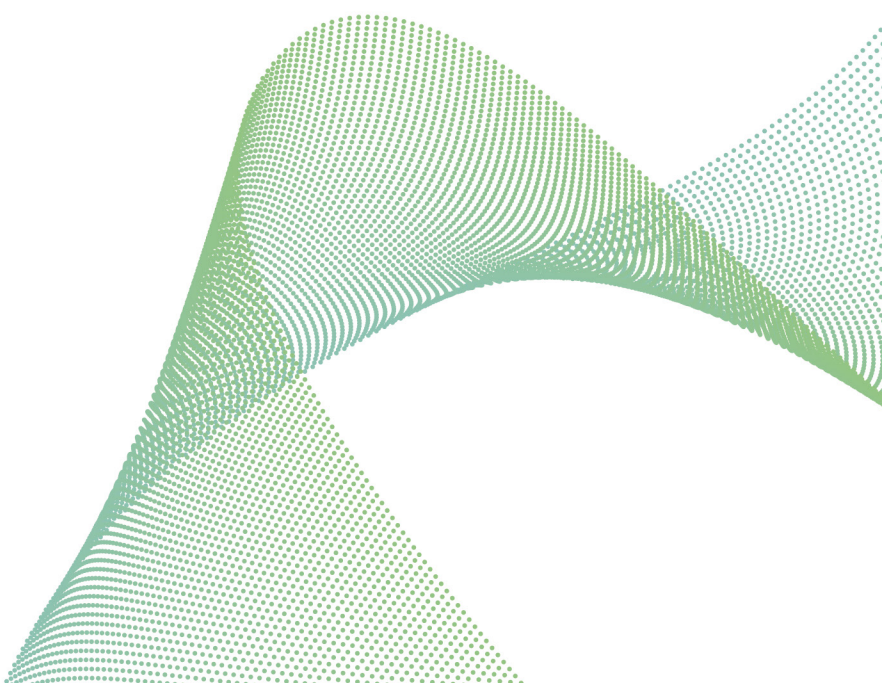


Backup and Disaster Recovery: What you need to know about data reduction techniques!



Contents

- Executive summary3**
- What data reduction techniques are available?.....4**
- What is deduplication?5**
- What is deduplication designed to solve?6**
- How is dedupe often implemented?7**
- Is this too good to be true?8**
- New uses for backup data9**
- Which data reduction technique is optimal?10**
 - The Daily Backup Storage Strategy10
 - Weekly Backup Storage Strategy11
 - Monthly Backup Storage Strategy11
- Conclusion.....12**
- Appendix.....13**
 - Links13
- About Veeam Software14**

Executive summary

It is important to note:

Data deduplication is not an optimal storage technique for every backup-related workload. Applications which make use of backup data as the source to power business continuance, development, testing, quality assurance and data warehouse population, all benefit from faster back-end storage I/O than is available with deduplication. Furthermore, the cost assumptions which once necessitated deduplication for disk-based storage of backup data have changed as the industry has advanced.

As a technique used to reduce the amount of storage required for backup data, data deduplication uses sophisticated algorithms to minimize redundancy. Deduplication has proven particularly valuable for backup applications which historically contain large amounts of duplicate data from one backup to the next. The technique reduces the amount of storage required for backup applications to a fraction of what was formerly required in the tape era. Lowering storage capacity requirements reduced costs associated with storing backup data, making it economically feasible to back up data to disk.

Since the introduction of deduplication over ten years ago, the industry has advanced, changing the math affecting the cost / benefit associated with deduplication. What has changed?

- The cost of disk storage has dropped significantly
- Storage density increases allow many times the amount of data in the same physical data center space
- Electricity requirements for power and cooling per gigabyte of stored data have dropped
- Sub-file, block-level change-tracking backup techniques reduce data amounts needed for incremental backups

At the same time, business requirements for faster recovery and the use of backup data for new applications increase performance demands on the backup infrastructure.

- Business has less tolerance for long recovery times than ever before
- Instant recovery of many simultaneous systems requires storage performance and throughput unachievable on legacy dedupe platforms
- Automated disaster recovery (DR) testing offered by technologies like Veeam® SureBackup requires fast storage for recoverability testing of more mission-critical applications and systems
- Test, development and quality assurance applications which use backup data work best when backup storage is fast
- Data warehouses and decision support systems may source their daily updates extracting daily changes from a copy of the backup data, requiring near production-class storage

This paper will examine the trade-offs between data reduction techniques and the business value associated with the use of backup data for purposes other than backup restoration. We will look at the monetary and time costs associated with various available techniques. And finally, suggest an approach to take advantage of the best available technologies for a typical modern use-case.

