



Veeam Cloud Connect v9

**A reference architecture to learn,
design, implement and manage
your Veeam-powered cloud services**

Luca Dell'Oca

EMEA Cloud Architect, Veeam Software,
vExpert, VCAP-DCD, CISSP

AVAILABILITY
for the Always-On Enterprise™

Contents

Introduction	4
Audience	6
How this book is designed	7
Acknowledgements	8
Architecture	9
Veeam Backup & Replication server	9
Standalone Console	13
Enterprise Manager	15
Cloud Portal	19
2.5 Cloud Gateways	21
2.6 WAN Accelerators	23
Backup Repositories	25
Proxies	32
Network Extension Appliances	34
Additional components	39
Regular maintenance of the components	42
Reference design for backup services	44
Network diagram for backups	45
Security zones	46
Veeam Cloud Connect Backup deployment	49
Reference design for replication services	68
The virtualized environment	68
Target Veeam proxies	72
Networking	72
Veeam Cloud Connect Replication deployment	77

Cloud Connect operations	85
Initial setup	85
Create hardware plans	86
Customer creation and backup resources	90
Backup and backup copy jobs	93
Restore data from Cloud Connect backups	97
Assign replication resources	100
Replication jobs	106
Failover plans	112
Partial failover	115
Full Failover	120
Monitoring Cloud Connect with Veeam ONE	125
What you can do with Veeam ONE Free Edition	125
Monitoring, alarms and reporting	127
APPENDIX A: SSL Certificates generation	132
APPENDIX B: How encryption works	141
APPENDIX C: Customize your Veeam Cloud Connect Portal	145
APPENDIX D: Advanced Registry settings	148
About the Author	150
About Veeam Software	150

Introduction

In 2014, with the release of Veeam® Backup & Replication™ v8, Veeam introduced a new technology — named Veeam Cloud Connect — developed specifically for service providers to create and serve remote backup repositories. In 2016, Veeam Backup & Replication v9 added replication functionality to Veeam Cloud Connect.

Service providers who are part of the Veeam Cloud & Service Provider (VCSP) program can use Veeam Cloud Connect to offer customers Backup as a Service (BaaS) and Disaster Recovery as a Service (DRaaS). Every Veeam Backup & Replication v8 or v9 customer can then consume these services from their service provider of choice to send backups off site or to replicate (v9 users only) virtual machines (VMs).

With Veeam Cloud Connect, service providers can build their own Veeam-powered services offering, leveraging a technology built from the ground up to be multi-tenant and scalable.

Veeam Cloud Connect removes the main hurdles that such services required in the past by implementing different design concepts in its architecture.

No VPN tunnels

It is not easy to configure a VPN automatically, and it usually requires interaction between the service provider and the customer. Even when it is properly configured, it requires ongoing monitoring and management to guarantee it is always up and running. Otherwise, customers cannot consume the service offered via VPN.

With Veeam Cloud Connect, every connection happens directly over the internet using a single TCP/UDP port protected by SSL/TLS encryption. This is possible thanks to a new and dedicated Veeam component called a **cloud gateway**. A cloud gateway is responsible for the transfer of all backup and replication traffic over the single port connection. The connection uses the public internet and guarantees complete confidentiality of the data traversing the connection.




Figure 1-1 : General overview of Veeam Cloud Connect Backup

Multi-tenancy

The second design principle is **complete support for multi-tenancy**. Service providers create competitive services by sharing their resources among their customers. This allows for price reduction but cannot happen at the expense of security. Each tenant needs to be completely isolated from everyone else and in total control of his or her slice of the environment, just like in a dedicated environment. This is possible in Veeam Cloud Connect thanks to two different components: a cloud repository and a cloud host. For backups, service providers expose a **cloud repository** to customers. The cloud repository creates an abstraction layer over an existing backup repository so multiple customers can store backups inside the same shared repository with the same level of confidentiality they have with a dedicated repository.

For replicas, service providers can offer a **cloud host**. A cloud host is an abstracted view of the virtualized environment — either VMware vSphere or Microsoft Hyper-V — confined by a hardware plan that sets limits on CPU, memory, storage and networking that the customer can consume.

Another component, called a network extension appliance (NEA), stretches the network connections between the customer and the service provider sites and guarantees complete isolation of customer networks at the service provider site. Multi-tenancy is built into Veeam Cloud Connect and doesn't require additional components.

Security

A service exposed via public internet connection and shared between multiple tenants cannot ignore **security**. Veeam Cloud Connect offers different levels of security:

- **At source:** By leveraging the encryption capability first introduced in Veeam Backup & Replication v8, data is immediately encrypted by Veeam components on the customer side using industry standard AES-256-bit encryption, and encryption keys generated by the customer. Customers can choose encryption, but service providers can make it mandatory in the software.
- **In flight:** The connection between a tenant and the cloud gateway(s) is encrypted using SSL certificates (technically, it's TLS 1.2). This way, no man-in-the-middle attack will happen unnoticed, and even unencrypted data can traverse the public internet securely.
- **At rest:** Backup files are stored in an encrypted format at the service provider using customer keys. There is no possibility for the service provider to read the content of a customer's backups if the customer doesn't share the passwords with the provider. This guarantees complete confidentiality to customers and removes any liability issue from service providers.

NOTE: Due to the format of VM virtual disks, native encryption is not available for replicated VMs.

Even when encryption is enabled, it doesn't affect the data reduction ratios of Veeam's built-in WAN acceleration, as is the case with general-purpose WAN acceleration. In fact, Veeam specifically designed its purpose-built WAN acceleration to work in conjunction with encryption. For more information, see Appendix B.

Automation

Veeam Backup & Replication has always been renowned for its simple and powerful graphical interface. However, when service providers need to manage their environments at scale, no one can avoid looking into automation. Veeam Cloud Connect can be managed in all its aspects and automated with PowerShell or integrated in an existing customer portal, thanks to RESTful API.

Integration

The two abstracted components Veeam Cloud Connect creates appear in the customer's local Veeam installation like local resources. This is to guarantee a consistent user experience and greatly improves ease of use. Customers do not have to learn new tools or processes to consume the resources exposed via Veeam Cloud Connect; they can simply configure backup copy jobs (toward a cloud repository) or replication jobs (toward a cloud cost) as before. The ease of use and complete integration makes Veeam Cloud Connect an easy-to-onboard and easy-to-consume solution for service providers.

Any customer with a paid license of Veeam Backup & Replication v8 (for off-site backups) or v9 (for off-site backups and replicas) will have the client component of Veeam Cloud Connect available in the same user interface. Directly inside the Veeam backup console, customers can find a service provider who offers Veeam Cloud Connect and select the desired service provider by country and other parameters. Once the customer subscribes to the service the service provider offers, the customer will receive the parameters needed to activate the Veeam Cloud Connect service.

Veeam Backup & Replication installed at the customer site will connect via the cloud gateway(s) at the service provider authenticates the customer, and it will enumerate and expose the subscribed resources as if they were local.

Once the new resources are added to the console, customers can start using them just like regular local resources: Make them targets for any backup, backup copy or replica job, directly within the user interface.

Audience

This guide is intended for individuals who work at service providers and are responsible for the architecture, design, deployment and management of Veeam Cloud Connect. Readers of this book should be familiar with concepts pertaining to Veeam Backup & Replication and virtualized environments.

This guide does not replace the official Veeam User Guides. For any additional information about Veeam solutions, please refer to the relevant User Guides available at:

Veeam Help Center (<https://www.veeam.com/documentation-guides-datasheets.html>)

You may also engage other peers working with Veeam Cloud Connect and Veeam experts over the Veeam Forums at <https://forums.veeam.com/>. There is a private forum dedicated to Veeam partners registered with the Veeam Cloud & Service Provider (VCSP) program. To join, go to the User Control Panel, select Usergroups, select Cloud & Service Providers and then choose Join Selected.

How this book is designed

This book describes a possible architecture created by a service provider offering Backup Storage as a Service (BaaS) and Disaster Recovery as a Service (DRaaS) to its customers using Veeam Cloud Connect in Veeam Backup & Replication v9.

This design is not supposed to be the only right one or the best possible one — it is intended as a reference guide to help service providers to design and deploy their services with Veeam Cloud Connect. Service providers can also architect and deploy other possible designs following their specific requirements and business objectives.

This document refers to Veeam Cloud Connect as available in Veeam Backup & Replication v9 Update 1 (build 9.0.0.1491). Veeam recommends that service providers always use the latest version; the software allows for backward compatibility up to one major version, so v9 can receive backups from customers' installations using both v9 and v8 (because replica capabilities are new in v9, there is no backward compatibility). However, the installation on the service provider side **must** be at least at the same version as the connecting customer's, even in terms of minor updates. Because of this, Veeam invites service providers to install any update as soon as it is available.

This book is account of an Architect working at a service provider starting a journey in the creation of Veeam Cloud Connect services. The book first covers the different components available in Veeam Cloud Connect, characteristics of the components and how to use them.

Then, the book covers designing Backup as a Service and illustrates all the needed components, suggested configurations and tips to make the design a successful solution.

Later, the book outlines the design for a Disaster Recovery as a Service environment, using VMware vSphere as the virtualized platform. Again, all the different components and configurations will be explained in detail from the viewpoint of a service provider willing to deploy Veeam Cloud Connect.

In the Appendixes, you will find additional specific information about more technical topics.

Terms and abbreviations

Along the book, you will find different acronyms and abbreviations. Here you can find a map to better understand what they mean:

VM: Virtual machine

BaaS: Backup as a Service, as related to Veeam Cloud Connect backup services

DRaaS: Disaster Recovery as a Service, as related to Veeam Cloud Connect replication services

VBR: The Veeam Backup & Replication server

NEA: Network extension appliance, the virtual appliance deployed by Veeam to manage networking for DRaaS

Acknowledgements

This book is the result of many weeks of hard work, equally divided between writing the actual content and working in my Veeam Cloud Connect lab to create what I've described in the book. During this journey, I've become a service provider myself, to the point that the examples you see — even comprising IP addresses, dns names and SSL certificates — are all real.

However, many people worked hard alongside me to guarantee that this book would exist and be of the highest quality.

First, I'd like to give a huge thanks to my colleagues at Veeam — from Product Management and Quality Assurance to R&D — who took the time to review the different drafts of the book, give me feedback and verify every technical detail. Say thanks to all of them if you find this book accurate and useful.

Second, thank you to every other team in Veeam I've worked with: my peers in Technical Product Marketing, the Creative Team, our solutions architects and everyone else who contributed at any level to this book.

Another huge thanks goes to the guys at PhoenixNAP. They gave me the entire lab you see described in this book. They also helped me during the creation of the book by configuring and tuning both the infrastructure and the networking to my needs.

Finally, I'd like to thank my family. This book has taken a large part of my time for many, many weeks, and it's all because of them that I was able to focus on this project without worrying too much about other matters.

Luca

Architecture

Veeam Cloud Connect is a modular architecture made up of several components. Each component has a precise function, and together they work to provide the overall functionality.

Some of them can and should be deployed in multiple instances for High Availability and scalability purposes; in each section will clearly state if the described component can be deployed multiple times.

In addition to specific Veeam components, the architecture requires (or it can benefit from) additional general components, which will also be listed and described. For each component, you will see a description of its function, how you can monitor it, what level of protection it requires and how it can be protected over the network.

Veeam Backup & Replication server

As in every Veeam Backup & Replication deployment, this is the central component. Veeam Backup & Replication holds the main Veeam backup service, which manages all configurations and saves them into the back-end Microsoft SQL Server. You can manage using the standalone console, which is installed locally on the same Windows server or in a remote Windows machine. You can also use either PowerShell or RESTful API to manage Veeam Backup & Replication.

Veeam Backup & Replication requires a 64-bit Windows operating system.

NOTE: *Veeam requires every service provider to deploy a dedicated Veeam Backup & Replication deployment to Cloud Connect services, without mixing Cloud Connect with other Veeam-powered services.*

If you are only using Veeam Cloud Connect backup, Veeam Backup & Replication does not involve local activities on the service provider's hypervisor hosts. Instead, it only receives backups from customers that are already processed at the customer's sites. For this reason, the requirements for its installation are lower than usual: A simple VM with 2 vCPU and 4 GB of RAM will suffice to hold both the Veeam backup service and Microsoft SQL Server.

Regarding the SQL Server, the default Microsoft SQL Server Express can be enough unless the Veeam Cloud Connect infrastructure will host a very large amount of customers because activity logs can fill the maximum size of an Express database (10 GB). If this is the case, you should plan to use a regular SQL installation (Standard or Enterprise) either in the same machine or in a dedicated one.

However, if you are going to deploy Veeam Cloud Connect for DRaaS also, the Veeam Backup & Replication service is going to manage a proper virtualized environment with many virtual machines belonging to all the different hosted customers. In this case, please refer to Veeam best practices to properly size the Windows server hosting the service, and plan on using at least Microsoft SQL Standard.

Service account


Security best practices suggest using a dedicated account to run the different Veeam services. This is usually referred as a service account because it is a user that will not be used for interactive logins, but rather only to run the different Veeam services.

The use of a service account has some advantages that providers should consider:

- The account can be configured with a very complex password, which only the minimum amount of administrators that will manage the service will know
- Regular accounts can follow security rules about changing their passwords regularly, without the risk to stop any service because the service account can use a dedicated user ID with an exception to the password expiration policy
- It is easier to trace and log activities for the different services over the network, both for debugging and for auditing purposes. For example, instead of seeing the same administrator account in every log, a service provider can create a service account as veeam-service, and whenever a log will report this user, administrators will know that the traced activity is related to Veeam services.


By default, the installation wizard of Veeam Backup & Replication Server uses LOCALSYSTEM as the service account to execute the service.

It is better to create and use a dedicated account to run the services. Once the account has been created, either as a local account or an Active Directory account, service providers need to add this user to the local administrators of the server that will host the Veeam Backup & Replication server. Then, they can use the account during the installation by selecting **Let me specify different settings**:



2.1: Specify different settings during Veeam Backup & Replication setup


In the following step of the wizard, administrators need to specify the service account:



2.2: Specify a service account for Veeam Backup & Replication

The service account is also used for the authentication in the locally installed SQL Server Express.

Then, administrators will see this option in the last step:



2.3: Check for updates once the product is installed and periodically

This option allows the Veeam Backup & Replication server to connect to the Veeam update notification server (<http://dev.veeam.com>), so that it will notify administrators about the availability of updates for the software. See the later chapter **Regular maintenance of the components** for additional details.

Firewall

Once deployed, Veeam Backup & Replication has different components, listening over different TCP ports:

Service	Port
Catalog Service	9393
Veeam Backup Service	9392
Veeam Backup Service over SSL	9401
Veeam Cloud Connect Service	6169

Veeam Cloud Connect does not need the catalog service because there is no local backup activity that stores file-level information in the catalog itself. However, different Veeam components rely on the catalog for their operations so you should install it anyway to avoid undesired results.

Monitoring

Once deployed, Veeam Backup & Replication Server has different services installed in the Windows machine that you should monitor to guarantee the best Availability for the service:

Service Display name	Service Name	Startup Type	Log On as
SQL Server (VEEAMSQL2012)	MSSQL\$VEEAMSQL2012	Automatic	Local System
Veeam Backup Service	VeeamBackupSvc	Automatic (Delayed Start)	CLOUDCONNECT\svc vbr
Veeam Cloud Connect Service	VeeamCloudSvc	Automatic (Delayed Start)	CLOUDCONNECT\svc vbr
Veeam Data Mover Service	VeeamTransportSvc	Automatic	Local System
Veeam Guest Catalog Service	VeeamCatalogSvc	Automatic (Delayed Start)	CLOUDCONNECT\svc vbr
Veeam Installer Service	VeeamDeploySvc	Automatic	Local System

Note: There are additional Veeam services deployed as part of the default installation. They are not in this list because they are not involved in a Veeam Cloud Connect infrastructure.


Note: Microsoft SQL Service is available in the server only if during the installation has been chosen to install the default Express version. If a dedicated SQL Server is created during the design, this service may not be available in the Veeam Backup & Replication server.

