WAN Acceleration

What is the Veeam WAN accelerator and how does it work?

Veeam Data Loss Avoidance

Paul Szelesi
Veeam Solutions Architect Manager EMEA

Contents

Executive summary

Businesses today have high demands on IT departments, as well as high expectations and high penalties if they fail. Software can meet those demands through design, exceeding the requirements of most organizations. However, to do so, it has to be correctly incorporated into the infrastructure like a right-fitting jigsaw piece. To meet those demands, we have to ensure the overall design meets the needs of the business while providing the opportunity to grow. This should be accomplished by pushing every component to the maximum without excessive expenditure, meeting the expectations and delivering to the organization. Part of the solution should be the protection of valuable data and retaining Availability.

Veeam® is a global leader in data Availability and protection. The flagship product is Veeam Backup & Replication™, a product which offers unrivalled Availability for all business types.

This paper provides practical advice to help administrators, architects and IT professionals understand how they can maximize the data flow between locations in their infrastructure when using Veeam Backup & Replication WAN accelerator tools, so that they can confidently deploy their Veeam services.

Target audience

This paper is written for the infrastructure architect or engineer who is looking to understand how Veeam WAN acceleration functions. Veeam WAN accelerator is not a fit for all technology. By understanding how to implement it with foresight of its functions is to get the best from it. This way it can be a positive influence to any Veeam design.

This paper is not a how-to-configure Veeam WAN acceleration document. Its primary goal is how WAN acceleration works, design principles, technology and will it work for you.

Introduction

Veeam's strategy is underpinned with five important capabilities that guide and help our customers achieve success in data Availability and protection.

- High-Speed Recovery
- Data Loss Avoidance
- Verified Recoverability
- Leveraged Data
- Complete Visibility

Veeam WAN accelerator is from the Veeam data loss avoidance capability and is included in both the Enterprise and Enterprise Plus editions. The Enterprise edition is limited to using the WAN accelerator for cloud connectivity optimization. The Enterprise Plus edition enables additional features for replication and backup copies across sites. WAN acceleration can dramatically reduce the bandwidth required for transferring backups and replicas from one location to another, as well as reduce the amount of redundant data that clogs up links. This eliminates the need to acquire additional network bandwidth or deploy a general-purpose WAN accelerator appliance in the right circumstances for replication and backup copies.

Moving data quickly from one location to another requires resources and bandwidth, which can be costly and complex. Other factors enforce limits on both, such as budget or infrastructure, which makes data movement a serious discussion. While protecting data with the use of backup software, the window of opportunity to move the data is critical in balancing a functional and profitable business while ensuring the right amount of data is moved at the right time.

Many businesses face the challenge of growing data, while the ability to protect the data reduces. This is generally caused by advancements in technologies and desires of the consumer. The need to operate 24.7.365 is normal in today's business world, and if you cannot keep up the supply, customers will look elsewhere. It's not just about the price to buy a service or item any more, but the ability to process and deliver in a timely manner.

While businesses struggle to keep up with the trends and requirements behind the scenes, we have a duty to protect the data that is collected or the transactions that are completed just in case anything goes wrong. Compliance and regulations are creeping into every corner of the markets and have to be adhered to in order to stay in business.

The Veeam portfolio can help with data protection, and WAN acceleration is a tool within the software stack that can reduce the administration and complexity, while providing optimizations for data movement in relevant circumstances.

The Veeam WAN accelerator

***Accelerator:** A person or thing that causes something to happen or develop more quickly (Oxford English Dictionary)*

Within Veeam Backup & Replication are many tools, which provide Availability, enhance data protection, and reduce complexity and administration. In the Enterprise and Enterprise Plus editions, one tool that stands out is the WAN accelerator. The accelerator does not speed up the data, as the word accelerator suggests in our common use of the term. Instead, it optimizes the data by only transporting new relevant traffic. Data flow can be to many locations in an infrastructure. Backup data can be accelerated to another storage location, another type of media or a replica to another site, maybe for DR. The flow of the data is crucial, especially when bandwidth is not optimum or locations have a limited option on the medium used to move data.