What data reduction techniques are available?

On its face, data reduction techniques for backup are designed to minimize the amount of data required to represent a full system image, consisting of the operating system, applications and data. Because backup applications typically keep many historic versions of the system image being protected, much redundancy exists from one version to the next. **Compression** as a data reduction technique looks for patterns in the backup data stream which may be reduced. **Deduplication** reduces data across an entire file system by comparing new backup data against that which has been previously stored. When an exact duplicate to a portion of the data is found, rather than writing that same portion of data again, the file system creates a pointer to what has already been stored. **Incremental backup** techniques look to record changes from one backup to the next using a complete image, or full image, as a starting point, then keeping track of incremental changes typically across set periods of time.

The available data reduction techniques may be combined for greater effect. For example, a backup application will typically use an incremental technique which writes data to a deduplication file system. The deduplicated data will then typically be compressed as it is written to storage. The net effect on the size of the overall data reduction is significant.

Modern incremental backup techniques function at a block, or sub-file level. Only the changed blocks of data are passed along to be deduplicated and compressed. The amount of redundant data from one backup to the next is greatly reduced, leaving the deduplication process with many fewer opportunities to eliminate redundancy. The net effect of block-level incremental backup techniques are data reduction results close to those of deduplication. Because deduplication technology is proprietary between vendors, and the process of deduplication requires significant processing resources for both backup and restore operations, any additional data reduction benefit must be weighed against the added cost of the deduplication technology and the business impact on slower recovery from dedupe files systems and the practicality of dedupe for supporting applications which use backup data.

What is deduplication?

Data deduplication, often called dedupe, is a method for reducing what is written to data storage media by eliminating duplicate data. Deduplication techniques may be seen in most backup software, operating systems, storage arrays and dedicated deduplication appliances. There is no standard for deduplication between vendors; each offering is unique and incompatible with that of another vendor.

Deduplication operates first by breaking up data to be written to storage into chunks or blocks. Each block of data is compared against what has already been written to determine whether it is unique or is a duplicate of what has already been stored across the entire file system or logical volume. If the block is unique, it is written to disk, if it is an exact match for what has been previously stored, the deduplication file system records a pointer to what has been previously stored, thereby eliminating a redundancy.

Much of the value of the deduplication system is the process for ensuring change does not result in data loss. For example, over time files are added and deleted. When a deletion happens, a dedupe file system cannot just erase the file, the system must first determine whether any other files are using any block associated with the file marked for deletion and adjust the pointers to correctly reflect which blocks are valid and which may be deleted during the garbage collection process.

Both compression and deduplication are often used together to reduce the size of stored data and, in many implementations, the algorithms used have some similarities. Deduplication is normally applied across an entire volume or even entire storage system whereas compression is applied at the block or file chunk level. The extra reduction comes at the cost of more required resources, including additional memory, CPU, and disk I/O, which is why many systems that provide global dedupe are expensive, purpose-built physical devices, or virtual appliances carving out significant resources from the virtual environment itself.

What is deduplication designed to solve?

Deduplication reduces storage amounts as its primary objective. It was designed to allow for a cost-effective way to keep more backup data on disk for the faster recovery of individual files than was possible using tape. Random access to individual files on disk offered a vastly superior time to recover individual files. Individual tapes needed to be mounted and fast-forwarded to the correct position from where that same individual file may be recovered. Deduplication reduced the time to recover from minutes and hours to seconds and minutes.

Deduplication made more efficient the movement of backup data across networks. When the deduplication process of data reduction occurs at the source of the data, significantly less infrastructure was required to support the movement of data associated with backup processes. Similarly, moving deduplicated backup data across a wide-area network to a disaster recovery facility proved more efficient than moving the same un-deduplicated backup image.

Deduplication solved the challenge of large data volume storage requirements typically associated with database backups. Databases do not benefit from traditional file-level incremental backup techniques because they “touch” every file, every single day. Since incremental backups capture the changed files, each incremental backup of a database is essentially a full backup, requiring lots of storage space. Because deduplication operates at the sub-file level, only the sub-file changes need to be stored, even though the full database would need to be copied to the dedupe system. Database space savings and recovery performance improvements over tape made dedupe system expenditures easier to justify.