Protection

From a protection standpoint, this machine is the most important piece of the environment. Since it cannot be installed in multiple instances, a good way to protect it is to run it as a VM and then rely on the underlying hypervisor for High Availability. Features like VMware vSphere HA or Hyper-V Failover Clustering can protect it and guarantee quick recovery times if the single hypervisor node where the VM is running fails. If a service provider needs an additional level of protection, he can also plan to use Veeam Backup & Replication itself and replicate this virtual machine every few hours; if anything happens, he can power up the replicated machine in a few minutes. In addition, service providers can and should use Veeam configuration backup in order to back up the overall configuration of the Cloud Connect environment, and plan to have a restore plan if anything happens to this machine and the corruption is replicated to the replica.

Standalone Console


One of the new features of Veeam Backup & Replication v9 is the standalone console. With this new option, service providers can install only the minimum required services in the Veeam Backup & Replication server — and control and manage it from a different workstation — by using the console. This removes the need to have a remote desktop connection towards the Veeam Backup & Replication server, and allows for more than two concurrent connections to the Veeam Backup & Replication server itself.



2.4: The Veeam Backup & Replication installation splash screen

When choosing the installation of the different components, a service provider can install Veeam Backup & Replication on the Veeam Backup & Replication server and then use the installer media to deploy the console on a different machine.

Once the administrator installs the two components, he can connect the console to the Veeam Backup & Replication server by filling in the required fields on the connection screen:



2.5: Veeam Standalone Console

Firewall

Incoming: No incoming connection is needed.

Outgoing: The standalone console connects to the Veeam Backup & Replication server over port TCP/9392 (if it has not been modified during Veeam Backup & Replication Service installation)

Monitoring


Because the standalone console is a client component, there is no need to monitor any service on the workstation running it.

Protection

Another great advantage of the standalone console is that it doesn't hold any data but rather shows only information retrieved from Veeam Backup & Replication server upon connection. Because of this, no protection is necessary on the workstation running the standalone console. If needed, providers can simply re-install the console on another workstation or laptop and connect again to the Veeam Backup & Replication server.

Enterprise Manager

Veeam Enterprise Manager is the service responsible for exposing to users the web interface of Veeam Backup & Replication and RESTful API. In a Veeam Cloud Connect environment, the latter is an important component if the service provider plans to develop and offer to users a custom portal for managing their Veeam Cloud Connect subscriptions.



2.6: Veeam Enterprise Manager

Veeam Enterprise Manager is a Windows Service; Veeam requires a modern 64-bit OS, like Windows Server 2008 R2 and above. It can be deployed on the same machine as Veeam backup service or on a dedicated machine. The choice to create and operate a separated machine for Veeam Enterprise Manager involves scalability considerations: If many users are going to interact with Veeam Cloud Connect via RESTful API, a service provider should plan to have a dedicated machine.

Furthermore, a dedicated machine is an additional effective layer of security: Because an optional custom portal will only connect to Veeam Enterprise Manager, a service provider can have additional firewall rules for the communications between Veeam Enterprise Manager itself and the Veeam backup server. When offering DRaaS services, the Cloud Portal is installed as an additional component of Veeam Enterprise Manager and exposed to internet so Veeam customers can reach it. Having this server separated from the Veeam Backup & Replication server can increase the overall security.


If a service provider chooses a dedicated machine, it should also have a dedicated Microsoft SQL Server locally installed to manage data stored by Veeam Enterprise Manager itself.

Because of the light load created by Veeam Cloud Connect, the default SQL Express installation is fine to use. However, you should carefully evaluate the amount of expected data to decide which edition of Microsoft SQL Server (Express, Standard or Enterprise) is best suited for Veeam Enterprise Manager.

Service Account


By default, the installation wizard of Veeam Enterprise Manager uses LOCALSYSTEM as the service account to execute the service. As explained in the previous chapter, it's better to create and use a dedicated account to run the services.

Once the account has been created — either a local account or an Active Directory account — service providers need to add this user to the local administrators of the server that will host Veeam Enterprise Manager. Then, they can use the account during the installation by selecting **Let me specify different settings**:



2.7: Specify custom configuration settings during Veeam Enterprise Manager installation

In the following step of the wizard, administrators will need to specify the service account:



The screenshot shows the 'Service Account Credentials' step of the 'Veeam Backup Enterprise Manager Setup' wizard. The window title is 'Veeam Backup Enterprise Manager Setup'. The subtitle is 'Specify Veeam Backup Enterprise Manager service account credentials.' There are two radio buttons: 'LOCAL SYSTEM account (recommended)' and 'The following user account:'. The second option is selected. Below it, there is a text box for 'User name:' containing 'CLOUDCONNECT\svc-em' and a 'Browse...' button. Below that is a 'Password:' field with masked characters and a visibility icon. At the bottom, there are '< Back', 'Next >', and 'Cancel' buttons.

2.8: Specify a service account for Veeam Enterprise Manager

The service account is also used for the authentication in the locally installed SQL Server Express.

Firewall

Once deployed, Veeam Enterprise Manager has different components, listening over different TCP ports:

Service	Port
Catalog Service	9393
Enterprise Manager Service	9394
Web UI over http	9080
Web UI over https	9443
RESTful API over http	9399
RESTful API over https	9398
Cloud Connect Portal	6443

For maximum security, you should enable only the HTTPS connections on the firewall and not the unprotected HTTP ones. Veeam Cloud Connect will not need the catalog service because there is no local backup activity that stores file information in Veeam Enterprise Manager.

In the table, you can see the suggested ports to open in bold.


Monitoring

Once deployed, Veeam Enterprise Manager has different services installed in the Windows machine that you should monitor to guarantee the best Availability of the service:

Service Display name	Service Name	Startup Type	Log On as
SQL Server (VEEAMSQL2012)	MSSQL\$VEEAMSQL2012	Automatic	Local System
Veeam Backup Enterprise Manager	VeeamEnterpriseManagerSvc	Automatic (Delayed Start)	CLOUDCONNECT\svc vbr
Veeam Guest Catalog Service	VeeamCatalogSvc	Automatic (Delayed Start)	CLOUDCONNECT\svc vbr
Veeam RESTful API Service	VeeamRETSvc	Automatic (Delayed Start)	Local System
World Wide Web Publishing Service	W3SVC	Automatic	CLOUDCONNECT\svc vbr

Web service

In the list of services, there is the World Wide Web Publishing Service, better known as IIS (Internet Information Services). This is the native Windows web server, and Veeam Enterprise Manager uses it to publish two web interfaces:



Name	ID	Status	Binding	Path
VeeamBackup	2	Started (http)	*:9080 (http);:9443 (https)	C:\Program Files\Veeam\Backup and Replication\Enterprise Manager\WebApp\
VeeamCloudConnectPortal	3	Started (http)	:6443 (https)	C:\Program Files\Veeam\Backup and Replication\CloudPortal\

2.9: Enterprise Manager and Cloud Portal are published via IIS


For any additional configuration of these two web sites, service providers can use IIS-native options.

Protection

Veeam Enterprise Manager does not hold any Veeam Cloud Connect information, and only communicates to Veeam Backup Service. If anything happens to the latter, Veeam Enterprise Manager cannot operate. You should have Veeam Enterprise Manager running on a VM, protected with an image-level backup of the entire VM. What needs protection is the underlying SQL database, plus optional customization done to the websites.


Cloud Portal

Veeam Cloud Connect Portal for Service Providers (in short, Cloud Portal) is an additional and optional component included in the installation of Veeam Enterprise Manager:



2.10: Veeam Cloud Connect Portal for Service Providers is an additional component of Veeam Enterprise Manager


If a service provider plans to offer replication/DRaaS to its customers, this component can be selected and installed during the installation of Veeam Enterprise Manager. In fact, the Cloud Portal cannot be installed as an isolated component, even if there is a dedicated MSI installer file in the installation media. Any attempt to install the Cloud Portal on a machine without Veeam Enterprise Manager leads to this error:



2.11: Cloud Portal cannot be installed without Veeam Enterprise Manager

Because of this, it makes even more sense to deploy Veeam Enterprise Manager together with the Cloud Portal on a dedicated machine..

Once installed, Cloud Portal can be accessed over an HTTPS connection over port TCP/6443. Service providers can re-map this port either by reconfiguring the website in IIS or by using a firewall.



2.12: The Cloud Connect Portal login page

Firewall

Once deployed, the Cloud Portal listens for incoming connections over a single TCP port:

Service	Port
Cloud Connect Portal	6443

This port is already listed among the open ports of Veeam Enterprise Manager. No further action is needed.

Monitoring

Cloud Portal is an additional component of Veeam Enterprise Manager. Please refer to the previous chapter related to Veeam Enterprise Manager for additional information about monitoring.

Protection

Cloud Portal is an additional component of Enterprise Manager. Please refer to the previous chapter related to Veeam Enterprise Manager for additional information about protection.

2.5 Cloud Gateways

Cloud gateways are the components responsible for receiving external connections from customers and tunneling all data transmissions over a single TCP port (also UDP for DRaaS services), protected by a SSL certificate.

Cloud gateway is a Windows service, so the best platform is a modern 64-bit OS like Windows Server 2012 R2. The correct sizing of a cloud gateway depends on the amount of traffic the service provider expects to receive, and on the redundancy design to be realized. You should note that Veeam encryption for backup and backup copy jobs is not managed by the cloud gateways, but directly by the target data movers (WAN accelerators and/or backup repositories). Cloud gateways are responsible for the SSL communications and data transfers, and their compute requirements are low.

A single connection from a customer consumes around 512 KB of memory. So, 1 GB of memory in a cloud gateway can be used to receive up to 2,000 concurrent connections.

A group of cloud gateways can work in concert to create a pool. They can all receive and manage incoming connections from customers and can balance these connections between them without the help of any external load balancer. If any gateway fails, another gateway can take care of the existing connections, giving continuity to customers' operations. This book will explain the interaction with external load balancers later.

In order to offer a reliable connection to customers, a service provider will deploy multiple cloud gateways following a N+K redundancy design. N is the minimum number of always- available gateways, and K is the number of gateways that can be lost. A typical redundancy design is N+1, where there is one more gateway than required to manage all incoming connections, so the service provider can lose up to one cloud gateway at any given time and still guarantee the planned level of service. Additional designs can be N+2 or others. Any service provider can find the right balance between the desired level of redundancy and the need to deploy additional gateways in advance.

Firewall

The cloud gateway is the external component of a Veeam Cloud Connect infrastructure that is responsible for interconnecting the different customers to the services offered by the service provider over the internet. Because it faces the internet, it is imperative to properly configure the firewall rules:

Service	Port
Veeam Cloud Gateway Service	TCP / UDP 6180 from outside
Veeam Cloud Gateway Service	TCP 6168 from VBR Server
Veeam Installer Service	TCP 6160 from VBR Server

While the cloud gateway has to talk with the managing Veeam Backup & Replication server over TCP ports 6168 and 6160, it has to be reached only with the single TCP/UDP port used by the service over the internet. By default, this port is 6180, but the service provider may customize this, like using 443 if customers have strict egress firewall rules.

Furthermore, additional connections are needed for the proper operation of a cloud gateway:

From	To	Destination Port	Notes
Cloud Gateway	SP VBR server	TCP 6169	SP VBR server listens to cloud commands from the tenant side. Tenant cloud commands are passed to the Veeam Cloud Connect Service via the cloud gateway
Cloud Gateway	SP backup repository	TCP 2500 to 5000	Default range of ports used as transmission channels for backup jobs. For every TCP connection that a job uses, one port from this range is assigned
Cloud Gateway	SP backup proxy	TCP 2500 to 5000	Default range of ports used as transmission channels for replication jobs. For every TCP connection that a job uses, one port from this range is assigned
Cloud Gateway	Provider- side network extension appliance	UDP 1195	Port used to establish secure VPN connection for network extension during partial site failover. If a tenant has several IP networks, additional odd ports should be opened starting from 1195 — one port per tenant's IP network. For example, 1195, 1197, 1199 to connect 3 different networks of the same tenant.

NOTE: When an SMB share or a deduplication appliance is used, the service provider backup repository is considered the backup gateway server directly connected to the device.

Monitoring

Once deployed, the cloud gateway has different services installed in the Windows machine that should be monitored to guarantee the best availability of the service:

Service Display name	Service Name	Startup Type	Log On as
Veeam Cloud Gateway Service	VeeamGateSvc	Automatic	Local System
Veeam Data Mover Service	VeeamTransportSvc	Automatic	Local System
Veeam Installer Service	VeeamDeploySvc	Automatic	Local System

Note: There are additional Veeam services deployed as part of the default installation. We ignored them in this list, as they are not involved in a Veeam Cloud Connect infrastructure.

Protection

From a protection standpoint, a cloud gateway does not need to be saved, because there is no permanent data on it. Additionally, a new cloud gateway can be deployed in a few minutes while other existing cloud gateways are serving customers.


2.6 WAN Accelerators

WAN accelerators are optional components that can be deployed at the service provider to improve bandwidth utilization of remote backups and replicas sent by customers. Even if any Veeam Cloud Connect operation can be executed without WAN accelerators, WAN accelerators become mandatory components for a service providers willing to offer remote backup or replication services. Several customers will probably have Veeam WAN accelerators in their infrastructure, so in order to leverage them, the service provider will need to deploy and configure them. Also, WAN accelerators are enabled in the Veeam Cloud Connect license given to service providers without need for further licensing, so there is no licensing concern for the service provider when deploying them.

In addition, starting from Veeam Backup & Replication v9, when the target of a job is Veeam Cloud Connect, customers using the Enterprise license are entitled to use WAN acceleration, while previously they had to have an Enterprise Plus license (and it's still like this in v9 for jobs not involving Veeam Cloud Connect).

WAN accelerators at the service provider sit between cloud gateways and repositories (for backup and backup copy jobs) or proxies (for replica jobs). They help improve the bandwidth utilization by caching blocks internally, avoiding the need to transmit every block over the wire.

WAN accelerator is a windows service, so the best platform is a modern 64-bit OS like Windows Server 2012 R2. The same design considerations made for local Veeam Backup & Replication deployments can be applied also in a Veeam Cloud Connect scenario when it comes to WAN accelerators: 8 GB of RAM at least, a fast disk for the cache (a SSD disk or SSD-backed volume is not mandatory, but preferred), and the correct sizing for the cache itself. In addition to the global cache configured during its deployment, a WAN accelerator consumes 20 GB per 1 TB of source data. A good choice is to use a dedicated volume for caching, so when it is filled, it does not create problems for the Windows OS and its running services.



2.13: Prefer a dedicated drive for the WAN accelerator cache

A single WAN accelerator can saturate links up to 150 Mbps on average, depending on the workload. However, some users choose to use WAN acceleration on much faster links in order to optimize bandwidth consumption of a shared WAN link. If bandwidth consumption is not a concern, using direct transfer mode usually achieves a better data transfer performance and a shorter job completion time on faster links.

When a service provider configures a new cloud repository for a customer and assigns a WAN accelerator, this relationship is fixed. Even if a service provider has multiple WAN accelerators, only one is used for a given cloud repository, until this configuration is changed. So, when adding new customers or assigning new resources, a service provider will need to balance the assignment of WAN accelerators to customers manually. When sharing one WAN accelerator among multiple customers, a service provider will have to take into account the total bandwidth of the customers and the expected storage consumption for the cache. For example, one WAN accelerator with a 50-Mbits bandwidth could be the target of five customers having each a 10-Mbits upload speed. Usually the maximum ratio for sharing a WAN accelerator is 5:1. Additionally, you should always take into account the failure domain: the more tenants that are connected to the same WAN accelerator, the more that are affected when one of WAN accelerator is not available.

Firewall

Once deployed, a WAN accelerator has different components, listening over different TCP ports:

Service	Port
Veeam Installer Service	6160
Veeam Data Mover Service	6162
Veeam WAN Accelerator Service	6164
Veeam WAN Accelerator Service	6165

Ports 6160, 6162 and 6164 need to be open towards the Veeam Backup & Replication server controlling the WAN accelerator. There are also communications between the different WAN accelerators (source and target) happening over port 6164 (the controlling port for RPC calls) and 6165 (data transfer between WAN accelerators). This last communication is tunneled by the cloud gateway.

