easy

vizioncore™
A QUEST SOFTWARE COMPANY

2300 Barrington Road, Suite 700 | Hoffman Estates, IL 60169 | Toll-free US: (866) 260-2483 |
International: (847) 589-2222 | Fax: (847) 279-1868

Vizioncore Inc, founded in 2002, is a fully owned subsidiary of Quest Software Vizioncore offers a comprehensive set of solutions that work well together, are easy to use and affordably address the range of issues that IT managers face in trying to improve the performance of their environments, both virtual and physical. Globally, over 18,000 Enterprise and SMB organizations use Vizioncore's products, with an extended partner network of 2,200 value-added resellers. For more information please visit www.vizioncore.com.

backup 2.0