Deduplication was intended to keep more online backups, deal with bandwidth issues and support the idea of applications which always need full backups. It was designed in a time when storage was costly and backup operations required the storage of massively redundant data.

How is dedupe often implemented?

Deduplication systems consist of two primary logical components, the deduplication engine and the deduplication storage platform.

The storage platform may be deployed in both software and hardware appliance form-factors. Regardless of the form-factor, deduplication storage systems require significant resources for its underlying CPU, memory and storage that hold the deduplication metadata.

The deduplication engine functions to reduce and compress the backup data. Depending on the vendor, the dedupe engine may be deployed at the backup data's source (client), integrated into the backup server, or as a dedicated dedupe appliance. The point at which deduplication is processed will impact the performance of the device on which it is operating and the other processes also operating on that particular machine. While pushing the dedupe engine out to the source (backup client) device has the greatest impact on production applications, it also distributes and parallelizes the dedupe function and minimizes data movement across the internal network. Client-side deduplication, however, does not typically push the restore rehydration process to the client; meaning the process takes place at the dedupe storage system point. The implication being, because rehydration takes significant resources on the dedupe system, the number of simultaneous recoveries will be limited. Also, because the recovery process takes place at the dedupe storage server, fully rehydrated and uncompressed data will need to move across the network and potentially the wide-area network. In the occurrence of multiple simultaneous outages or a true disaster, planning and prioritization of the recovery process will be required and recovery time objectives (RTO) understood. A dedupe system with large capacity may not be the best choice when the recovery of many systems is required as the system itself may become the recovery bottleneck affecting many systems. Some people choose to sacrifice some global deduplication efficiency for better backup and recovery performance using multiple smaller dedupe systems or multiple dedupe pools on the single system to achieve faster parallel ingest and rehydration. The trade-offs associated with the point at which deduplication takes place must be understood and balanced against the business impact of backup and recovery performance.

Another method used to solve the problem of deduplication ingest performance is running the deduplication process as a post-backup process. Post-process deduplication involves first staging the backup job to a non-deduplication storage partition followed by moving the data to dedupe partition in off hours. Post-process deduplication allows for fast backup without requiring the level of CPU processing resources needed to stream and deduplicate data at top speed. The staging area used by post-process deduplication, if sized right, may be used for fast recovery of non-deduplicated data and other advanced applications using backup data as their source.

Most modern operating systems, storage arrays, backup software and hardware offer some form of data dedupe. Microsoft Windows 2016, as an example, offers deduplication in their Resilient File System (ReFS). ReFS scales to data volumes measured in zettabytes (dedupe is currently limited to 64 TB) allowing for generic Intel-based hardware and back-end storage.

Is this too good to be true?

This seems to be all good, less data on my storage so my storage can last longer. If deduplication is so good at eliminating redundancy and reducing data amounts, then why is it not used everywhere? Is there another side to know about?

It is important to understand that there are use-cases where dedupe is not optimal. For example, deduplicating previously deduped, compressed, or encrypted data may result in data sizes actually growing. Deduplication can also introduce performance issues due to the processing required for the initial ingest and deduplication process, as well as the retrieval and re-hydration process. Dedupe everywhere is not realistic, or necessary, due to cost, complexity and performance degradation. In addition, it makes some of the amazing Veeam technologies work much slower or not at all.

Backup restorations from a deduplicated source will be much slower than those coming from a traditional file system, often rendering techniques like instant recovery (running the restored machine directly from backup) unusable and file and application restore unusable due to high I/O latency. But standard VM restore workloads can be affected as well, because the maximum throughput can be only achieved by parallelization. While dedup storages often support a maximum throughput of 800–950MB/s overall for parallel streams, a single stream works on just 100–200MB/s over time. Ergo while you do not experience any issues while backing up to the storage (high parallel stream count), a single huge VM disk restore can take longer than expected.

The poorer performance not only impacts data restoration but also more frequently done operational tasks of duplication of the backup data to alternate and off-site disk storage to tape. Specifically, the output to tape will lead to a high rehydration load on the deduplication system, bringing those often to their maximum. This can be the bottleneck for any tape processing and could lead to stop-rewind-start operation on the tape side as not enough data is rehydrated in parallel to keep all tape drives in stream mode. Alternatives for these tape scenarios are the usage of software-based deduplication and compression that works on backup file level and does not have to be rehydrated when transported to tape (Veeam's own methods).