Finally, ports 2500 to 5000 need to be open between WAN accelerators and backup repositories for data transfers of WAN-accelerated jobs.

Monitoring

Once deployed, the WAN accelerator has different services installed on the Windows machine that you should monitor to guarantee the best Availability of the service:

Service Display name	Service Name	Startup Type	Log On as
Veeam WAN Accelerator Service	VeeamGateSvc	Automatic	Local System
Veeam Data Mover Service	VeeamTransportSvc	Automatic	Local System
Veeam Installer Service	VeeamDeploySvc	Automatic	Local System


Protection

From a protection standpoint, WAN accelerators need to be protected properly. A backup job that is WAN accelerated cannot failover to a direct connection if the WAN accelerator fails.

The job itself fails until the WAN accelerator is restored or the job is reconfigured for direct mode, and this needs to be done at both ends (service provider and customer). For this reason, having WAN accelerators hosted as virtual machines on a hypervisor with High Availability capabilities is a best practice. There is no need to back up a WAN accelerator because its cache can be populated from scratch when it is redeployed. In order to avoid low performance while the cache is warming up after a redeployment, the service provider can warm the cache before placing the new WAN accelerator into production.

Backup Repositories

Backup repositories are the destination of backup and backup copy jobs in Veeam Backup & Replication. You can create them using Windows or Linux machines with local attached or remote storage carved out of a SAN, or they can be a storage appliance exposing its disk space via SMB protocol. They can also be one of the supported deduplication appliances.



2.14: Veeam Backup & Replication supports different repository types

Once a backup repository is configured and registered in the Veeam Backup & Replication console, a new **cloud repository** is created and assigned to the user during the creation of a new Veeam Cloud Connect customer consuming backup resources, using a portion of an existing backup repository. From a service point of view, a cloud repository is a remote repository exclusively assigned to its user. From an infrastructure point of view instead, a cloud repository is a sub-folder of a backup repository with an applied quota.

For a Veeam Cloud Connect deployment, there are no special requirements for repositories, but the general rules of Veeam Backup & Replication are still valid. You should use a Windows or Linux server instead of an SMB share so that a proper Veeam data mover service can be deployed on the repository machine. With this service running, all write/read operations are delegated to this service.

Additional caution should be taken for the use of deduplication appliances: As tenants have the option to encrypt their backup files, and service providers cannot forcefully disable this option (but they can force mandatory encryption if needed), encryption itself can nullify any advantage of deduplication appliances, which are going to be filled with encrypted backups that cannot be deduplicated. If a service provider can control the configuration of incoming backup or backup copy jobs, a deduplication appliance may be a good choice, but for those service providers offering Veeam Cloud Connect Backup publicly, a deduplication appliance may not be the best choice.

Firewall

Once deployed, a repository has different components, listening over different TCP ports:

Service	Port
SSH Server (Linux only)	22
Veeam Installer Service (Windows only)	6160
Veeam Data Mover Service (control)	6162
Veeam Data Mover Service (data)	2500–5000

Ports from 2500 to 5000 need to be open between WAN Accelerators and Backup Repositories for data transfers of WAN accelerated jobs, or between Backup Repositories for direct jobs.

Monitoring

Monitoring considerations are different for Windows and Linux repositories. The latter in fact has no permanent component installed: Instead, a temporary component is loaded dynamically every time a job is executed. For this reason, the monitoring information is split for the two options:

Windows repository

Service Display name	Service Name	Startup Type	Log On as
Veeam Data Mover Service	VeeamTransportSvc	Automatic	Local System
Veeam Installer Service	VeeamDeploySvc	Automatic	Local System

Linux repository

Administrators need to verify that the Linux repository has the SSH server enabled and running, and Perl subsystem available. The Veeam Backup & Replication server connects to the Linux machine via SSH, copies the temporary binaries and executes them using Perl. No permanent Veeam component is installed in the repository, so there is no Veeam component to monitor.


Service providers may want to monitor the SSH server to guarantee that is up and running.

Protection

From a protection standpoint, backup repositories need proper protection. They are the components storing customers' backup data, and losing them means losing the customers' data. Because of the many available technologies used to build a repository, there are no universal considerations that apply to every scenario. A service provider must carefully evaluate the available options in regards to the technology used to create the backup repository.

Concurrency

The concurrency limits of a repository should be carefully evaluated by the service provider, otherwise customers could be stuck with their jobs waiting for available resources at the service provider.




2.15: Configure carefully the repository load control

Server sizing

Storage space sizing is not covered in this book. There are too many options available on the market to build a Veeam backup repository, and each solution has its own limits in terms of number of disks, stripes, volumes, etc. The only limit on Veeam Cloud Connect Backup is 2 PB (petabyte) for a single cloud repository.

In regards to the memory sizing of a backup repository, it is important to know how a Veeam repository uses memory. Veeam Backup & Replication v9 has four different levels of storage optimization for a backup job:



2.16: Storage optimization options for a backup job

A repository uses memory to store incoming blocks. This queue collects all blocks coming from source data movers, caches them in memory and after some optimization, it is flushed to disk. This reduces the random I/O affecting the backup files to a minimum, while trying to serialize as many write operations as possible. The amount of memory consumed by the queue is simple to calculate: It uses **2 GB of memory per active job**.

However, this is not the only amount of memory consumed by the repository: Veeam backup files contain deduplicated information of the saved blocks. As with any deduplicated storage, there are metadata information stored along the file itself in order to keep track of stored blocks.

To improve performance, the repository loads dynamically this metadata information into memory. Starting from Veeam Backup & Replication v8 Update 2, the cache accelerates both write and read operations, but there are also differences in the way the cache is populated and used. The amount of consumed memory for metadata depends on the selected block size for deduplication:

VBK size	Optimization	VBK block size	Memory consumption for VBK metadata
1 TB	WAN target	256 KB	700 MB
1 TB	LAN target	512 KB	350 MB
1 TB	Local target	B1024 K	175 MB
1 TB	Local target 16+ TB	4096 KB	44 MB

Note: Starting from Veeam Backup & Replication v9, the new block size for Local target 16+ TB is 4 MB instead of 8 MB. The previous value for memory consumption was 22 MB.

By adjusting these values to a real scenario, service providers can estimate how much data a given repository will be able to process at a certain point in time (how much memory will be needed for an expected amount of processed data).

For example, if the memory is 8 GB, and you assume the OS and all other running processes use 1 GB, 7 GB of memory can handle around 41 TB of backup files at the default block size. This also includes the additional incremental files of a backup chain.

If a given backup repository is assigned to 10 different customers and all of them are executing their jobs at the same time, the total memory must be divided among all the jobs. The Veeam repository constantly consume the same amount of RAM, because it can dynamically load and offload metadata, but planning for the maximum possible consumption is a good choice to be prepared for the worst-case scenario.

Finally, backup and backup copy jobs are configured by the customer and not by the service provider. There is no direct way for the service provider to plan for an accurate utilization of the backup repository memory, because he does not know in advance which block size will be used and what the total size of a backup set will be. However, the quota configured for a tenant in Veeam Cloud Connect can also be considered the maximum possible size of a backup file of a customer. For these reasons, proper monitoring of the backup repository is paramount, so the provider can quickly identify when the system is too stressed.


The maximum possible size of a single cloud repository is 2 PB (2,097,151 GB to be precise); the memory required to manage this amount of data at the default block size would be theoretically around 350 GB. This value will never be reached because there are mechanisms in place to flush the cache. However, it is up to the service provider to design a single large backup repository, or decide to have multiple pods (which will be discussed later) and size their memory accordingly.

Because of the creation of several cloud repositories on top of the same backup repository, some additional design principles should be considered.

Two main design choices are suggested. Both solutions are effective and can be used for a Veeam Cloud Connect Backup infrastructure. The choice between the two depends, among other technical and business reasons, on the technical skills of the service provider's IT department and their knowledge of the described technologies.

Pod Design

The first design is what can be called a pod. A pod is a single repository built with any supported storage (local disks in a Windows or Linux machine, a SAN, a NAS or a deduplication appliance) that has a fixed size or expandable up to a certain limit.




2.17: Pod design for repositories

With this kind of repository, a service provider needs to plan how to distribute customers among the several repositories he could have. Above all, however, he must keep some free space for a future increase in the cloud repositories' quotas and transform operations. A customer may start with a small amount of space, but after some time, the customer could ask for an increase in the storage quota. If there is no free space left in the repository, the service provider will be able to satisfy the customer's request only by migrating the customer's backup files into another repository. This can be done almost transparently, but it involves some manual activities on the part of the service provider and some downtime in the customer's Veeam Cloud Connect service. Cloud repository quotas are strictly applied, but as long as the customer is not using the entire amount of the assigned quota, the service provider can use some over-commitment. However, the service provider should carefully evaluate level of over-commitment to avoid any interruption of the service.

A pod design can be expanded by adding additional pods aside of the first one (thus the meaning of the name). The different pods do not share their storage resources with each other, the service provider will manually balance cloud repositories among them, and move any customer from one to another as needed.

Single namespace scale-out design

The second type of design is **single namespace scale-out design**. This design is more complex than the previous one, but it has some advantages. To create it, a Veeam repository is connected and uses storage resources from a scale-out storage solution that can be expanded over time without changing the configuration of the exposed resources. There are several solutions — both open-source and commercial — with these capabilities, and Veeam does not promote any one above the others.



2.18: Single namespace scale-out design for repositories

The important aspect of this design is the **single namespace**. Instead of adding additional storage with a new path to the Veeam console, the addition of a new node to the scale-out array in this scenario does not change the path Veeam needs to use to save data. Simply put, once a new node with some capacity is added, the repository is going to expose the same path with a transparently increased capacity.

This solution helps service providers avoid capacity problems in their repository design, especially when enabling self-service capabilities to their customers. If a customer can set up his storage quota freely, proper capacity planning cannot be effective: A scale-out approach helps to react quickly to a capacity shortage without changing any configuration to the repository structure.

If concurrency becomes a problem with this approach, a service provider can deploy additional repositories using the same scale-out storage. Even if the same storage path cannot be exposed by more than one repository at the same time, a service provider can create multiple paths (directories) in the same storage, and then use multiple repositories at the same time. This will add more concurrency to accommodate customer's activities.

NOTE: Veeam Scale-out Backup Repository™, a new technology introduced in Veeam Backup & Replication v9, cannot be used to host cloud repositories. The scale-out design can be used as an alternative.


Proxies

Proxies, which are deployed at the service provider, are the components responsible for receiving replication data from customers (from their source proxies) and writing data onto the virtual disks of the VMs hosted on the hypervisor where the service provider offers the DRaaS service.

The Veeam proxy is a Windows service, so the best platform to use is a modern 64-bit OS like Windows Server 2012 R2. The correct sizing of a proxy depends on the expected amount of traffic the service provider will receive and on the redundancy design to be realized. At least one proxy per hypervisor cluster must be deployed so that this proxy has access to all the available underlying storage and can then write the received data onto it. To improve performance and resiliency, multiple proxies should be deployed.

A group of proxies can work in concert to create a **pool**. They can all receive and manage incoming replication data from customers and balance these connections between them. If any proxy fails, another proxy can take care of the existing connections, giving continuity to customer operations.

To streamline the replication process, service providers can deploy a backup proxy as a virtual machine. The virtual backup proxy must be registered on an ESXi host that has a direct connection to the target datastore. In this case, the backup proxy will be able to use the virtual appliance transport mode to write replica data to the target. This is the most effective transport mode for the target-side proxy:




2.19: Backup proxy transport mode

During the first run of a replication job, Veeam Backup & Replication creates a replica with empty virtual disks on the target datastore. If the virtual appliance mode is used, replica virtual disks are mounted to the backup proxy and populated through the ESXi host I/O stack. This results in increased write speed and fail-safe replication to ESXi targets.

If the backup proxy is deployed on a physical server or if the virtual appliance mode cannot be used for other reasons, Veeam Backup & Replication uses the network transport mode to populate replica disk files. For these reasons, you should deploy proxies VMs and leave the transport mode configuration as **automatic selection**.

Concurrency

The service provider should carefully evaluate the concurrency limits of a proxy. Supposing that a completely shared environment exists and every proxy deployed in the environment can access every available datastore, the amount of replication jobs a service provider can receive concurrently is the sum of the maximum concurrent tasks of all the proxies:



2.20: Proxy maximum concurrent tasks

Service providers should follow Veeam best practices of a 1:1 ratio between available CPU in a proxy and the number of maximum concurrent tasks. For example, a service provider can have five virtual proxies, each with four virtual CPUs, and thus be able to receive as many as 20 concurrent replication jobs from tenants.

NOTE: Every tenant can perform one replication job targeted to the cloud host simultaneously. Veeam Cloud Connect does not support parallel processing. The number of concurrent tasks of proxies equals the maximum amount of incoming tenants a service provider can receive at a given point in time.

Firewall

Once deployed, a proxy has different components listening over different TCP ports:

Service	Port
Veeam Installer Service	6160
Veeam Data Mover Service (control)	6162
Veeam Data Mover Service (data)	2500–5000

Ports from 2500 to 5000 need to be open between proxies and cloud gateways and between WAN accelerators and proxies for WAN-accelerated data transfers.

Monitoring

Once deployed, a proxy has different services installed in the Windows machine that should be monitored to guarantee the best Availability of the service:

Service Display name	Service Name	Startup Type	Log On as
Veeam Data Mover Service	VeeamTransportSvc	Automatic	Local System
Veeam Installer Service	VeeamDeploySvc	Automatic	Local System

Protection

From a protection standpoint, a proxy does not need to be saved because there is no permanent data on it. Also, a new proxy can be deployed in a few minutes while other existing proxies receive customer data.


Network Extension Appliances

To enable communication between production VMs on the tenant's side and VM replicas at the service provider, Veeam Backup & Replication uses network extension appliances. A network extension appliance (NEA) is a tiny, hardened, Linux-based VM (1 VCPU, 512 MB RAM, 44 MB ISO file + virtual floppy disk for configuration) deployed automatically by Veeam Backup & Replication server on the virtualization hosts on which tenant VMs and their replicas reside.

NEA have two main purposes:

- During **full failover**, it provides both access to internet for replica VMs and access to replica VMs from the internet
- During **partial failover**, it extends customer network to the service provider environment, using L2 (Layer 2) technologies, so that production VMs can communicate with replica VMs.


For every tenant who plans to replicate VMs to the service provider and use the built-in networking and failover capabilities described above, at least two NEAs should be deployed: one on the service provider's side and the other on the tenant's side.



2.21: General overview of network extension appliances

The network extension appliance on the service provider side is deployed on the virtualization platform that acts as a replication target. The NEA has at least two network interfaces:

- The external interface is connected to the service provider external network, ideally where also cloud gateways are deployed. This network uses public IPs assigned to each NEA so that the appliance itself can act as a firewall or gateway for every replica VM
- One or more internal interfaces are created once a hardware plan is assigned to a new tenant. Veeam Cloud Connect uses VLANs as a way to isolate network resources assigned to every tenant. Each network of each tenant is created in the virtualized platform as a new port group, tagged with a unique VLAN ID (taken from a VLAN pool preloaded during the Veeam Cloud Connect initial configuration). VLANs are not routed between each other at the physical switching layer. The only way a VLAN segment can communicate with another one is by the NEA, which guarantees the complete isolation of tenants.




2.22: VLANs are used to create network isolation and multi-tenancy

Full failover

During a full failover, a network extension appliance loads on the internal interfaces the same IP addresses that gateways use by each tenant for their virtual machines, one per network or port group. This way, no change to VM IP configuration is needed during the failover and the NEA can act as the new default gateway.

The tenant allows access to replica VMs during the configuration of a cloud failover plan:



2.23: Setting public IP destination rules in a cloud failover plan

Each rule created in the public IP step of the wizard becomes a new ingress rule on the NEA, from the public IP address to the private IP of the selected VM. Chapter 5 will give you additional information about this topic.

Partial Failover

During partial failover, only single replica VMs are selected by a tenant to be powered on at the service provider and those VMs will be able to connect to their original network at the tenant side, thanks to the network extension appliance.