The job of an accelerator is to enhance the process or speed up the process in some way to make it more efficient. The Veeam WAN accelerator takes incoming data from the source and processes the contents, removing the unrequired duplicated data and transferring only new or previously unsent data. In this way, we can reduce duplication and wasted packet traffic to enhance only the relevant data flow. By removing unneeded data, the remaining information can be sent more efficiently to the target location.

Where should it be used

Before discussing how the Veeam WAN accelerator works and how to size for it, we can discuss if and where a WAN accelerator is best used and where it is not. Many times it has been implemented in infrastructures where the actual WAN accelerator ends up being the bottleneck. It's not suitable for all infrastructures, and if used incorrectly, can slow down the data flow rather than enhance it.

WAN acceleration is designed to optimize high latency and low bandwidth links between locations. There is a natural overhead and resource requirement when this is in operation, and there will come a breaking point in regard to deciding if WAN acceleration works efficiently for your infrastructure. There are a number of ways to determine this breaking point, based around speed and your available resources.

WAN acceleration can be a one-to-one or many-to-one connection. The first thing you should consider is the bandwidth available between the locations to see if the cost of optimizing your traffic is outweighed by the speed of your link. This is very important in deciding how useful it will be in your infrastructure.

Bandwidth considerations for WAN accelerations

- **Link less than 3 Mbps**
  The WAN accelerator is likely saturated at this point. WAN acceleration will work, but other storage efficiency methods should be employed here, such as more data reduction (dedupe or compression).
- **Link more than 3 Mb/s and less than 40 Mb/s**
  WAN acceleration is a good fit here. Processing rate and data flow will see the biggest benefit from this bandwidth spread.
- **Link more than 50 Mb/s**
  The WAN accelerator will not be fully utilized and other forms of data movement, such as direct mode, will be more effective.

These numbers are to be considered as a baseline. "Your mileage may vary." The performance of the underlying storage can greatly impact the performance of the WAN accelerator function.

There are no significant performance differences in using spinning disk drives as storage for the target WAN accelerator rather than flash storage. However, when multiple WAN accelerators are connected to a single target WAN accelerator (many-to-one deployment), it is recommended to use flash or equivalent storage for the cache, as the I/O is now the sum of all the different sources. What is also relevant and should be taken into consideration is performance of the source cache, which could affect the overall processing performance. The disks chosen should be balanced to achieve a satisfactory level of performance.

When thinking about the repository used at the target WAN accelerator end of the job, data may be taken from this repository at the target WAN accelerator from the first full backup or if the data is not found in the cache but is known to exist in a previous transfer. If slow disks are used, it can have an impact on the speed of the completion of the job and overall processing rate. A suitable level of performance disks should be used at the repository for this reason, or it can slow down the job.

Other factors are also present, such as if this link is going to have bi-directional data flow when using the WAN accelerators and how many jobs will be using the link at the same time.

Measure your overall saturation of the link prior to using WAN acceleration to ensure that it meets your needs.

The source WAN accelerator can also consume CPU, memory and disk resources while processing the workload. The overhead of this can be not only intensive, but the time for processing can be much slower than the direct mode used on bandwidth over 40 Mbps. Trial and measure should be used around this figure to get the right type of data transfer (WAN acceleration or direct mode).

Testing is always recommended to ensure you can achieve your desired results.

How Veeam WAN acceleration works

WAN acceleration works by combining multiple technologies, such as network compression, multi-threading, dynamic TCP window size, variable block size deduplication and global caching. WAN acceleration provides sufficient capability when the network bandwidth is low or dramatically reduced when performing backup copy and replication jobs.

The WAN accelerator uses its own digests (discussed later) based on the hashes of the data blocks inside a VM disk. It reads data from the backup files, re-hydrating them on the fly (or it reads directly from the source VM in case of replication). The WAN accelerator component will process those data blocks with much more efficient data deduplication and compression algorithms. Because of this, the WAN accelerator can consume significant amounts of CPU and RAM resources, and an ample amount of both should be provided for its consumption.