Cost is another factor which can make deduplication less than ideal. While deduplication is increasingly built into many software and hardware platforms as a no-cost feature, there is a cost associated with the CPU and memory resources needed for performance. Dedicated dedupe platforms are sized to support deduplication and a premium is paid for the proprietary functionality. In many cases this premium can make the disk cost-saving because of reduced data amounts, making deduplication insufficient and bringing the costs near to a standard storage system built for massive storage scale outs. At the same time, these standard storages do not have the deduplication random read performance penalty and can work multiple times faster with these workloads like Instant VM Recovery®, Guest File restore and application restore.

Modern file systems like ReFS and storages specifically built for massive storage scale out make it very easy to manage huge capacities without dealing with hundreds of LUNs. This is a domain which deduplication storages used in the past as one of the main selling points.

Because the dedupe techniques of every vendor vary slightly, lock-in to any particular vendor can restrict future flexibility, increasing the difficulty of replacing one dedupe hardware platform with another.

Sometimes vendors force their customer to use global deduplication systems, as they offer a real price benefit, leaving their customers a performance penalty by design. In many cases it comes with a per-TB capacity backend pricing after their global software deduplication. This forces the usage of their global deduplication (a performance penalty because of costs) and makes hardware deduplication storage systems unreasonably expensive. On the other hand, this software needs huge investment into the database backend as the deduplication metadata is stored there.

Finally, there's reliability aspects of an "all eggs in one basket" design. By fully eliminating any and all redundancy, you are risking your entire archive should you run into a data corruption issue. Specifically, this is true if systems use global deduplication and just replicate the storage as secondary "backup."

New uses for backup data

There is value in the use backup data for applications beyond traditional restoration operations. Vendors like Veeam have developed a set of applications which make use of the backup data to: Run production workloads, facilitate faster and better testing and development, automate disaster recovery (DR) testing, populate data warehouses with fresh data without affecting production operations and more. Effective use of backup data for new applications requires an architecture closer to production performance characteristics than traditional backup, however. The strategy for back-end storage requires a performance tier for the most recently backed up data, intentional insertion of physical gaps between backup copies and the use of cost-effective storage strategies for longer-term and off-site retention requirements.

Agnosticism with regard to storage hardware is extremely important as it allows for great flexibility in designing a backup architecture which meets the performance and Availability requirements of the business. For example, Veeam automates backup storage policies which, depending on the need, may use: Highly-performant flash storage for the performance tier, a tier of fast disk for the next level of performance, low-cost commodity storage, deduplication appliances, tape libraries and cloud-based object storage. Veeam may also automate the replication of data and backup data to off-site locations. Combinations of storage types may be created and mixed to form storage pools segmented by use-case and performance requirements. This offers all the best options available in the industry for mixing and matching storage to meet cost, performance and retention goals.

Veeam's Instant VM Recovery® allows the use of backup storage for running production workloads in case of disaster recovery situations. In the event of a primary VM failure or corruption of some sort, Veeam starts up the replacement production VM using the backup data as the source. Veeam supports hundreds of simultaneous production instances of Instant VM Recovery running off the backup data, which obviously requires fast back-end storage to meet the expectations of the business. Thus, a high-performance tier is recommended containing the most recently backed up data. Deduplication is not optimal for running production workloads, but is useful for older backup data with longer-term retention.

Veeam SureBackup® automates the backup recoverability testing of backups. SureBackup validates and reports on the ability to recover from backup by starting up the VM from backup data in the isolated lab, pinging the server, testing the applications and even allowing for custom scripts which for example may issue SQL queries to validate database integrity. The number of SureBackup operations which may run concurrently every evening is also dependent solely on the back-end infrastructure performance. The more highly performant the infrastructure, the more SureBackup tests may be run.

The Veeam On-Demand Sandbox™ allows for one or more VMs operating directly from backup in an isolated environment to troubleshoot, test and train on a working copy of the production environment, all without impacting business operations. The On-Demand Sandbox has proven useful for testing and development which allows for testing against the use of the most recent production applications and data. Similarly, the On-Demand Sandbox allows for the startup of a database server from a backup image, as an example, from which SQL queries may be issued to extract daily changes to populate data warehouses and decision-support systems, without impacting a production copy of the application. A fast tier of storage for current backup data is likewise essential to good performance for On-Demand Sandbox functions.