In this scenario, there are two appliances involved in the process:

- The network appliance deployed at the tenant has one network interface, connected to the same port group as the virtual machines (multiple appliances are deployed when the end user has multiple port groups with VMs). This appliance starts the OpenVPN client and connects to the cloud gateway, and from here to the network appliance deployed at the service provider side
- The network appliance at the service provider side receives the connections from the OpenVPN client thanks to an OpenVPN server running inside it. The result is the creation of a L2 tunnel between the two sites.

On top of the L2 tunnel, a Proxy ARP solution running inside both network appliances forwards L2 packets from one side to the other, and vice versa. The result is that VMs can use the same subnet or broadcast domain, regardless of the site where they are powered on.

Firewall

Because of the different networking configurations that a network extension appliance can have, a general firewall configuration for the appliance itself is not possible. There is one TCP port that is always open on the external interface:

Service	Port
SSHD	22

In addition, every rule created by a tenant in a cloud failover plan becomes a new firewall port open in the network appliance.

The only difference is that while SSHD needs to be reached only from Veeam Backup & Replication server in order to control and reconfigure the network appliance, every other rule is meant to have no defined source IP address, as end user may need to connect from different locations to their replica VMs.

NOTE: Because the network extension appliance is already a firewall and opens only the minimum amount of ports required by an end user, there is little sense in putting an additional firewall in front of it. Service providers can filter and monitor incoming connections, if needed, using firewalls operating at L2 (or in transparent mode). This allows to use the real public IP of an appliance instead of complicated multiple NAT levels.

Monitoring

There is no need to monitor a network extension appliance because it is powered on on- demand by the Veeam Backup & Replication server at the service provider when needed in reaction to end user activities like the start of a partial or a full failover.

Protection

From a protection standpoint, a network extension appliance does not need to be saved because there is no permanent data on it. A significant part of configuration (the content of the floppy image) is created during appliance deployment; the managing VBR server passes additional configuration upon boot. In both cases, data is stored in the Veeam Backup & Replication server and reconfigured upon a redeployment of the NEA.

Additional components

Even if they are not part of the Veeam Cloud Connect infrastructure, the following components help to successfully create and operate the infrastructure.

Active Directory domain controllers

Active Directory is the directory service developed by Microsoft years ago, starting with Windows 2000. Directory services allow central authentication and authorization for all users and computers in the network based on a Windows domain, assigning and enforcing security policies for all computers and users and installing or updating software.

The ability to centrally authenticate and authorize access to resources is an important solution to guarantee optimal security of any IT environment. In addition, all Veeam Cloud Connect components (with the exception of Linux repositories and the network extension appliances) are designed to be executed on Windows machines, so having Active Directory in place makes perfect sense.

Finally, Active Directory offers integrated DNS services. Any IT infrastructure heavily relies on proper DNS configuration (with both forward and reverse resolution correctly configured) to reach all the different components.

For these reasons, Active Directory is recommended in a Veeam Cloud Connect environment.

Firewalls

By its nature, Veeam Cloud Connect is a service that needs to be exposed over the public internet to serve users. Because of this, network security solutions like firewalls should be deployed and properly configured in order to protect Veeam Cloud Connect.

Different technical solutions and business requirements lead to different security designs. For this reason, it makes little sense to describe a detailed firewall design for Veeam Cloud Connect. Instead, this book suggests two high-level design concepts to use when protecting Veeam Cloud Connect:

- **Separate different logical components in different security zones:** For example, keep cloud gateways and NEAs in different and separated areas. Because they are components that need to be exposed over the internet, a compromise on these machines will not lead to a compromise of the entire Veeam Cloud Connect environment.
- **Reduce network connections to a minimum:** You should have firewalls authorizing any communication between components. You can do so by opening the minimum number of required TCP/UDP ports. Chapter 3 and 4 describe these ports in detail.

Load balancers

As explained before, different cloud gateways work as one logical pool to share the load and guarantee High Availability. They are designed to balance themselves without the help of any additional load balancer.

To better understand this design principle, remember this important design consideration: **Each cloud gateway needs to have its own public IPv4 address, regardless of whether it is directly configured on the cloud gateway itself (direct mode) or with a firewall in front of it (NAT mode).**

This is a mandatory configuration. For service providers worried about the consumption of public IP addresses, there is no need to have a large amount of cloud gateways even on large installations, so the use of public addresses should not be an issue for most service providers.

These requirements have a direct consequence on load balancing: A service provider cannot use a load balancer with shared IP address to publish multiple cloud gateways.

Only the initial connection from a tenant to the Veeam Cloud Connect environment needs balancing. This can be accomplished by using simple DNS Round Robin: The public FQDN (fully qualified domain name) of Veeam Cloud Connect can be configured in the DNS to have multiple A (host) records. This way, when a tenant connects to the assigned resources in Veeam Cloud Connect, he connects to one of the registered public IP addresses, realizing a simple balancing between the cloud gateways. An example configuration would look like this:

DNS name	record type	IP address
cc.virtualtothecore.com	A	185.62.37.98
cc.virtualtothecore.com	A	185.62.37.99
cc.virtualtothecore.com	A	185.62.37.100

NOTE: This example uses a personal but real domain name because when creating a real SSL certificate, every certificate authority checks the information of the applicant and the registered domain. A fake domain would create errors during the SSL verification phase.

For proper operations when using wildcard certificates (like *.virtualtothecore.com), the public hostname that resolves to each public IP assigned to a cloud gateway must use the same domain of the common DNS name used to publish the service. The same must be used in the SSL certificate. Therefore, in the example, the gateways must have their host names mapped in DNS as:

DNS name	record type	IP address
gtw1.virtualtothecore.com	A	185.62.37.98
gtw2.virtualtothecore.com	A	185.62.37.99
gtw3.virtualtothecore.com	A	185.62.37.100

One limit of this design may be that DNS does not have any notion of the state of the different cloud gateways, and users may receive information regarding failed or disabled gateways. This is not an issue in reality because Veeam Cloud Connect's client component reads all the A records from the DNS resolution and tries to connect to each of them until it can establish an initial connection. Once it has reached one of the cloud gateways, it receives a list of all the existing and available cloud gateways directly from the service provider's Veeam Backup & Replication server. The Veeam cloud service installed on the Veeam Backup & Replication server maintains and updates this list.

NOTE: *In order to optimize the use of DNS Round Robin and avoid connection problems caused by DNS caching, you should configure low TTL (time to live) values for the A records.*

During regular operations, the Veeam cloud service keeps a list of all existing activities happening on all cloud gateways, and it instructs new incoming tenants to use the less-used cloud gateway. As a result, Veeam Cloud Connect load balances directly without any need for external load-balancing solutions. Load balancing is based on the number of active connections per gateway.

When one of the cloud gateways fails, all existing connections are lost. Depending on the type of job that was going through this gateway, two scenarios can occur:

- Backup jobs are sensitive to network interruptions. Running jobs will fail, but subsequent retries will be sent to surviving cloud gateways. Customers will see a failed job and then a successful retry. Retry attempts are configured by default in any backup job; service providers should advise their customers to not change these parameters.
- Backup copy jobs can survive network interruptions. Depending on the duration of a network interruption, Backup copy jobs are likely to restore the connection in place, or if the TCP timeout threshold has been reached, backup copy jobs will be redirected to a surviving cloud gateway without any notification to the user about a failed connection.

Finally, a note on the failover process of cloud gateways from an end-user perspective: The list of available gateways is retrieved by the end-user component of Veeam Cloud Connect upon any job start or retry. The available gateways are listed in a specified order in which the first usable gateway is assigned #1, the second #2 and so on. The number assignment and the priority depends on the actual load (number of active tasks) of all gateways.

As long as the gateway marked as #1 is available, the end user keeps using this one. As soon as this gateway is not available, a new connection is automatically tried against #2; if this is available, the connection is automatically established and any running job is continued. If not, a connection is tried against the next gateway on the list. When all the gateways have been tried unsuccessfully, the running job fails and a new list is retrieved for the following retry.

Regular maintenance of the components

A critical environment designed to offer a service to external customers like Veeam Cloud Connect needs to guarantee the best possible experience and the highest possible uptime. For this reason, different operational criteria should be applied when managing Veeam Cloud Connect.

Splitting components

As outlined in this chapter, Veeam Cloud Connect is designed as a distributed architecture. To guarantee the best performance, each component can and should be deployed on a separated server, physical or virtual.

Even if multiple components can be installed in the same server, when a maintenance activity is needed or an issue is found, the downtime created can negatively affect each component deployed in the same server. If the Veeam Backup & Replication server is deployed together with one of the cloud gateways, the Veeam server will need stopped when maintenance is needed for a gateway, which creates downtime for the entire server.

By splitting each component over a different server and deploying several instances of those components that allow multiple instances, the chances to interrupt the entire Veeam Cloud Connect service because of a single unavailable server are extremely reduced.

Patches and upgrades

Veeam Cloud Connect is deployed for the vast majority of its components over Microsoft Windows operating systems. In order to guarantee stability and protection from vulnerabilities (especially for those components exposed to internet like the cloud gateways), service providers should check regularly for the availability of Microsoft patches and upgrades, and apply them as soon as possible. Again, the distributed architecture allows for a stacked patching activity of different components at different times.

Service providers also need to address the Veeam Backup & Replication updates. If the Veeam Backup & Replication server is able to reach Veeam Update Notification server (<http://dev.veeam.com>), the software will notify users about the availability of a new update.

Service providers that subscribe to receive email notifications from the VCSP program should also receive notifications about new updates a few weeks in advance when compared to the notification server. This notification system enables service providers to plan upgrades to their environment before each user sees the available update from the notification server.

Veeam Cloud Connect is backwards-compatible to previous versions up to one major release. For example, service providers using v9 can receive backups from customers using v8. However, a service provider cannot run a version older than what the customer uses regardless of this compatibility. The service provider should deploy any update before it's generally available to end users.

Time sync and DNS resolution


Two important aspects of every network are time sync and DNS resolution. In a Veeam Cloud Connect environment, these two services must be properly configured and monitored to guarantee correct operations.

Every server of the environment — both those running Veeam Cloud Connect components and the additional machines like Active Directory servers or the hypervisor hosts — must be synchronized with the same time source, and providers must be sure that time and time zone are correctly configured. Differences in time between the different components can lead to unexpected errors when operating Veeam Cloud Connect.

The same is true for DNS: Each and every component of the environment must be reachable by using forward DNS resolution. The DNS servers must be correctly configured and populated with every existing record (for both forward and reverse zones), and they must check to verify that they are operating correctly and are reachable over the network by any other server. Redundant DNS servers should be deployed to guarantee optimal uptime of the overall DNS service. If you follow this guide, you'll see two Windows Servers as both Active Directory controllers and DNS servers for the entire infrastructure.

Reference design for backup services

This chapter will describe a complete Veeam Cloud Connect Backup deployment at a service provider. By impersonating a provider willing to offer backup services, you will learn about design and deployment all the servers, networks and network rules necessary to run Veeam Cloud Connect Backup.



3.1: Cloud Connect Backup overview

All the components will be deployed in VMs running on top of a hypervisor in order to leverage the quick deployment times of new VMs starting from templates and to protect components that cannot be executed in multiple instances, like the Veeam Backup & Replication server. Because of this, sizing rules will be based on vCPUs (virtual CPUs) rather than sockets and cores. Additionally, the size adjustments of virtual disks will be easier than in a physical server, like when it will be needed to increase the WAN accelerator cache size for example. The only exception is backup repositories: For better performances, you should have physical backup repositories.

During this chapter (and the following one dedicated to replication services), these VMs are referred to as servers; please remember that physical servers only refer to the repositories.

Note: If you plan to use physical servers, adjust the CPU considerations to existing and available CPU models and specifications.


Network diagram for backups

The first activity that a service provider has to do when starting the deployment of Veeam Cloud Connect is design the overall infrastructure. This is paramount in order to better understand the relationships between the different components of Veeam Cloud Connect, network ports and services, and the way they communicate with each other.

The first activity is the creation of the following network diagram. This diagram is specifically about Veeam Cloud Connect Backup. In the next chapter, you will see another diagram dedicated to replication services.

Veeam Cloud Connect Backup

Service and Network diagram



3.2: Network diagram for Veeam Cloud Connect Backup

The diagram depicts the different Veeam Cloud Connect areas and the communication happening between the different components. Later in the chapter, a similar diagram will depict the detailed layout of the different servers. For additional information about network connections between the different Veeam components, you can refer to the Veeam Backup & Replication User Guide or the knowledge-base article KB1518 (<http://www.veeam.com/kb1518>).

The list of ports used does not change once the services are deployed, but there are two specific use cases in which additional firewall rules are needed temporarily, even if they are not directly related to Veeam Cloud Connect components:

1. **Disabling a gateway:** The Veeam Cloud Service running in the Veeam Backup & Replication server (management zone) needs to access the installer service running on the gateway (DMZ zone) on TCP/6160 in order to disable a gateway. If this port is not open, the gateway can still be disabled but the UI will freeze for a while.
2. **Installing updates:** It's recommended to temporarily disable firewall rules between the different security zones during updates because operations require multiple open ports, like SMB access to upload new .MSI installers to Windows machines, RPC access to restart services remotely and others.

The following parts of this chapter will explain the network diagram.

Security zones

The Veeam Cloud Connect environment is divided into different security zones, and different server types are placed in each zone. All the zones are protected from each other and from the outside by firewalls.

By applying different firewall rules to allow only the minimum amount of necessary connections between the different zones, the level of security is improved.

As described in the network diagram, there are four different areas:

- **DMZ:** This area hosts the cloud gateways and an optional web portal to offer users self-service capabilities. The portal is not a Veeam component, but it can be developed by a service provider to offer self-service operations to customers. This is the only area connected and reachable from users via a public internet connection (directly or via firewall or NAT).
- **Management:** This area hosts the management components of Veeam Cloud Connect. This area is not reachable from outside.
- **Storage:** This area hosts the WAN accelerators and the backup repositories. This area is not reachable from outside. A more complex design can also have WAN accelerators and repositories divided in two separated areas.
- **WAN (public):** This area is the public internet or, in general, the network outside of the Veeam Cloud Connect infrastructure where tenants are supposed to connect to the cloud gateways and their cloud repositories to consume the web portal, if available.

Firewall considerations

One of the reasons to separate the environment in several distinguished security zones is because of the possibility to limit the network connections between them to a minimum.

Inside the same zone, all servers are free to communicate with each other. For example, the Veeam Backup & Replication server can freely connect to Veeam Enterprise Manager.

Assume that all connections between security zones are denied unless explicitly allowed via a firewall rule. For a complete list of the required network ports, please refer to the network diagram and the additional general required ports in the Veeam Backup & Replication User Guide or in the knowledge-base article KB1518 (<http://www.veeam.com/kb1518>).

Management zone

The two domain controllers are contacted by the Veeam Backup & Replication server and the Veeam Enterprise Manager server. Outside of the management zone, no server needs to connect to Active Directory services. All cloud gateways, WAN accelerators and Windows-based repositories will use local authentication only. This way, any security breach in these zones (especially the DMZ) will not expose Active Directory to any risk.

Additionally, this design will keep the management components of Veeam Cloud Connect isolated..

However, for better management, all servers will be registered in the DNS services running on the domain controllers. Even the servers using local authentication will be reachable using their hostname and the domain suffix, cloudconnect.local. For the same reason, the only connections to the domain controllers that is allowed will be toward the DNS servers over ports TCP/UDP 53.

DMZ zone

This security zone hosts the cloud gateways. These components are the only ones directly reachable via public internet connections. For the best protection, a service provider should isolate this zone from both public internet (allowing only the single TCP port needed for publishing the service) and the rest of the Veeam Cloud Connect infrastructure.

The cloud gateways need to communicate with the management zone for DNS resolution using the domain controllers (and for Active Directory operations if they were joined to the Cloud Connect internal domain). They also need to communicate with the Veeam backup server to operate the backup services and to the storage zone to allow communication between the data mover components at the customer site and the WAN accelerators and repositories at the service provider site.