We can see from the diagram below how WAN acceleration traffic flows, and the next section will go into the detail on what each accomplishes.

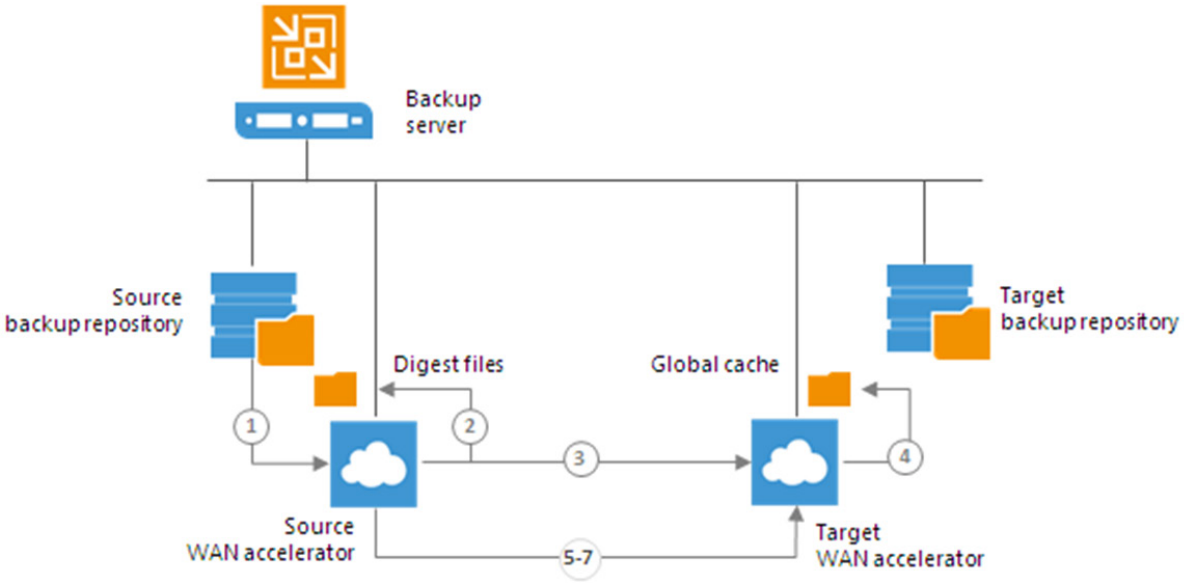

1. [For Backup Copy Job] Veeam Backup & Replication uncompresses the backup file to analyze its content.
2. The Veeam WAN accelerator service on the source WAN accelerator analyzes data blocks of the file to be transferred and creates a file with digests for these data blocks. The created file with digests is stored to the VeeamWAN folder on the source WAN accelerator.
3. Veeam Backup & Replication compresses the backup file data (for backup copy) or VM data (for replication) and copies it to the target side.
   At this point, Veeam Backup & Replication can perform deduplication within the VM itself — that is, deduplicate identical data blocks in every VM disk.
4. During the data transfer process, the Veeam WAN accelerator service on the target WAN accelerator populates the global cache storage with data blocks from the copied file.
5. During the next job cycle, the Veeam WAN accelerator service on the source WAN accelerator analyzes data blocks in the file that must be transferred this time and creates digests for these data blocks.
6. The Veeam WAN accelerator service compares the created digests with the digests that have been previously stored to the VeeamWAN folder on the source WAN accelerator. If duplicate data blocks are found, the actual data block in the backup file is not copied over WAN. Instead, it is taken from the global cache and written to the restore point in the backup copy folder or on the target data volume.
7. Additionally, Veeam Backup & Replication analyzes restore points that have been previously copied to the target side.
   If duplicates are found, Veeam Backup & Replication does not copy such blocks over WAN but takes them from the previously copied restore points.