Which data reduction technique is optimal?

Optimizing the cost/benefit of various backup techniques and storage technologies means we will need to understand business requirements associated with the use of backup data and retention requirements. There are trade-offs associated with each of the major techniques; a hybrid approach may be the best choice for most organizations and will be discussed below. Here are some assumptions for an architectural design for "optimal performance," which will of course vary depending on each organization's requirements.

- First, we will assume we are following a typical retention strategy of seven (7) daily backups, three (3) weekly backups and three (3) monthly backups. Furthermore, our strategy will introduce intentional gaps in the physicality of the backup data.
- Second, we will follow a 3-2-1 strategy, meaning three copies of the data on at least two types of storage with one copy off-site. Two types mean Vendor 1 and Vendor 2 so that if a firmware upgrade, for example, disables one storage vendor you still can pull restores from the other storage.
- Third, we will assume the organization will take advantage of Veeam applications which use backup data and require a higher level of performance than typical backup recovery.
- Fourth, cost is a factor in the overall architecture and capital and operational costs are to be minimized while achieving the first three assumptions.

Achievement of the objectives will require a hybrid storage approach and the use of additional available Veeam data reduction techniques.

The Daily Backup Storage Strategy

The seven daily backups should live on highly-performant storage as this will be the tier from which the majority of application demands are made, including the typical restores.

Why is this a cost-effective use of storage? Veeam recommends the use of an incremental strategy which captures daily block-level changes to minimize the amount of data storage at this performance tier. Storage requirements for block-level incremental backups are comparable in size to deduplicated storage while offering a significant performance advantage. This is because deduplication works by capturing and storing only those blocks of data which have changed since the last backup and eliminating redundancy which may exist. Since block-level incremental backup also captures only the changed blocks of data since the last backup, storage requirements are comparable.

Sizing assumptions for dedupe; assume a 1 TB file is backed up. Using deduplication and compression, this 1 TB file is typically reduced to one-third of its original size, or 333 GB. Each file-level incremental backup we will assume is 10-percent of the size of the original file, or 100 GB. Using the same rule of thirds, each subsequently deduped incremental backup will require 33 GB of additional storage capacity. Over the seven days of retention, deduplication will require $(333 \text{ GB} + (33 \text{ GB} \times 6 \text{ days})) = 531 \text{ GB}$ for seven days of retention.

Sizing assumptions for block-level incremental backups. Assume the same 1 TB file is backed up, compression will reduce this file to 50-percent of its original size, or 500 GB. Each block-level incremental backup operates at a sub-file level, recording only the blocks of data which change from one day to the next. Assuming the same 10-percent daily file change rate or 100 GB, we will assume 50-percent of the blocks contained within the file have changed, or 50 GB per day compressed to 50-percent of its size, or 25 GB per day. Over the seven-days of retention, block-level incremental backups require $(500 \text{ GB} + (25 \text{ GB} \times 6 \text{ days})) = 650 \text{ GB}$.

Actual results for each technique may vary significantly based on the data being backed up and the actual change rate. Reading data from a deduplicated storage pool is typically half the performance of its write speed. Furthermore, the number of simultaneous read-streams coming from dedupe storage quickly reduces the overall performance due to the calculations needed to rehydrate the deduplicated backup data; CPU and memory become the performance bottleneck. A non-deduplicated storage pool will not face the CPU / memory performance bottleneck and will typically be many times faster in read performance. When the business needs require fast access to backup data, "global" deduplication is not the optimal platform.

Weekly Backup Storage Strategy

Following the 3-2-1 rule, the weekly storage will live on different physical storage than daily storage. The retention requirement at this tier maintains three full copies of weekly backup data. In other words, the backup data must reflect the state of the data across each of three weeks. How we speak of this is important because dedupe platforms do not actually store three full copies of data, they store one master copy with additional blocks and pointers which reflect the full image for each of the other two weeks. The deduplication algorithms effectively minimize the amount of storage required to reflect each of those weekly backup images. Conventional backup techniques, which retain three full weekly backups, have significant redundancy, a 1 TB file would require approximately 3 TB of storage, or 1.5 TB using compression. Whereas, a 1 TB file stored in dedupe would typically require half of that, depending on the amount of change from one week to the next.