Storage zone

This security zone hosts the data movers managing all the inbound and outbound data streams. Backup repositories are the foundation to create the logical cloud repositories customers use, while the (optional) WAN accelerator technology allows important bandwidth savings for those customers who have WAN accelerators on their own side.

Both components need to communicate with cloud gateways and through the cloud gateways to the customers. Additionally, they will communicate with the Veeam backup server and with the domain controllers (if the storage components have been joined to Active Directory or to simply use their DNS services).

Direct access should be limited to few authorized people, because an administrator can see all customers' backup files on the on the backup repositories. If those are not encrypted, unauthorized access to customers data is possible.

Subnets

Each security zone is isolated thanks to dedicated VLANs and firewalls that are the only entry points to and from each security zone to another, with rules limiting connections to the minimum required to operate a Veeam Cloud Connect environment.

The service provider uses one IPv4 subnet for each security zone. This allows you to write firewall rules per subnet more easily.

Security Zone	Subnet	VLAN	Gateway
WAN	185.62.37.96/28	6	185.62.37.97
DMZ	10.10.111.0/24	111	10.10.111.254
Management	10.10.51.0/24	51	10.10.51.254
Storage	10.10.110.0/24	110	10.10.110.254

All subnets have a gateway address; these IP addresses are configured and managed by one or more firewalls. This way, every communication between the security zones is filtered.

Notice that WAN and DMZ are two distinct subnets. Because of more advanced configurations required by replication services in Chapter 4, cloud gateways will have two network connections: one in the WAN subnet and the other in the DMZ subnet. Because they are reachable from the internet and thus vulnerable to attacks, their connection to the other subnets is blocked by a firewall that only allows the minimum required ports to the other services.


Veeam Cloud Connect Backup deployment

Once VLANs and subnets are created, and the firewalls are in place to protect communications between the different security zones, it's time to deploy the different servers needed and suggested to build the complete Veeam Cloud Connect Backup environment.

This is the schema of the servers:

Veeam Cloud Connect Backup

Servers detailed diagram



3.3: Cloud Connect Backup servers details

Active Directory

The internal domain is named **cloudconnect.local** and is managed by two domain controllers:

dc1	
server name	dc1.cloudconnect.local
IP Address	10.10.51.21
Operating System	Windows Server 2012 R2
Installed roles	AD, DNS, Global Catalog, FMSO roles
vCPU	2
RAM	4 Gb
Disk	40 Gb

dc2	
server name	dc2.cloudconnect.local
IP Address	10.10.51.22
Operating System	Windows Server 2012 R2
Installed roles	AD, DNS, Global Catalog
vCPU	2
RAM	4 Gb
Disk	40 Gb

Active Directory should use at least Windows Server 2008 level and be configured with no backward compatibility with older domain controllers. This allows you to reach an additional level of security can. If possible, use native Windows 2012 R2 Active Directory level.

10.10.51.21 and **10.10.51.22** are also the DNS servers to be configured in all other servers of the Veeam Cloud Connect infrastructure. For those servers that will use local authentication, DNS records should be configured manually.

Veeam management servers


There are two windows servers so you can separate Veeam Backup & Replication and Veeam Enterprise Manager.

EM	
server name	em.cloudconnect.local
IP Address	10.10.51.41
Operating System	Windows Server 2012 R2
Installed components	Veeam Enterprise Manager + Cloud Portal (used only in replication services)
vCPU	2
RAM	4 Gb
Disk	40 Gb

This server holds the installation of Veeam Enterprise Manager and its related database. By having a separated installation, a service provider can better manage the different performance requirements of Veeam Enterprise Manager and the Veeam Backup & Replication server and configure a specific security rule to allow access to the RESTful API service running on the Enterprise Manager from an optional web portal only.

The installation has no specific requirements, and you can follow the default wizard from start to finish. A dedicated Microsoft SQL Server 2012 Express is installed locally as part of the installation wizard, and Veeam Enterprise Manager itself will use it. If the service provider is also going to offer replication or DRaaS services, the optional cloud portal should be selected during the installation.

Once the installation of Veeam Backup & Replication is completed on **vbr.cloudconnect.local**, the configuration of Veeam Enterprise Manager can be completed by adding this server to the list of managed backup servers.




3.4: Add Veeam backup server to Veeam Enterprise Manager

VBR	
server name	vbr.cloudconnect.local
IP Address	10.10.51.40
Operating System	Windows Server 2012 R2
Installed components	Veeam Backup & Replication
vCPU	4
RAM	8 Gb
Disk	40 Gb

This server holds the installation of Veeam Backup & Replication. In a Veeam Cloud Connect infrastructure, this server is the central location for daily activities.

The installation has no specific requirements, and you can follow the default wizard from start to finish. A dedicated Microsoft SQL Server 2012 Express is installed locally as part of the installation wizard, and the Veeam Backup & Replication server itself will use it. During the component selection, a service provider should also choose to install the optional PowerShell SDK: Cloud Connect can be heavily automated via RESTful API or PowerShell, so having both available is a good choice.

Once the setup is completed and the license to enable Veeam Cloud Connect is installed (directly or pushed via Veeam Enterprise Manager), the initial management interface can be reached by opening the Veeam console and selecting the node Cloud Connect:



3.5: Veeam Cloud Connect start screen

From here, the required steps to have a fully functional Veeam Cloud Connect Backup infrastructure are:

1. Create and install a proper certificate (see Appendix A)
2. Deploy and configure the required cloud gateways
3. Deploy and configure the optional WAN accelerators
4. Deploy and configure at least one backup repository

Once all the configurations steps are completed, a service provider will be able to create and manage users/tenants.

Cloud gateways

A Veeam Cloud Connect infrastructure requires at least one cloud gateway, but as explained previously, multiple gateways are mandatory to deploy a reliable solution. In this scenario, you will deploy three cloud gateways, to satisfy a 2+1 redundancy: three gateways will be available to accept and manage incoming connections, and in case of a failure of one of them, there will always be two available gateways, thus guarantying load balancing and redundancy even in a degraded situation. Furthermore, the use of three gateways allows maintenance activities to any of the gateways (patching, hardware maintenance or upgrades, etc.) while always leaving two running gateways.

GTW1	
server name	gtw1.cloudconnect.local
IP Address	10.10.111.98
IP Address	185.62.37.98
Operating System	Windows Server 2012 R2
Installed components	Veeam Cloud Gateway
vCPU	2
RAM	2 Gb
Disk	40 Gb

GTW2	
server name	gtw2.cloudconnect.local
IP Address	10.10.111.99
IP Address	185.62.37.99
Operating System	Windows Server 2012 R2
Installed components	Veeam Cloud Gateway
vCPU	2
RAM	2 Gb
Disk	40 Gb

GTW3	
server name	gtw3.cloudconnect.local
IP Address	10.10.111.100
IP Address	185.62.37.100
Operating System	Windows Server 2012 R2
Installed components	Veeam Cloud Gateway
vCPU	2
RAM	2 Gb
Disk	40 Gb

For cloud gateway sizing, a service provider should follow these recommendations:

CPU: 2 vCPU or core can manage a bandwidth up to 10Gbit/s.

RAM: Around 512 KB of RAM are consumed per single connection. From a load perspective, it is suggested to limit a gateway to 1,000 connections by adding multiple instances when the total amount of connections goes above this value.


With 1,000 connections, the total memory requirement for the cloud gateway service is around 512 MB; the requirements of the underlying OS must be taken into consideration and added to this value, hence the 2 GB suggested value.

Cloud gateways networking

Cloud gateway networking is configured in a specific way in this guide. Service providers may decide to follow this example or to create a different configuration.


Each cloud gateway server has two network connections linked to the public WAN network and the DMZ. This way, the external interface can be configured with a public IP and be reachable directly from the tenants' side. The internal windows firewall of the machine and the external firewall protecting this subnet allow only connections towards TCP/UDP 6180, the default port of the Veeam Cloud Connect service.

As it can be observed in one of the gateways, the configuration of the network connection looks like this:



3.6: Cloud gateway WAN link configuration

Only TCP/IP v4 has been enabled; every other protocol and service available as a default in the Windows connection has been disabled. The WAN connection has the default gateway enabled, but no DNS configuration because this is configured in the DMZ connection:



3.7: Cloud gateway DMZ link configuration

There are some permanent routes configured on the Windows machine:

Persistent Routes:			
Network Address	Netmask	Gateway Address	Metric
0.0.0.0	0.0.0.0	185.62.37.97	Default
10.10.51.0	255.255.255.0	10.10.111.254	1
10.10.110.0	255.255.255.0	10.10.111.254	1

The first one is the default route. The other two are created to allow the cloud gateway to connect back to the management network (10.10.51.0/24) and to the storage network (10.10.110.0/24). To create these two rules, these commands can be executed in an elevated Windows command prompt:

```
route add 10.10.51.0 mask 255.255.255.0 10.10.111.254 -p
```

```
route add 10.10.110.0 mask 255.255.255.0 10.10.111.254 -p
```

10.10.111.254 is the IP address of the firewall connecting and segregating the DMZ subnet from the other subnets. Only the minimum amount of required connections are allowed from DMZ to management and storage:

NOTE: To make rule creation easier, some aliases have been created:

VCC_gateways: 10.10.111.98, 10.10.111.99, 10.10.111.100

Domain_controllers: 10.10.51.21, 10.10.51.22

VBR_Server: 10.10.51.40

WAN_accelerators: 10.10.110.11, 10.10.110.12

Linux_repositories: 10.10.110.51

Windows_repositories: 10.10.110.52

Proto	Source	Port	Destination	Port	Description
IPv4 TCP/ UDP	VCC_gateways	*	Domain_controllers	53 (DNS)	Allow gateways to use internal dns
IPv4 TCP	VCC_gateways	*	VBR_Server	6169	Gateways pass tenant VBR commands to SP VBR
IPv4 TCP	VBR_Server	*	VCC_gateways	6160	Veeam Installer from VBR to VCC gateways
IPv4 TCP	VBR_Server	*	VCC_gateways	6162	Veeam Transport from VBR to VCC gateways
IPv4 TCP	VBR_Server	*	VCC_gateways	6168	Cloud gateway listen for cloud commands from SP VBR
IPv4 TCP/ UDP	VBR_Server	*	VCC_gateways	137– 139	Veeam SMB share access from VBR to VCC gateways

The last rule can be disabled and enabled only when a new Veeam component needs to be installed or upgraded because Veeam uses SMB shares to deploy the installer packages into remote Windows servers.

DoS protection

The cloud gateway is directly exposed over internet. In order to be protected by DoS (Denial of Service) attacks trying to saturate all the available connections, this component has default limits on the amount of connections it can accept:

number of connections from the same IP address = 16

number of total connections = 256

To change these values, a service provider needs to create two new DWORD registry keys in each cloud gateway in:

HKEY_LOCAL_MACHINE\SOFTWARE\Wow6432Node\Veeam\Veeam Gate Service and configure them as follows (and restart the service to apply the new numbers):

PeerCloudConnectionsLimit (per IP, default is 16)


MaxSimultaneousCloudConnections (total, default is 256)

Remember that a cloud gateway is a failure domain when evaluating the impact on connections caused by its loss. One thousand connections lost on a failed cloud gateway will impact several customers. A service provider should carefully evaluate this scenario and deploy multiple cloud gateways to spread the connections over a larger number of smaller failure domains.

Installation

Once the three cloud gateways are added to the backup infrastructure as managed Windows servers, the service provider will deploy on each of them the cloud gateway component. The procedure is quick and easy, and should be repeated for all the three gateways.


1. From the Cloud Connect node, go to Cloud Gateways and select Add Cloud Gateway
2. Select one of the previously added servers:



The screenshot shows the 'New Cloud Gateway' wizard in the Veeam Cloud Connect console. The 'Name' step is active, indicated by a green plus icon and a cloud icon. The wizard title is 'New Cloud Gateway'. Below the title, there is a message: 'Choose a Microsoft Windows server to set up cloud gateway service on. We recommend that you set up multiple cloud gateways for high availability and automatic load balancing.' On the left, there is a sidebar with tabs: 'Name', 'Networking', 'Review', 'Apply', and 'Summary'. The 'Name' tab is selected. The main area contains a 'Choose server:' dropdown menu with 'gtw1.cloudconnect.local' selected and an 'Add New...' button. Below this is a 'Description:' text box. Further down is an 'External port:' spinner box set to '6180'. A note below the port says: 'TCP/UDP port for external connections. All traffic between you and your users will go through this port. Your users will need to specify this port when establishing the initial connection to your service.' At the bottom, there are navigation buttons: '< Previous', 'Next >', 'Finish', and 'Cancel'.

3.8: Add a new cloud gateway

1. Configure the desired networking mode:



The screenshot shows the 'New Cloud Gateway' wizard in the Veeam Cloud Connect console, now on the 'Networking' step. The title is 'New Cloud Gateway'. Below the title, there is a message: 'Specify how this server is connected to the internet.' On the left, the sidebar has tabs: 'Name', 'Networking', 'Review', 'Apply', and 'Summary'. The 'Networking' tab is selected. The main area contains two radio button options. The first option, 'This server is connected directly to the internet', is selected. Below it is a 'Select network adapter (NIC) connected to the internet:' dropdown menu with '185.62.37.98 (Ethernet1)' selected. The second option, 'This server is located behind NAT', is unselected. Below it is a note: 'Type in the external DNS name or IP address of the NAT device. Port forwarding rule must be set up on the NAT to forward network traffic from the external port to the internal port specified below.' There are two input fields: 'DNS name or IP address:' and 'Internal port:' (a spinner box set to '6180'). At the bottom, there are navigation buttons: '< Previous', 'Next >', 'Finish', and 'Cancel'.

3.9: Select the desired networking mode for the cloud gateway

This guide suggests and explains the direct move. Direct mode is available to directly expose a cloud gateway over the Internet with a public IP address configured on the gateway machine itself. Veeam Cloud Connect fully supports both deployment modes, and service providers should properly protect the cloud gateways behind a firewall, regardless which mode is used.

When configuring the cloud gateways in NAT mode, the wizard needs to be filled with the expected DNS name that will be used to connect to the gateway itself. Following the example, the mappings would be:

HOST	INTERNAL IP	DNS HOST	NAT IP
gtw1.cloudconnect.local	10.10.111.98	gtw1.virtualtothecore.com	185.62.37.98
gtw2.cloudconnect.local	10.10.111.99	gtw2.virtualtothecore.com	185.62.37.99
gtw3.cloudconnect.local	10.10.111.100	gtw3.virtualtothecore.com	185.62.37.100

Once the DNS A (host) records are configured with all the public IP addresses in order to enable round robin, the internet-facing part of Veeam Cloud Connect is ready.

WAN accelerators

WAN1	
server name	wan1.cloudconnect.local
IP Address	10.10.110.11
Operating System	Windows Server 2012 R2
Installed components	Veeam Cloud Gateway
vCPU	4
RAM	8 Gb
Disk	40 Gb, OS disk
Disk	200 Gb, cache disk

WAN2	
server name	wan2.cloudconnect.local
IP Address	10.10.110.12
Operating System	Windows Server 2012 R2
Installed components	Veeam Cloud Gateway
vCPU	4
RAM	8 Gb
Disk	40 Gb, OS disk
Disk	200 Gb, cache disk

As explained previously, WAN accelerators are optional components, but any service provider should deploy them. Veeam Cloud Connect licensing enables the use of WAN accelerators at no additional cost for service providers, and most of all, their presence allow a service provider to offer a complete solution to those customers owning Veeam licenses with WAN acceleration enabled.

The 200 GB size for the cache disk is mainly a suggested starting point. One hundred gigabytes are assigned to the general cache, plus an additional 100 GB are allocated for each job cache requirements. Depending on the amount of customers assigned to a specific WAN accelerator and thus the total amount of managed data, the cache should be then increased to guarantee optimal performance to all customers connecting to a given WAN accelerator. Please refer to Veeam User Guide and Best Practices to learn how to properly size the WAN accelerator cache.

Finally, the use of multiple WAN accelerators is a good design solution in terms of High Availability (HA). Even if only one WAN accelerator can be assigned to a given customer, the presence of additional servers eventually allow to quickly reconfigure all customers linked to a failed WAN accelerator to use another one.