WAN accelerator job flow

For a job, the flow is as follows:

1. Data is uncompressed and the contents are analyzed as it is transferred through the accelerator process.
2. The data is divided into chunks of variable length between 1 Kb and 11 Kb (4 KB average size). Each chunk is then hashed using an MD5 128-bit algorithm to create a digest file, which is stored in the source digest cache. Once a full pass of a disk is complete, we have a fingerprint of that particular disk. Veeam stores two fingerprints in the source cache for each disk: the previous job run and the current job run. There is no need to keep any more. The two files are then used as comparisons for the WAN accelerator to define how the data will be treated.
3. The Veeam WAN accelerator will then run a data redundancy algorithm on the data to establish if it has previously been sent over the wire to the target WAN accelerator. If an identical block is found, it will be marked as a duplicate.
4. Veeam Backup & Replication then compresses the backup file data (Backup C opy Job VM) or VM (replication) and starts to copy the data to the target WAN accelerator. Deduplication will also take place at this point, removing unnecessary data from the file / VM (content of the VM).
5. As the data is being transferred, new data will start to populate the global cache on the target side from the source that has not previously been stored. The global cache consists of OS-specific data as well as Exchange data.
6. During the Veeam Backup & Replication job, the source WAN accelerator creates new digests stored on the source. The new digests become the current fingerprint and are stored in the source digest cache folder. The WAN accelerator service checks the current digest against the last job run digests. If the digests are found to be duplicates, then the WAN accelerator will exclude the data block from the transfer (not the metadata). The data is then written to the target from the global cache or the data repository of the duplicate.

If the target WAN accelerator is a target for a many-to-one relationship (many source accelerators connect to the same target accelerator), the global cache cluster or local repository may contain required data blocks for a particular VM type. In this case at the start of the job, the target WAN accelerator will copy the files into the relevant global cache cluster from the relevant location as part of the job. You can check if this has occurred in the job history.

The sources used for data deduplication are as follows:

- The WAN accelerator will use the data blocks from inside a VM backup. If duplicates are found, they will be deduplicated.
- When using WAN acceleration, Veeam Backup & Replication will examine restore points from the target backup located in the target repository. If a duplicate is found, the new data is not transferred during the process.
- At the target location, a global cache is created. This stores all relevant data (OS, Exchange and root drive data) that is sent over the wire to the target. It is examined by Veeam Backup & Replication, and any data found to be duplicated will not be transferred from the source data or sent over the wire.

In the first two in the list above, WAN acceleration processes data from one VM disk at a time. The third source (global cache) will deduplicate based off the data stored from any cached data of the VM disks previously sent over that source WAN accelerator.

First about Veeam replication

Veeam Backup & Replication is built for virtual environments. It operates at the virtualization layer and uses an image-based approach for VM replication. It does not install agent software inside the VM guest OS to retrieve VM data, rather it leverages VMware vSphere snapshot capabilities. When you replicate a VM, Veeam Backup & Replication requests VMware vSphere to create a VM snapshot. The VM snapshot can be thought of as a cohesive point-in-time copy of a VM, including its configuration, OS, applications, associated data, system state and so on. Veeam Backup & Replication uses this point-in-time copy as a source of data for replication.

In many respects, replication works similarly to forward incremental backup. During the first replication cycle, Veeam Backup & Replication copies data of the original VM running on the source host and creates its full replica on the target host. Unlike backup files, replica virtual disks are stored uncompressed in their native format. All subsequent replication cycles are incremental. Veeam Backup & Replication copies only those data blocks that have changed since the last replication job session to keep track of changed data blocks.

Veeam Backup & Replication lets you perform on-site replication for High Availability (HA) scenarios and remote (off-site) replication for disaster recovery (DR) scenarios. To facilitate replication over the WAN or slow connections, Veeam Backup & Replication optimizes traffic transmission. It filters out unnecessary data blocks, such as duplicate data blocks, zero data blocks, blocks of swap files and blocks of excluded VM guest OS files, and compresses replica traffic. Veeam Backup & Replication also allows you to use WAN accelerators and apply network throttling rules to prevent replication jobs from consuming the entire network bandwidth.