The block-level incremental technique may be used for representation of the three weekly full backups as well. Using the same math from the daily backups, a 1 TB file will compress down to 500 GB for the first full backup. But rather than creating two additional full backups of the data, we can create twenty (20) additional incremental backups from which to recreate any backup over the same three-week period. Extrapolating the math over twenty-one days of retention (three weeks), block-level incremental backups require $(500 \text{ GB} + (25 \text{ GB} \times 20 \text{ days})) = 1 \text{ TB}$. In other words, the three weekly backup images require approximately the same amount of storage space as the original data amount. While requiring 25- to 50-percent more space than with typical deduplication, the cost per terabyte of the storage platforms will be close due to raw storage being significantly cheaper than purpose-built deduplication appliances.

Monthly Backup Storage Strategy

The location for the monthly backups should normally be off-site, when following the 3-2-1 rule. The type of storage will be dictated by the off-site requirements. If the data is not expected to be retrieved for anything but the worst disasters, replication of the backup data to a cloud service provider offering the lowest cost will often suffice for smaller customers. For enterprises, backup to tape with off-site vaulting offers a good air-gap strategy to protect company data from possible on-line manipulation by hackers or a malicious insider.

DR from backup data will require more work to determine the optimal storage type. Optimal use of wide-area networks to replicate backup data to an off-site location will require minimizing data amounts using WAN acceleration and/or deduplication techniques. Deduplication appliances are very good at minimizing data movement but typically require identical hardware at the DR site. Veeam replication also offers built-in WAN acceleration to minimize the amount of data which must be moved. Keeping three monthly full backups in dedupe storage will be comparable in cost to using change-block incremental techniques. However, when the requirement exists for retention exceeding three months, the storage advantage shifts to deduplication platforms, which are best suited for long-term retention.

Conclusion

Veeam supports and encourages the use of the industry's leading hardware-based dedupe appliances when the use-case makes sense. Veeam also integrates deduplication as a standard part of its core backup engine. In addition to dedupe, Veeam uses other advanced data reduction technologies like BitLooker™ to create space efficient images of VMs, specifically designed to support fast backups and fast restores. New backups typically follow a block-level incremental forever strategy with deleted data block zeroed out, which means there will be far less data to dedupe and write to disk. Because the use of these advanced techniques reduces data comparably to that of deduplication, the economic cases for deduplication as the Veeam backup target are fewer and typically limited to long-term retention.

The business requirement to use backup data for production operations, development and test usage, backup and replica integrity verification and other uses requires storage faster than typical deduplication and tape platforms can offer. Where the business uses backup data for advanced applications, faster storage is optimal. When long-term retention and data archival is more important, deduplication makes more economic sense.

Appendix

Links

Wikipedia definition of Data Deduplication: https://en.wikipedia.org/wiki/Data_deduplication

Veeam Support: <https://www.veeam.com/support.html>

Veeam University: <https://www.veeam.com/university.html>

Veeam and ReFS working together: <https://notesfrommwhite.net/2017/03/17/veeam-and-refs-how-to-make-things-work/>

Resource Library has videos on different features: <https://www.veeam.com/resource-library.html>

About Veeam Software

Veeam® recognizes the new challenges companies across the globe face in enabling the Always-On Business™, a business that must operate 24.7.365. To address this, Veeam has pioneered a new market of Availability for the Always-On Enterprise™ by helping organizations meet recovery time and point objectives (RTPO™) of < 15 minutes for all applications and data, through a fundamentally new kind of solution that delivers high-speed recovery, data loss avoidance, verified protection, leveraged data and complete visibility. [Veeam Availability Suite™](#), which includes [Veeam Backup & Replication™](#), leverages virtualization, storage, and cloud technologies that enable the modern data center to help organizations save time, mitigate risks, and dramatically reduce capital and operational costs.

Founded in 2006, Veeam currently has 47,000 ProPartners and more than 242,000 customers worldwide. Veeam's global headquarters are located in Baar, Switzerland, and the company has offices throughout the world. To learn more, visit <http://www.veeam.com>.

AVAILABILITY for the Always-On Enterprise™

VEEAM

A night-time photograph of the Chicago skyline, featuring the Willis Tower and other illuminated skyscrapers, with their lights reflecting on the water in the foreground.

Veeam makes the Fortune 500 Available. 24.7.365

To enable its **Digital Transformation**, 70% of the Fortune 500 rely on Veeam to ensure Availability of all data and applications. **24.7.365**