WAN accelerator networking


Additional firewall rules are required to correctly deploy and manage WAN accelerators:

Proto	Source	Port	Destination	Port	Description
IPv4 TCP/ UDP	WAN_ accelerators	*	Domain_controllers	53 (DNS)	Allow accelerator to use internal dns
IPv4 TCP	VBR_Server	*	WAN_accelerators	6160	Veeam Installer from VBR t WAN accelerator
IPv4 TCP	VBR_Server	*	WAN_accelerators	6162	Veeam Transport from VBR t WAN accelerator
IPv4 TCP	VBR_Server	*	WAN_accelerators	6164	Veeam WAN Control fro VBR to WAN accelerator
IPv4 TCP	VCC_gateways	*	WAN_accelerators	6165	Gateways transfer data to WAN accelerator
Pv4 TCP	VCC_gateways	*	WAN_accelerators	2500– 5000	Gateways transfer data to WAN accelerator
Pv4 TCP	VBR_Server	*	WAN_accelerators	2500– 5000	VBR transfers data to WAN accelerator
Pv4 TCP	VBR_Server	*	WAN_accelerators	49152– 65535	Veeam RP from VBR t WAN accelerator
IPv4 TCP/ UDP	VBR_Server	*	WAN_accelerators	137– 139	Veeam SMB share access fro VBR to WAN accelerator

You can disable the last rule and enable it only when a new Veeam component needs to be installed or upgraded because Veeam uses SMB shares to deploy the installer packages into remote Windows servers like WAN accelerators and Windows-based repositories.

WAN accelerator deployment


The deployment and configuration of a WAN accelerator component is a simple and quick process. When starting the wizard to deploy a new WAN accelerator, select the corresponding Windows server and accept the default values. If needed, you can increase the number of streams a second time to increase the utilization of the WAN accelerator:



The screenshot shows the 'New WAN Accelerator' wizard with the 'Server' step selected. The left sidebar contains 'Server', 'Cache', 'Review', 'Apply', and 'Summary'. The main area has a 'Choose server:' dropdown menu with 'wan1.cloudconnect.local' selected and an 'Add New...' button. Below is a 'Description:' text box. Further down, 'Traffic port' is set to 6165 and 'Streams' is set to 5. A note states: 'Using multiple upload streams helps to fully saturate WAN links.' At the bottom are '< Previous', 'Next >', 'Finish', and 'Cancel' buttons.

3.10: Choose a server to install WAN accelerator

The cache is placed in a dedicated disk, so any disk consumption problem will not affect the OS partition:



The screenshot shows the 'New WAN Accelerator' wizard with the 'Cache' step selected. The left sidebar contains 'Server', 'Cache', 'Review', 'Apply', and 'Summary'. The main area has a 'Folder:' text box with 'E:\VeeamWAN' and a 'Browse...' button. Below is a table showing disk paths, capacity, and free space:

Path	Capacity	Free
C:\	39.7 GB	29.3 GB
E:\	199.9 GB	199.8 GB

Below the table, 'Cache size' is set to 100 GB. A note states: 'We recommend at least 10GB per each operating system used in the environment. Larger cache improves data reduction ratio, but requires faster storage.' At the bottom are '< Previous', 'Next >', 'Finish', and 'Cancel' buttons.

3.11: Configure WAN accelerator cache

After the configuration of the two WAN accelerators is completed, they will be both listed in the corresponding section of Veeam Backup & Replication and be ready to use.

Cloud repositories

As suggested previously, it's preferred to use a Windows or Linux server as a backup repository so a proper Veeam data mover service can be deployed on the repository machine itself. With this service properly deployed and running, all read and write operations are delegated to this service and all the available compute resources can be used by the data mover deployed locally on the backup repository.

SMB shares are completely supported by Veeam Cloud Connect, but even in this scenario, it's advisable to deploy a dedicated Windows machine that will act as the gateway server (not to be confused with a cloud gateway) to directly communicate with the SMB share. The data mover will be deployed on this machine and not on other systems, especially the Veeam backup server, which should be only used as a management console:

The screenshot shows the 'New Backup Repository' wizard with the 'Share' tab selected. The 'Name' field is empty, and the 'Type' is set to 'Share'. The 'Shared folder' is set to '\\smb\share'. The 'This share requires access credentials' checkbox is unchecked. The 'Gateway server' is set to 'vbr.cloudconnect.local (Backup server)'. The 'Finish' button is highlighted.

3.12: Select a dedicated gateway server for SMB shares

This server will ultimately act like a proper repository machine.

There are several design choices for a backup repository, and to list them all here will be simply impossible because many will be surely left out. Instead, this guide will describe two options: a Windows Server and a Linux Server, both with local storage. This is not intended to suggest these are the best storage solutions; it is mostly an example to better describe the process of adding a backup repository to the Veeam Cloud Connect infrastructure.

Linux Backup Repository

repo1	
server name	repo1.cloudconnect.local
IP Address	10.10.110.51
Operating System	CentOS Linux 7.0
Installed components	none
vCPU	4
RAM	8 Gb
Disk	20 Gb, OS disk
Disk	1 Tb, Backups disk


Networking is easy to manage for a Linux repository because only a control port is required:

Proto	Source	Port	Destination	Port	Description
IPv4 TCP/ UDP	Linux_ repositories	*	Domain_controllers	53 (DNS)	Allow repositories to use internal dns
IPv4 TCP	VBR_Server	*	Linux_repositories	22	Veeam VBR connects to Linux repositories
IPv4 TCP	VCC_gateways	*	Linux_repositories	2500– 5000	Gateways transfer data to Linux repositories
IPv4 TCP	VBR_Server	*	Linux_repositories	2500– 5000	VBR transfers data to Linux repositories

The entire management of a Linux repository is done using SSH. Because of this connection, the Veeam Backup & Replication server deploys a runtime component every time that starts the Veeam Data Mover using Perl (the only other requirement together with SSH), then transfers data using the usual Veeam ports 2500 to 5000.


Once the Linux server is added among the managed servers, the configuration of a new repository follows the usual process.

NOTE: Starting from CentOS 7.0, together with SSH and Perl you need to install manually also Perl-Data-Dumper. You can do so by running `yum install Perl-Data-Dumper`



3.13: Select the path to be used as a backup repository

Managing the ingestion rate of the repository is an important configuration aspect of the repository:



3.14: Configure path and load control

A typical Veeam Cloud Connect customer will be limited by the upload bandwidth that is available; this will be the main bottleneck in most of the use cases. However, this does not mean it will be the primary bottleneck for the service provider. Because the service provider is accepting several concurrent connections, the number of concurrent tasks connecting to the repository could be notable.

For this reason, a service provider needs to check the performance of a given storage solution, and configure the limits for concurrent tasks and/or data rate accordingly beforehand. On the other hand, a service provider needs to have room for enough concurrent connections so that customers do not end up waiting for available resources for their jobs.

Because vPower NFS is not supported to date in Veeam Cloud Connect, a service provider can safely disable the configuration of this component during the repository creation wizard and complete it.

Windows backup repository


repo2	
server name	repo2.cloudconnect.local
IP Address	10.10.110.52
Operating System	Windows Server 2012 R2
Installed components	Veeam Backup Repository
vCPU	4
RAM	8 Gb
Disk	40 Gb, OS disk
Disk	1 Tb, Backups disk

Networking requires a few more ports in order to manage a Windows repository, compared to the previous Linux repository:

Proto	Source	Port	Destination	Port	Description
IPv4 TCP/UDP	Windows_repositories	*	Domain_controllers	53 (DNS)	Allow repos to us intern
IPv4 TCP	VBR_Server	*	Windows_repositories	6160	Veeam Instal from Wind repos
IPv4 TCP	VBR_Server	*	Windows_repositories	6162	Veeam Trans from Wind repos
IPv4 TCP	VCC_gateways	*	Windows_repositories	2500–5000	Gate transf data t Wind repos
IPv4 TCP	VBR_Server	*	Windows_repositories	2500–5000	VBR transf data t Wind repos
Pv4 TCP	VBR_Server	*	Windows_repositories	49152–65535	Veeam from Wind repos
IPv4 TCP/UDP	VBR_Server	*	Windows_repositories	137–139	Veeam SMB acce VBR Wind repos


You can disable last rule and enable it only when a new Veeam component needs to be installed or upgraded because Veeam uses SMB shares to deploy the installer packages into remote Windows servers like WAN accelerators and Windows-based repositories.

Once the Windows server is added among the managed servers, the configuration of a new repository follows the usual process.




3.15: Select the path to use as a backup repository

Again, as in the Linux repository, managing the ingestion rate of the repository is an important configuration aspect of the repository:



3.16: Configure path and load control

Finally, because vPower NFS is not supported to date in Veeam Cloud Connect, a service provider can safely disable the configuration of this component during the repository creation wizard and complete it.




3.17: Disable vPower NFS

Once the backup repositories are deployed and configured, Veeam Cloud Connect is ready to consume for backup services.

Reference design for replication services

This chapter describes a complete Veeam Cloud Connect Replication deployment at a service provider. By impersonating a provider willing to offer DRaaS, you will design and deploy all necessary servers, networks and network rules in order to run Veeam Cloud Connect Replication.



4.1: Veeam Cloud Connect Replication overview

The virtualized environment

In order to offer replica resources for a DRaaS solution, a service provider needs to build and connect to Veeam Cloud Connect a virtualized environment, based on one of the supported hypervisor technologies: VMware vSphere or Microsoft Hyper-V. Veeam doesn't perform any conversion between the two supported platforms, so it can be a good choice for a service provider willing to onboard more customers to have both the platforms available for his customers.

NOTE: In this book, replica resources will be offered using only a VMware vSphere platform. Even if many concepts can be applied also to Microsoft Hyper-V, please refer to the official Veeam User Guide for more information.

vSphere environment


To serve customers, the service provider builds a VMware vSphere environment. There are so many possible configurations and options, that it would be impossible to suggest the best one. Every VMware architect can design and deploy a working vSphere solution that is in line with the business requirements of the service provider he works for.

Few notes however can be listed as suggestions:

- Even if Veeam Cloud Connect can work with single hosts, the correct choice for a true DR environment should be a proper cluster, with all the necessary redundancy in place (shared storage, HA, DRS). This is to avoid that any problem to a single host can interrupt the entire service.
- Veeam Cloud Connect uses Veeam replication technology to create replica VMs in the DR environment. What's important in order to guarantee a complete compatibility between the tenant's production environment and the Veeam Cloud Connect platform is not the version of vSphere, but rather the virtual hardware of the replicated VMs. As long as the version of vSphere the service provider uses can accept VMs created in the tenant's vSphere environment, any combination of vSphere versions between the two sites is accepted. Obviously, in order to guarantee the best compatibility, a service provider should use the latest version of vSphere. As of today, Veeam Cloud Connect supports vSphere 6.0.x.
- If available, use distributed switches. Veeam Cloud Connect supports both standard and distributed switches, and Veeam Cloud Connect configures new port groups automatically in both scenarios. However, distributed switches are easier to manage on large infrastructures.


Service account

Veeam Backup & Replication needs to connect to the vSphere environment in order to create new VM's, port groups, and other activities. During the setup of the vSphere environment in the Veeam console:



4.1: Add VMware vCenter server

a credential to connect and operate on vCenter is required. A dedicated account for Veeam Cloud Connect should be created in the vSphere environment:



4.2: Add a new user in vCenter Server

Enter values for this user, including the password.

User name: cloudconnect

Password: *****

Confirm password: *****

First name:

Last name:

Email address:


Description: User dedicated to connections coming from Veeam Cloud Connect

OK Cancel

4.3: Create a dedicated account for Veeam Cloud Connect in vCenter server

With this dedicated account, it is easier to filter logs and identify activities belonging to Veeam Cloud Connect in vCenter, so problem solving should be easier. Additionally, it is extremely easy to disable any access of Veeam Cloud Connect to the vSphere environment, if needed, by simply disabling this user.


Once the user is created and is added to the vCenter Administrators group, the user can be configured in Veeam when registering the vCenter server itself:



The screenshot shows the 'New VMware Server' wizard with the 'Credentials' tab selected. The left sidebar has 'Name', 'Credentials', 'SSH Connection', and 'Summary'. The main area contains instructions to select an account with local administrator privileges. A dropdown menu shows 'cloudconnect@vsphere.local' with an 'Add...' button. Below it is a 'Port' field set to '443'. At the bottom are '< Previous', 'Next >', 'Finish', and 'Cancel' buttons.

4.4: Specify dedicated user account to connect to vCenter Server

The registration process is correctly completed, and vCenter is added to the managed platform:



The screenshot shows the 'New VMware Server' wizard with the 'Summary' tab selected. The left sidebar has 'Name', 'Credentials', and 'Summary'. The main area displays a summary of the configuration: 'VMware vCenter Server 'vcsa.cloudconnect.local' was successfully created.', 'Host info: VMware vCenter Server 6.0.0 build-3634794', and 'Connection options: User: cloudconnect@vsphere.local, Port: 443'. At the bottom are '< Previous', 'Next >', 'Finish', and 'Cancel' buttons.

4.5: vCenter is added to Veeam Backup & Replication

Target Veeam proxies

Veeam Cloud Connect leverages Veeam replication technology to create replica VMs at the DR site. For this reason, a Veeam Cloud Connect environment offering replica resources to tenants not only needs to have a virtualized environment, but also needs one or more Veeam proxies to be used as targets of the replica process.

Outside of the general rules for Veeam replication best practices and the information in Chapter 2, there are some additional notes in regards to target proxies deployed in Veeam Cloud Connect:

- Among the possible transport methods, hotadd is probably the best suited for Veeam Cloud Connect. For this reason, one of the best choices to deploy target proxies is to have them as VMs running on top of the service provider's virtualized environment.
- The amount of necessary proxies vary based on the amount of concurrent replicas the service provider is going to receive. Please refer to Chapter 2 for concurrency considerations.
- Regardless of a given target proxy's processing slots, a service provider should deploy multiple proxies also to guarantee operation continuity should one proxy fail.


Networking

Outside of the virtualized platform, there's a complete range of additional hardware resources that a service provider has to setup in order to guarantee Veeam Cloud Connect operates properly.

In particular, two components are of interest for Veeam Cloud Connect.

VLANs

Veeam Cloud Connect guarantees multi-tenancy at the network level using VLANs. Each port group created for any service provider is tagged with a unique VLAN ID so communications between port groups, VLANs and networks are only possible by traversing a network extension appliance and thus are regulated. Veeam Backup & Replication, upon creating a new network for a tenant, automatically creates a new port group on the selected virtual switch and sets a VLAN tag. The available VLAN IDs need to be configured in advance both in Veeam Cloud Connect:



4.6: Configure VLAN pool settings

In the example, the service provider has a vSphere cluster named VCCR Cluster. There is a distributed virtual switch named vcc-dvs, and VLANs 112 to 149 have been already configured in the hardware switches on the uplinks registered in the distributed switch. No routing has been configured between the VLANs so this property will be exclusively managed by Veeam network extension appliances.


Two settings control the way NAT is applied to VMs belonging to a given VLAN:

- **"No Internet"/"With Internet"** toggles Source NAT. VMs belonging to a VLAN "With Internet" can reach internet.
- **"Public IP"** setting enables Destination NAT. Both VMs belonging to "No Internet" and "With Internet" VLANs can be reached from internet if a Public IP publishing rule is created.

These two settings are independent from each other, so that service providers can satisfy different customers' needs.

Public IPs

During a full failover, replicated VMs are powered on and need to be reached from the outside so that a tenant can consume their services. All communications happening in a failover are managed by Veeam network extension appliances; in order to be reached from internet and to allow failed over VMs to reach the internet, an NEA acts like a firewall or gateway. To do so it needs to have a public IP loaded on its external interface. Just like VLAN pools, public IPs are not manually assigned to each tenant; instead, whenever an additional public IP is needed, Veeam Cloud Connect uses one of the available IPs that have been pre-loaded into the configuration:



4.7: Public IP pre-loaded in Veeam Cloud Connect

Service providers can consume public IPs coming from different pools, as long as the external interface of an NEA is connected to the right VLAN where this public IP can be used.