You can use WAN accelerators if you replicate VMs over a slow connection or over WAN.

In the replication process, WAN accelerators are responsible for global data caching and deduplication. To use WAN acceleration, you must deploy two WAN accelerators in the following way:

- The source WAN accelerator must be deployed in the source side, close to the backup proxy running the source Veeam Data Mover.
- The target WAN accelerator must be deployed in the target side, close to the backup proxy running the target Veeam Data Mover.

Disconnections

During replication, a VM snapshot is used on the source host. This becomes the working snapshot for the duration of the data transfer on the target. If a job is running and the connection between WAN accelerators fails for any reason, it will wait up to 30 minutes for a reconnection to continue that session. If a reconnect happens, it will continue where it left off with the transfer. If 30 minutes expires before reconnection, Veeam will still not fail a job. A new session can be established. At this point, the last used working snapshot will be merged into a new working snapshot when using replication or the same payload. If the source is a Backup Copy Job, the cycle will attempt to transfer all data since the last successful one. In replication, Veeam will keep two previous snapshots, merging any others if disconnections continue to happen for the duration of the cycle. This keeps snapshot chains short.

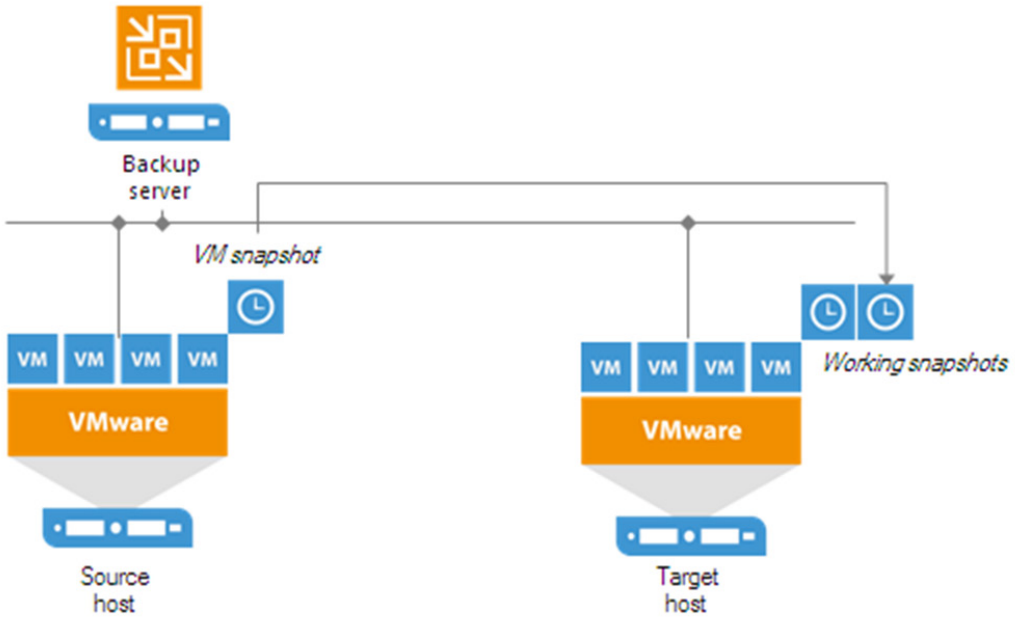

The end result is that only unique data blocks will be sent over the link. No corrupt or damaged data will be sent, and disconnections can be handled without restarting the job or having multiple failures. Data blocks that are available already on the target will not be sent.

When the WAN accelerator collects data to be transmitted to the target, it creates a 1 GB payload. This payload ensures that if a transfer disconnects, becomes corrupt or otherwise needs to be retransmitted, it will never be more than a 1 GB file.

Understanding how to size a Veeam WAN accelerator

The first thing to understand about how a WAN accelerator will impact your backup process is to look at the type of backup you have in effect in your schedule.