Public IP's or NAT-ed IP's

When Veeam Cloud Connect v8 was first released, only backup services were available. For this solution to work, cloud gateways could have been published either loading public IPs directly in their network interfaces or using NAT (network address translation) technologies, cloud gateways were using non-public IP addresses, and an external component like a firewall would have published the TCP ports needed to expose cloud gateways to the public internet.

This is still the case for backup services in Veeam Cloud Connect v9, but in order to correctly publish replica resources, it is highly recommended to load public IPs on both cloud gateways and NEAs.

There are two main reasons to do so:

- Cloud gateways and NEAs need to be able to communicate during a partial failover. If one of the two is behind a NAT system, the NAT device itself hides the real IP of the device. However, because DNS resolution is also in place, a service provider needs to build a complicated solution to hack DNS in order to correctly resolve Veeam Cloud Connect resources via their NAT-ed IPs, instead of the real ones. Simply using public IPs with no NAT makes this configuration much easier.
- Tenants can execute a full failover by themselves by accessing the Veeam Cloud Connect Portal. When at least one public IP address mapping rule is created, a service running in a VM is published on the outside using one of the public IPs assigned on the external interface of the NEA. Actually, any IP would be usable in this situation, like in this example:



4.8: A public IP address mapping rule using a non-public IP address

The public IP address 198.51.100.3 is a special subnet dedicated to create documentation and examples, but it can be compared with an internal IP used in a DMZ network, like 192.168.0.3. None of these IPs is reachable over internet. If a tenant attempts a failover using this configuration at the service provider, the service provider itself will need an additional NAT rule to publish the internal IP used in the NEA using a real public IP, like 185.62.37.102. The problem is that the customer can be confused because the Veeam Cloud Connect Portal suggests the user to reach his virtual machine over the IP address loaded in the NEA:

Cloud Connect Portal

You logged in as: **veeamon** | [Sign out](#)

SESSIONS HISTORY

NAME	STATUS	CREATED ↓	FINISHED
✓ ecommerce Validating VM Processing failover to restore point 4 days ago (3:30 PM Thursday 11/5/2015) for VM ecommerce Reverting VM to the restore point Powering on VM Processing full site failover settings Initialized network 192.168.11.0/24 (veeamon network 1 (internet)), internet access is enabled Enabled public endpoint 192.168.100.51:80 for service 192.168.11.112:80 () Failover to restore point 4 days ago (3:30 PM Thursday 11/5/2015) executed successfully for VM ecommerce	✓	11/9/2015 05:12:3...	11/9/2015 05:14:1...
✓ E-Commerce DR to VCC Job started at 11/9/2015 5:12:35 PM Building VM list. Setting up network extension for tenant veeamon, routing between networks is disabled Processing: ecommerce Failover plan executed, 1 VMs processed. Successes: 1, Warnings: 0, Errors: 0. Job finished at 11/9/2015 5:14:19 PM	✓	11/9/2015 05:12:3...	11/9/2015 05:14:1...

4.9: Failover plan using a non-public IP

This IP address (192.168.100.51) loaded in the NEA is not a public IP, and so the customer cannot reach his VM unless the service provider advises the tenant to use the public IP that for example a firewall is using to publish the NEA on the internet.

This scenario is a major complication in both the network design and the level of automation that can be possibly reached. The NEA is a hardened Linux system, exposing to the outside just the ports that the customer or tenant has configured in his failover plan. No additional protection or routing is needed on the external interface of the NEA. If a service provider desires to have additional control, he can deploy an external firewall working at L2 (Layer2) and thus be totally transparent to the network configuration of a NEA or use the firewall to also be the router of the NEA's public IP.

Veeam Cloud Connect Replication deployment

A service provider can decide to offer only replication services or to extend existing backup services by adding replication. This is the scenario of a service provider upgrading Veeam Cloud Connect from v8 to v9, and it's the scenario we will show in this guide.


For this reason, the starting point of the architecture of Veeam Cloud Connect replication is the one described in Chapter 3. Here, we will add the additional components needed to realize the replication service.

Again, this book will start from the design of the overall infrastructure. This is paramount in order to better understand the relationships between the different components of Veeam Cloud Connect, network ports and services, and how they communicate between each other.

So, the first activity is the creation of the following network diagram. This diagram is specifically about Veeam Cloud Connect Replication, but it shares different components with backup services.

Veeam Cloud Connect Replication

Service and Network diagram




4.10: Veeam Cloud Connect Replication diagram

Based on the network diagram, deploy the additional components needed to create a Veeam Cloud Connect Replication service:

Veeam Cloud Connect Replication

Service and Network diagram



4.11: Veeam Cloud Connect Replication servers diagram

Following the initial deployment of Veeam Cloud Connect Backup, you now have to add different components: the vSphere environment and the new proxy servers.

vSphere environment

To receive replicas of VM, deploy a new and dedicated vSphere 6.0 environment, built with several ESXi nodes and a vCenter appliance.


ESX1	
server name	esx1.cloudconnect.local
IP Address	10.10.51.11
Operating System	VMware ESXi 6.0
CPU	12
RAM	128 Gb

ESX2	
server name	esx2.cloudconnect.local
IP Address	10.10.51.12
Operating System	VMware ESXi 6.0
CPU	12
RAM	128 Gb

ESX3	
server name	esx3.cloudconnect.local
IP Address	10.10.51.13
Operating System	VMware ESXi 6.0
CPU	12
RAM	128 Gb

ESX4	
server name	esx4.cloudconnect.local
IP Address	10.10.51.14
Operating System	VMware ESXi 6.0
CPU	12
RAM	128 Gb

The four nodes are grouped into a vSphere cluster where a shared storage is available and visible to all nodes. Also, HA, vMotion and DRS are enabled, so that a failure in one of the nodes doesn't interrupt the cluster itself and the replication services can continue.




4.12: The vSphere cluster

The cluster is managed by a vCenter appliance:

VCSA	
server name	vcsa.cloudconnect.local
IP Address	10.10.51.30
Operating System	SUSE Linux Enterprise 11
vCPU	4
RAM	12 Gb

Finally, the networking part: In order to manage networking on the virtualized environment better, a distributed switch has been created:



4.13: Networking in the vSphere environment

Each ESXi host has 2 * 10 Gb uplinks, connected to the physical switches where the different VLANs are terminated. There are some notable port groups, tagged with VLAN IDs:

- **Management (id 51)**: This is the management network where vCenter, Veeam Backup & Replication and other management machines are deployed. The network is 10.10.51.0/24, vlan id is 51.
- **vcc_public (id 6)**: This is the network where the public IPs are published. Here there are the three external interfaces of the cloud gateways, and here the external interfaces of the NEAs will be connected.

Any additional port group assigned to a tenant will be created directly over this distributed switch, and a unique VLAN ID will be assigned to it.

Veeam proxies

In order to receive replication data from a tenant, at least one proxy is needed. This proxy needs to be able to talk with its controlling VBR server, the vCenter server and all the ESXi hosts. To increase the availability of the service, deploy two proxy servers. Any service provider should consider carefully how many proxies are necessary based on the specific design of the environment.

PROXY1	
server name	proxy1.cloudconnect.local
IP Address	10.10.110.101
Operating System	Windows Server 2012 R2
Installed components	Veeam Proxy
vCPU	4
RAM	4 Gb
Disk	40 Gb

PROXY2	
server name	proxy2.cloudconnect.local
IP Address	10.10.110.102
Operating System	Windows Server 2012 R2
Installed components	Veeam Proxy
vCPU	4
RAM	4 Gb
Disk	40 Gb


Proxies need to communicate with the different components: ESXi hosts and vCenter, Veeam Backup & Replication server, cloud gateways, WAN accelerators. For this reason, you should add firewall rules between the different networks.

NOTE: Some additional aliases have been added to the central firewall: Proxies: 10.10.110.101, 10.10.110.102

Proto	Source	Port	Destination	Port	Description
IPv4 TCP/ UDP	Proxies	*	Domain_controllers	53 (DNS)	Allow accelerators to use internal dns
IPv4 TCP	VBR_Server	*	Proxies	6160	Veeam Installer from VBR to Proxies
IPv4 TCP	VBR_Server	*	Proxies	6162	Veeam Transport from BR to Proxies
IPv4 TCP	VCC_gateways	*	Proxies	2500– 5000	Gateways transfer data to WAN accelerators
IPv4 TCP	VBR_Server	*	Proxies	2500– 5000	VBR transfers data to Proxies
Pv4 TCP	VBR_Server	*	Proxies	49152– 65535	Veeam RPC from VBR to Proxies
IPv4 TCP/ UDP	VBR_Server	*	WAN_accelerators	137– 139	Veeam SMB share access from VBR to WAN accelerators

You can disable the last rule and enable it only when a new Veeam component needs to be installed or upgraded because Veeam uses SMB shares to deploy the installer packages into remote Windows servers.

Once all the different firewall rules are in place, service providers can deploy the proxy component on the different proxy servers:




4.14: Install a new Veeam proxy

You should leave all the configuration parameters as the default ones; transport mode specifically should be left as automatic because the proxies are VMs and they will use hotadd mode, which is the preferred mode for a target proxy. However, leaving automatic selection on allows for the usage of network mode as a failover option should something not work for the hotadd mode.

Traffic rules are also left empty, as any bandwidth management is done directly by Veeam Cloud Connect when configuring a tenant.

One last step must be done once the different Veeam proxies have been deployed: By default, the Veeam Backup & Replication server itself is also configured as a proxy.



The screenshot shows the Veeam Backup & Replication console interface. At the top, there is a toolbar with icons for 'Add Proxy', 'Edit Proxy', 'Disable Proxy', 'Remove Proxy', and 'Upgrade Proxy'. Below the toolbar, the left sidebar shows the 'BACKUP INFRASTRUCTURE' tree with 'Backup Proxies' selected. The main pane displays a table of available proxies.

NAME	TYPE	HOST	DESCRIPTION
proxy1.cloudconnect.local	VMware	proxy1.cloudconnect.local	
proxy2.cloudconnect.local	VMware	proxy2.cloudconnect.local	
VMware Backup Proxy	VMware	vbr.cloudconnect.local	Created by Veeam Backup & Replication

4.15: List of available Veeam proxies

To guarantee that replication traffic follows the designed path from cloud gateways to WAN accelerators and proxies, you have to disable the default proxy role installed in Veeam Backup & Replication server or even choose to remove the role completely.

Cloud Connect operations

Once all the different components have been installed and the network is correctly configured, it's time to start using Veeam® Cloud Connect.

In this chapter, our fictitious service provider will configure Veeam Cloud Connect, create tenants, and start offering Cloud Connect services to his users.

Initial setup

Once the entire infrastructure needed for Veeam Cloud Connect is deployed and configured, it will be time for the service provider to configure the Cloud Connect environment itself.

License

The first operation the service provider has to complete is to obtain and install the correct license to enable Cloud Connect. This step is required to be able to deploy the different components of Cloud Connect, so it should be the first activity when deploying Cloud Connect.

Once the service provider obtains a license, the license file needs to be loaded into Veeam Backup & Replication™ or, if it's installed, load it into the Enterprise Manager, as this component pushes licenses to all the registered Veeam Backup & Replication servers. To do so, select "Configuration" from the main page of Enterprise Manager, then "Licensing," and finally the "Change License" button. Once you have loaded the correct license file, the license screen in the Veeam Backup & Replication server should look like this:



The screenshot shows a 'License Information' dialog box with a green header bar. The dialog contains the following fields and values:

Field	Value
Status	Valid (90 days until expiration)
License type	Rental
Licensed to	Veeam Software
Contact person	Luca Dell'Oca
Edition	Enterprise Plus
Cloud Connect Provider	Yes
Support ID	00000000
Support expiration date	6/20/2016 (90 days left)
Protected VMs	0
Cloud Connect (Backup)	100 (0 used)
Cloud Connect (Replica)	100 (0 used)

At the bottom of the dialog, there is a checkbox labeled 'Update license key automatically' which is currently unchecked. To its right is an 'Update now' button. Below these are three buttons: 'Licensed Hosts', 'Install License', and 'OK'.

5.1: Veeam Cloud Connect license

Important: *The service provider needs to have the proper licenses for Veeam Cloud Connect VMs to offer Backup as a Service (BaaS), and Veeam Cloud Connect Replication VMs to offer Disaster Recovery as a Service (DRaaS). This book will not go into details about licensing and the auto-update features, so please refer to the User Guide for additional information.*

Certificates

After the proper license is in place, the following configuration step is related to Certificates. Veeam Cloud Connect is reached by end users via public internet, over a single TCP/UDP connection where the endpoint is one of the cloud gateways.

The connection is protected with SSL certificates, thus the service provider has to obtain and use those certificates so that the end users can verify the identity of the provider they are connecting to. Self-signed certificates are intended primarily for test purposes, even if Veeam Cloud Connect offers that option. In a production environment, a service provider should acquire and use a proper certificate, validated by one of the Certification Authorities recognized by Windows OS. In this way, no alert or warning is raised upon connection, and the trust between the service provider and his or her users is optimal.

To learn how to create and load SSL Certificate, please refer to APPENDIX A: SSL Certificates generation.

Create hardware plans

Parameters of backup resources are defined directly in the tenant options; however, for replica resources Veeam Cloud Connect uses a different solution.

Hardware plans can be compared to mobile phone plans. Each plan is defined with a series of options, and each customer/tenant subscribed to any given plan will receive that amount of resources. Every time a hardware plan is modified, each and every tenant subscribed to the same hardware plan will be affected.

Even if hardware plans can be assigned to multiple tenants, we suggest service providers to create dedicated hardware plans for each tenant. This way, whenever a single tenant requests a change in his hardware plan, the change can be applied to his dedicated plan and no other customer is affected.

A hardware plan comprises the following resources in the service provider virtualization infrastructure:

- **CPU** — The maximum amount of CPU that can be used by all replicated VMs of a tenant subscribed to a hardware plan when powered on at the service provider;
- **Memory** — The maximum amount of RAM that can be used by all replicated VMs of a tenant subscribed to a hardware plan when powered on at the service provider;
- **Storage** — A quota on a datastore (for VMware hardware plans) or a volume (for Hyper-V hardware plans) that a tenant can utilize for storing replicated VMs;
- **Network** — The specified number of networks to which a tenant's VM replicas can connect. When the service provider subscribes a tenant to a hardware plan, Veeam Backup & Replication creates the corresponding number of network adapters (vNICs) on the network extension appliance that is deployed on the service provider side.

A hardware plan is created in advance in the dedicated section of Cloud Connect:

The screenshot shows the 'New VMware Hardware Plan' dialog box with the 'Name' tab selected. The dialog has a sidebar with tabs: Name, Host, Storage, Network, Apply, and Summary. The 'Name' tab contains the following fields:

- Name:** A text box containing 'HWP-tenant1'.
- Description:** A text box containing a single character 'l'.

At the bottom of the dialog are four buttons: '< Previous', 'Next >', 'Finish', and 'Cancel'.

5.2: Create a new VMware hardware plan

We suggest service providers choose and adhere to a naming policy for their hardware plan. It is problematic to figure out which hardware plan is assigned to which tenant as the number of tenants becomes significant. By using a dedicated naming policy (in our example is HWP- tenant_account_ name), it can be easier to filter hardware plans and identify the needed one.

Then, we select the vSphere cluster where the hardware plan will be created, and assign CPU and Memory limits:


The screenshot shows the 'New VMware Hardware Plan' dialog box with the 'Host' tab selected. The dialog has a sidebar with tabs: Name, Host, Storage, Network, Apply, and Summary. The 'Host' tab contains the following fields:

- Host or cluster:** A text box containing 'VCCR Cluster' and a 'Choose...' button.
- CPU:** A slider and a text box showing '10.0 GHz'.
- Memory:** A slider and a text box showing '12 GB'.

At the bottom of the dialog are four buttons: '< Previous', 'Next >', 'Finish', and 'Cancel'.