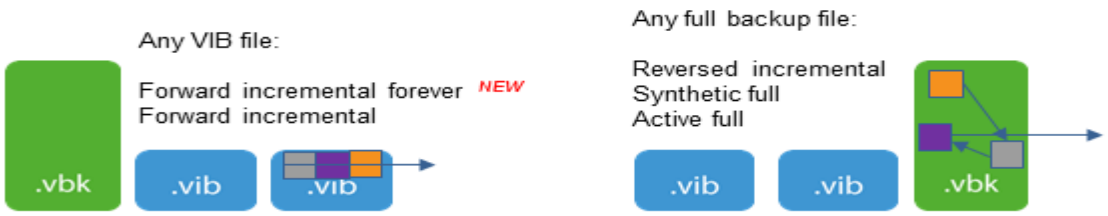

Reverse incremental, synthetic full and active full backup data is randomized due to its nature, which means it will increase source-side latency with disk reads. As a single workload, this may not be significant, however, when applied to multiple workloads, that latency will soon accumulate.

Relationships and WAN accelerators

There are two types of relationships for WAN accelerators, and there is a major difference in sizing when using one or the other:

- **One-to-one relationship**

One to one is the most common for small environments with a single remote office using a slow bandwidth high latency link. For this relationship, a single source accelerator is mapped to a single target accelerator. I/O requirements are usually not very heavy and 10 K disks are usually sufficient for this type of workload. This is also the easiest to size as the following section discusses.

- **Many-to-one relationship**

Many to one has one target accelerator, which services up to four source accelerators (best practice). This is commonly seen in organizations that utilize a hub and spoke remote office backup where all or some have slow bandwidth high latency connections. I/O loads can be significantly higher for these target accelerators, and faster speed disks or SSD for the global caches are advised to accommodate the performance demands.

Cache and digest sizing considerations

Sizing the WAN accelerator fingerprints and cache are crucial to running a successful and efficient job. Here we will discuss the sizing recommendations for both aspects. First, we must understand the data flow. Although we create WAN accelerators within Veeam Backup & Replication, we are not creating a dedicated source type or target type accelerator. The accelerator is a common object on creation, and a source for one job can be a target for another. In this way, the job design is important in sizing to avoid excessive disk purchasing on source accelerators when it is not required.

There are distinct differences between a source accelerator and target accelerator. Although they host the same files and folders, they are populated differently.

**Digest sizing and cache sizing are cumulative. You need to calculate both and add the sum together to get an accurate sizing. Digest sizing is not part of the cache sizing done in the WAN accelerator wizard in Veeam Backup & Replication.*

Fingerprint sizing

The fingerprints are stored on the source WAN accelerator, and each WAN accelerator can create fingerprint files when used as a source. The source will create the fingerprint based off data traveling through and then stored locally.

This is the normal operation, however, there are a few circumstances where this can work differently. For example, if a global cache is pre-populated but has not connected to a source, the target will create a digest from the files in the global cache. Once connected to a source accelerator, it will copy the digest to the source, saving time in fingerprinting the VM.

Another use is if a fingerprint becomes lost on a source accelerator, the target can recreate the fingerprints from the last restore point in the backup chain and send them back to the source as a way to speed up remediation of the host.

Sizing a digest on both is the same calculation. We can take 2% per 1 TB (20 GB) of the source data and size according to that output, including working space.

Source data = 10 TB
Digest space = 10 X 20 /100 =200 GB

The source accelerator

A source accelerator has a high I/O requirement. It has to read / write data and consume CPU cycles throughout the process, so it can be very resource hungry. It is recommended to use fast disks on the source to accommodate all the I/O loading it will have to meet while processing the data. The high CPU utilization is from decompressing data and recompressing / deduplicating the data once more as it sends it on, as well as MD5 generation usage .

For the source WAN accelerator, it should have, at minimum, four CPU cores and 8 Gb RAM alongside fast disks.

Sizing digest storage for a source accelerator should be a simple process as it holds only digests for the transferred data. The digest sizing should be a minimum of 2% or 20 GB per 1 TB of source data.

*If the source will be used also as a target accelerator, then target accelerator sizing should be applied on top of the source sizing.

The target accelerator

A source accelerator does not hold a global cache where a target accelerator does. The global cache is populated by the target and uses various data sources for its content. These include the disks that contain the operating system and objects in backup from a VM as it passes data. One of the most important, outside of the root drive, is Exchange data. Since Exchange 2010, single instance no longer exists. Exchange contains a vast number of small files in duplicate, triplicate or more. By storing this in the global cache, it is the first stop for the deduplication process and can have a big impact on the process.

When you create a WAN accelerator, it is best to pre-populate the global cache on the target accelerator from the backups in the repositories it will serve. This is a good starting point as the job will do this on the first run anyway. Doing it before the job runs will reduce time to complete the first run.

Let's say you have three repositories, and these have a mix of Windows Server 2012, Windows Server 2016 and Linux. No matter how many VMs they incorporate in the repositories, the three core system files will be copied into the global cache as a starting point. The below example shows that the count of OS types will be three, and we allocate at least 10 GB per OS type, so we will need 3 x 20 GB at minimum (We will use all the allocated space in the WAN accelerator settings proportionally.).

Windows Server 2012	20 GB
Windows Server 2012 R2	20 GB
CentOS	20 GB
Minimum global cache capacity	60 GB (3 X 20 GB)
Add at least 20 GB for working space	80 GB (60 GB + 20 GB)*

The starting point for sizing a target accelerator in this case would be 80 GB.

**Digest space is also required in addition to this.*

The above calculation is good for a one-to-one relationship where a single source is mapped to a target that services only that one source. In this case, the global cache is unique to that source data traffic. If we have multiple source accelerators to a single target (many to one), the new accelerators need a global cache for each source accelerator.

Let's take the calculation above and calculate based off four source accelerators rather than the one:

Windows Server 2012	20 GB
Windows Server 2016	20 GB
CentOS	20 GB
Minimum global cache capacity	60 GB (3 X 20 GB)
Add at least 20 GB for working space	80 GB (60 GB + 20 GB)

The starting point for sizing a single target accelerator in this case would be 80 GB*.

**Digest space is also required in addition to this.*

We would now make that available for the four source accelerators and create an additional cache for each source accelerator.

First accelerator cache:	
Windows Server 2012	20 GB
Windows Server 2016	20 GB
CentOS	20 GB
Minimum global cache capacity	60 GB (3 X 20 GB)

This is now our base cache taken from the OS calculation. We now need to add the minimum storage we calculated for each, which is usually the same as a single cache (in this case all four use the same size for simplicity).

1. Source 1:	60 GB
2. Source 2:	60 GB
3. Source 3:	60 GB
4. Source 4:	60 GB
5. Add our working space:	20 GB
Total space required as starting point:	**260 GB***

**Digest space is also required in addition to this for each source.*

In our example, a design with a one-to-one relationship and a design with a many-to-one relationship using four source accelerators have the same three operating system types / count for each, require much more space for the global cache and most probably need a much faster disk subsystem for I/O resources.

**If the target will be used also as a source accelerator, then source accelerator sizing should be applied on top of the target sizing.*

When using a many-to-one design, the global cache data can be copied to other global caches on the same target accelerator to save the retransfer of data over the WAN link when a new job starts. It can also populate other global caches in the same way. The technology design is meant to reduce data movement on the link to the minimum possible by using data in other global caches or even target repositories.

The WAN accelerator processes one VM disk at a time in a job. Sometimes this needs to be more as it can be a bottleneck in itself. This is the point where a many-to-one topology is very useful, as up to four source accelerators can be in use on any one target at any one time.

Calculating streams with WAN acceleration and how to saturate a link

By default, Veeam uses five TCP/IP connections to transfer data from the source accelerator to the target. Many streams can cause network traffic to be saturated if multiple jobs are running simultaneously. We can change the number of streams to optimize and coordinate with the link speed and latency monitored on the infrastructure. It is recommended to calculate a baseline and then monitor to see what effect it has on the link speed and job efficiency.

The stream number is set on the source accelerator during configuration or editing of the Backup Copy Job or replication job prior to the job running. When a flow of data starts from the source accelerator to the target, the source will set the stream count and communicate this to the target, which will work with the same setting. To this end, the target accelerator setting is not used (unless it is also a source).

In the data flow for a WAN acceleration job, each stream must form a connection at the target before any data can start to flow. If there is a high number of streams specified, each one must form a connection before it can move onto the next stage. A number such as 100 streams would take much longer to start than a job with five to 10 streams.

The default number of streams is five, which is a good optimal setting. For about 10 Mbps with around 100 ms of latency, this should achieve saturation at around 15 streams and 100 ms of latency.

Link (Mbit/s)	Latency (ms)	Packet loss (%)	Streams	Throughput (Mbps)
10	100	0	3	3.5
10	100	0	10	7.5
10	100	0	15	10
10	100	0	20	10

16 streams here will mean all the streams will not be capable of starting, which will have a negative effect on the performance.

The base calculation is usually around x1.5 per Mbps, but latency is a big factor here. If the setting is too high, then it will have a negative effect on the job. Too low can also have a negative effect on the workload. The key factor is link speed and latency and finding a happy medium for your workload.

Based off the default of five streams as a starting point, increase and monitor your change and then readjust up or down to get the best use of your link. Do not start too high. As mentioned before, this can have a dramatic effect on job times, and if you set 50 streams and only 30 can connect, the job will have a delay while the others time out. They then will try to connect in place of closing streams and only add to the latency of a job.

Conclusion

In this paper, I have discussed the Veeam WAN acceleration tool and how this helps with data loss avoidance, one of the Veeam five capabilities of Availability.

I have discussed the following in this paper:

- Is WAN acceleration the right solution for you?
- How the Veeam WAN accelerator works
- The Veeam WAN accelerator job flow
- What the source accelerator is and how to size for digests and cache
- What the target accelerator is and how to size for digests and cache
- One-to-one relationships
- Many-to-one relationships
- Disconnections during data transfer
- How to size for digest, why they are required, how they are created and the effect they have on data flow
- Data streams and how they affect data flow

Veeam WAN acceleration is a tool that is sometimes used in the wrong way. This leads to a misguided view of what its capabilities are. By sizing the cache and understanding how it works, it is much simpler to create well-formed and resourced source and target accelerators to make your data flow much more efficiently.

Data is very valuable to any business, making it work well in production is only half the story. When it becomes corrupt, disappears or is threatened, we need a reliable method to keep it available and protected. Veeam's five capabilities follow a strategy to accomplish this level of Availability and protection while retaining simplicity by reducing complexity or administration, allowing the business to focus on what's important.

About the Author

Paul Szelesi is the Solutions Architect Manager for Veeam in EMEA. He is based in the UK and is focused on technologies around VMware and Hyper-V.
He has more than 20 years of experience in the field working with storage and virtualisation products. He managed projects as small as single servers to large storage implementations in the enterprise space. His current focus is on Veeam product implementation and designing infrastructures for customers.

Specialties: VMware Virtualisation, Microsoft Virtualisation, Backup, EqualLogic, Dell Storage, HP Storage, NetApp

About Veeam Software

Veeam® recognizes the new challenges companies across the globe face in enabling *Intelligent Data Management for the Hyper-Available Enterprise*™, where business that must operate 24.7.365. To address this, Veeam has pioneered a new market of Availability for the 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 more than 294,000 customers worldwide, and adds an average of 4,000 new customers each month. 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.

Intelligent Data Management

fuels the ***Fortune 500***

VEEAM

Intelligent Data Management
for the Hyper-Available Enterprise