5.3: Specify cluster, CPU and Memory limits

When creating a hardware plan, no changes are made in the virtual infrastructure at that time. Dedicated resource pools (one per combination of tenant and Plan) will be created only when the hardware plan is assigned to a tenant:



5.4: Veeam Cloud Connect Replication resource pool


Under the starting resource pool, each new tenant with an assigned hardware plan is mapped against a dedicated resource pool:



5.5: Hardware Plan resource pool

The resource pool clearly shows the limits for CPU and Memory that were configured during the creation of the hardware plan (10 Ghz and 12 GB). Note that, since a hardware plan is identified by a resource pool, a hardware plan cannot span over multiple clusters.


Then, one or more storage resources are assigned to the hardware plan:



5.6: Assigned storage to a Hardware Plan

Using a shared datastore that is accessible by all the ESXi hosts, a storage resource is created. Service providers can give it a "friendly name" to hide the real datastore name, and assign a quota. The quota can be overcommitted, but in this case proper monitoring should be implemented, to guarantee storage is not totally consumed, thus creating issues to the service.

Then, network resources are specified:



5.7: Specify network resources in the Hardware Plan

Service providers will assign one or more networks to a tenant, selecting how many networks should be with internet access, and how many should be internal. The difference between the two is the possibility for the tenant to publish services hosted in a failed over virtual machine to the internet, but only if the virtual machine is connected to a network with internet access.

A single network with internet access is the simplest configuration a service provider can create. More complex designs can be done, closely replicating the environment of a tenant.

NOTE: A virtual machine in vSphere can have a maximum of 10 network interfaces, so the maximum number of assigned networks is nine, as one interface is reserved for the external interface of the Network Extension Appliance. This interface is not visible to tenants.

Once all the parameters have been configured, the hardware plan is configured and is ready to be assigned to a tenant.

Customer creation and backup resources

Once the Cloud Connect infrastructure is completely configured, the license is in place and hardware plans have been created (for replica services), a service provider can start to configure users and accept their incoming backups or replicas.

In the Cloud Connect node, select “Add Tenant” to start the customer creation wizard. A username and a password needs to be configured (manually or automatically generated), and a lease time can optionally be configured, for example for trial purposes:

The screenshot shows the 'New Tenant' wizard interface. The title bar is blue with the text 'New Tenant' and a close button. The sidebar on the left is light blue and contains a 'Tenant' icon and the text 'Specify tenant name, password, assigned resource types and optional contract expiration date.' The main content area is white and contains the following fields and sections:


- Tenant** (selected in sidebar):
 - Username:
 - Password: [Generate new](#)
 - Description:
- Assigned resources**:
 - ☒ Backup storage (cloud backup repository)
 - ☐ Replication resources (cloud host)
- Automatic expiration**:
 - ☐ Contract expires: Never [Calendar](#)

At the bottom of the window are four buttons: '< Previous', 'Next >', 'Finish', and 'Cancel'.

5.8: Create a new tenant

A customer of the service provider is uniquely identified as a tenant, and thus can rent both backup and replication resources, and connect to them with the same credentials. Both options are grouped in the same section of the configuration process for this reason. In this first step, we are configuring backup resources. We will also assign replication resources to the same tenant later.


The maximum allowed bandwidth can be configured for the tenant, or left unlimited:



5.9: Specify maximum allowed bandwidth

Note: This limit is an overall value for the tenant, regardless the number of backups or replicas are going to be received by the service provider. The bandwidth available to one tenant is equally split between all tasks performed by this tenant.

In the second step, at least one cloud repository needs to be configured. A user can have multiple cloud repositories, for example with or without WAN acceleration, or stored on backup repositories with different characteristics and price per GB:



The 'Set Quota' dialog box is used to configure backup resources for a new cloud repository. It includes fields for the cloud repository name, backup repository selection, user quota, and WAN acceleration settings.

Cloud repository name:
Cloud repository 1
Specify a friendly repository name that will be shown to the user.

Backup repository:
repo2 ()
1023.7 GB free of 1023.9 GB


User quota:
100 GB

☒ Enable WAN acceleration through the following WAN accelerator:
wan1.cloudconnect.local ()

OK Cancel

5.10: Configure a new Cloud Repository

The Cloud Repository is created and assigned to the tenant:



The 'New Tenant' dialog box shows the configuration for a new tenant, including backup resources. It features a sidebar with navigation options and a main area for adding and managing cloud repositories.

Backup Resources
Add one or more cloud repositories for this tenant to use.

Tenant
Bandwidth
Backup Resources
Summary

Cloud repositories:

Repository	Friendly name	Quota	Free	WA
repo2	Cloud repository 1	100.0 GB	1023.7 GB free of 1023.9 GB	war

Add...
Edit...
Remove

< Previous Next > Finish Cancel

5.11: Tenant list of backup resources


The service provider completes the wizard, and the tenant is ready to consume the assigned backup resources.

At any point in time, the service provider can edit the tenant configuration, and modify any parameter.

Backup and backup copy jobs

Once a customer has subscribed to Veeam Cloud Connect backup resources, he can start to consume them immediately.

In the Veeam Backup & Replication installation on the customer site, the customer goes into the backup infrastructure node, the service provider's sub-node, and selects to add a new service provider:




5.12: Add a service provider

In the first step, the customer inputs the DNS name configured with round robin by the service provider. Unless the TCP port has been changed by the service provider, no additional configuration is needed:

The screenshot shows a 'Service Provider' configuration window. It has a title bar with a close button. Inside, there's a header section with a plus icon and the title 'Service Provider'. Below the header, there's a text box with instructions: 'Type in DNS name or IP address and a port number received from the service provider. If you are using more than one service provider, add a custom description to more easily distinguish between them.' The main area is divided into two sections. The left section is a sidebar with 'Service Provider' selected, and below it are 'Credentials', 'Apply', and 'Summary'. The right section contains the configuration fields: 'DNS name or IP address' with a text box containing 'cc.virtualtothecore.com', 'Description' with a large text area, and 'Port' with a spinner box set to '6180'. At the bottom, there's a checkbox labeled 'Allow this Veeam Backup & Replication installation to be managed by the service provider' with a descriptive text below it: 'Select this check box if you have managed backup contract with the service provider you are adding, and want to allow it to manage your installation remotely.' At the very bottom, there are four buttons: '< Previous', 'Next >', 'Finish', and 'Cancel'.

5.13: Configure service provider DNS name and TCP port


By hitting the "Next" button, Veeam Backup & Replication connects at the tenant side to one of the cloud gateways based on the information retrieved from DNS, and checks the SSL certificate:



5.14: SSL certificate is verified upon connecting to Veeam Cloud Connect


As explained in more details in Appendix A, the certificate is issued by a recognized Certification Authority, so no security warning is raised. In the same step of the wizard, the customer will add the username and password created for him by the service provider.

By hitting "Next" again, Veeam Backup & Replication logs into the Cloud Connect infrastructure with the given credentials, and Cloud Connect returns the resources the tenant is entitled to consume:



5.15: VBR lists the resources a tenant is entitled to consume


When the customer completes the wizard, the service provider is registered in the corresponding section, and even more important the cloud repository is registered under the available backup repositories of the tenant and can be used as a target for backup and backup copy jobs:



5.16: Cloud repository is registered and ready to be used


A backup or a backup copy job is then configured in the same exact way, regardless of if the final destination is a local repository on the tenant's premises or a cloud repository exposed via Veeam Cloud Connect.

The tenant configures a new backup copy job, selects the virtual machines he wants to protect off site using Cloud Connect. And, when it comes to the selection of the repository, the only difference compared to a regular job is that a Cloud Repository is selected:




5.17: Select a cloud repository as the backup copy job target

Then, depending on the availability at the tenant's side of a WAN accelerator, the job can be configured to leverage this capability or to use "Direct" mode:



5.18: Choose between direct mode or through WAN accelerators

The job can be completed and enabled, and both the tenant and the service provider can see it while it's running:



5.19: The service provider overview of incoming jobs from tenants

The service provider doesn't have all the details the tenant has, but it can check the progress and list how many jobs and virtual machines is receiving at any given time.


Restore data from Cloud Connect backups

Once the virtual machine copies are safely stored in a remote cloud repository, the tenant has successfully obtained an additional off-site copy of his workloads and data. From here, any restore operation can be accomplished, even if the original data has been lost at the tenant's primary site.

Restore files


The most common operation a tenant would do is to restore files. If the files that need to be restored are not in a local backup anymore, but are stored in a cloud repository, then the tenant needs to start a file restore operation.

NOTE: As of today, only files stored in virtual machines running Microsoft Windows OS can be directly restored from a cloud repository. For other operating systems, a tenant needs to retrieve the entire virtual disk containing the required files first.



5.20: Restore Windows guest files from a Cloud Repository

In this scenario, the backup browser works exactly like in a restore from a local repository. The backup is mounted directly and transparently in Veeam Backup & Replication through Cloud Connect, and the tenant can browse the content of the windows file system like it was local:



5.21: Restore files from a Windows VM

Finally, the service provider can monitor the activities done by his users, looking at the “History Tab” in the console:

CLOUD CONNECT

Cloud Connect

Cloud Gateways

Tenants

Backup Storage

repo2

Replica Resources

Last 24 hours

Success

Failed

History

System


Type in an object name to search for

JOB NAME	SESSION TYPE	STATUS	START TIME ↑	END TIME	TENANT	DATA SENT	DATA RECEIVED
Restore 7becba3f-276f-43f8-a032-4744b816d971	Cloud Restore	Success	4/26/2016 8:14 AM	4/26/2016 8:35 AM	Tenant1	36.8 MB	19.4 KB
Backup Copy Job to VCC	Cloud Backup Copy	Success	4/26/2016 7:14 AM	4/26/2016 7:14 AM	Tenant1	0.0 KB	0.0 KB
Backup Copy Job to VCC	Cloud Backup Copy	Failed	4/26/2016 12:15 AM	4/26/2016 2:16 AM	Tenant1	0.0 KB	0.0 KB
Backup Copy Job to VCC	Cloud Backup Copy	Success	4/26/2016 12:14 AM	4/26/2016 12:15 AM	Tenant1	0.0 KB	0.0 KB
Restore a4b5169b-0ac5-4ecb-a9d0-9f78666a1c30	Cloud Restore	Success	4/25/2016 8:01 PM	4/25/2016 8:57 PM	Tenant1	86.1 MB	228.6 KB
Backup Copy Job to VCC	Cloud Backup Copy	Success	4/25/2016 7:14 AM	4/25/2016 7:14 AM	Tenant1	0.0 KB	0.0 KB
Backup Copy Job to VCC	Cloud Backup Copy	Failed	4/25/2016 12:15 AM	4/25/2016 2:15 AM	Tenant1	0.0 KB	0.0 KB
Backup Copy Job to VCC	Cloud Backup Copy	Success	4/25/2016 12:14 AM	4/25/2016 12:14 AM	Tenant1	0.6 KB	1.8 KB
Backup Copy Job to VCC	Cloud Backup Copy	Success	4/24/2016 8:36 AM	4/24/2016 10:14 PM	Tenant1	586.9 KB	9.6 GB
Backup Copy Job to VCC	Cloud Backup Copy	Success	4/24/2016 12:16 AM	4/24/2016 8:36 AM	Tenant1	362.2 KB	5.7 GB
Backup Copy Job to VCC	Cloud Backup Copy	Success	4/23/2016 10:55 PM	4/24/2016 12:16 AM	Tenant1	221.2 KB	980.5 MB
Backup Copy Job to VCC	Cloud Backup Copy	Success	4/23/2016 10:42 PM	4/23/2016 10:54 PM	Tenant1	198.5 KB	135.9 MB
Backup Copy Job to VCC	Cloud Backup Copy	Success	4/23/2016 10:29 PM	4/23/2016 10:41 PM	Tenant1	197.8 KB	137.8 MB
Backup Copy Job to VCC	Cloud Backup Copy	Success	4/23/2016 10:17 PM	4/23/2016 10:29 PM	Tenant1	196.2 KB	126.5 MB
Backup Copy Job to VCC	Cloud Backup Copy	Success	4/23/2016 12:33 PM	4/23/2016 10:16 PM	Tenant1	536.6 KB	7.1 GB
Backup Copy Job to VCC	Cloud Backup Copy	Success	4/23/2016 12:32 PM	4/23/2016 12:33 PM	Tenant1	0.0 KB	0.0 KB
Backup Copy Job to VCC	Cloud Backup Copy	Success	4/23/2016 12:32 PM	4/23/2016 12:32 PM	Tenant1	0.0 KB	0.0 KB
Backup Copy Job to VCC	Cloud Backup Copy	Success	4/23/2016 12:12 PM	4/23/2016 12:13 PM	Tenant1	0.0 KB	0.0 KB
Backup Copy Job to VCC	Cloud Backup Copy	Success	4/23/2016 8:45 AM	4/23/2016 12:11 PM	Tenant1	412.6 KB	2.5 GB

5.22: Cloud Connect history


Restore virtual machines

A less common activity that can be done using data stored in a cloud repository is a virtual machine restore. This can be either of the entire virtual machine or only some of its files.



5.23: Select a virtual machine to be restored from a cloud repository

Once a tenant selects a virtual machine from the cloud repository, he can choose to restore the entire VM, for example. A dedicated wizard starts, and the steps are exactly the same as a restore from a local backup file: Restore to the original location or to a new one, configure target options like host, datastore, folder and so on. The user experience does not change at all. Once every option is configured, the restore process starts. Depending on the line speed between the tenant and the service provider, the operation could be completed in a few minutes or hours:



5.24: VM restore process from a Cloud Repository

Assign replication resources

When the Cloud Connect Replication infrastructure is correctly in place, the service providers can start to offer replication resources to their customers.

The service provider can offer the service to existing customers that are already consuming backup resources, or to a new tenant willing to only consume replica resources. In this example, we are showing the first option, but remember that both are possible:

Edit Tenant1

Tenant
Specify tenant name, password, assigned resource types and optional contract expiration date.

Username:

Password: [Generate new](#)

Description:

Assigned resources

- ☒ Backup storage (cloud backup repository)
- ☒ Replication resources (cloud host)

Automatic expiration

☐ Contract expires: Never [Calendar](#)

< Previous Next > Finish Cancel

5.25: Assign replication resources to a tenant

Once replication resources are assigned to the tenants, the new steps in the tenant wizard are enabled. First, the service provider configures replica resources:

Edit Tenant1

Replica Resources
Add one or more hardware plans for this tenant to use.

Hardware plans:

Name	CPU	Memory	Storage	Networks
HWP-tenant1	10 GHz	12 GB	datastore-tenant1 (500 GB)	1

[Add...](#)
[Edit...](#)
[Remove](#)

[Manage network settings](#)

☒ Use built-in network management capabilities during failover
We will deploy network extension appliances to allow tenants perform partial or full site failover. Select this checkbox if you are not using 3rd party solution to manage failover networking.


< Previous Next > Finish Cancel

5.26: Add a hardware plan to a tenant

Here, the service provider selects the hardware plan that has been previously created for a given tenant and assigns that plan to this tenant.

NOTE: We assume during the entire guide that the service provider is going to use the Veeam built-in network capabilities via Network Extension Appliances. If a service provider disables this option for a tenant in this step of the wizard, different technologies need to be put in place, placing them out of the scope of this guide.

Then, the service provider configures the network extension appliance that is going to be deployed on the service provider side:



5.27: Configure networking for the Network Extension Appliance

There is some information here that the service provider has to correctly fill in:

- **name:** every appliance should have a unique name in order to easily identified and managed. We suggest to use a naming convention that is followed throughout all deployments. In our example, we used "NEA-nameofthetenant";
- **External network:** Every NEA has one external network and one or more internal network. The internal networks are configured by creating different port groups, each tagged with a unique VLAN ID, as explained in this guide. The external network is instead the internet-facing network of the appliance and the NEA allows virtual machine to reach internet, and to be reached from internet (if the tenant has configured his public IP publishing rules) via this interface. This external interface has to be connected to a port group that can reach the internet, and where the subnet in use can be applied. In our case, we are connecting this interface on the same port group used for the external interface of the cloud gateways. During a partial failover the two components can communicate with each other, as this is a requirement to make partial failover work;

- **IP address:** after the port group has been selected, an IP address should be configured in accordance with the port group. We are using a public IP address in the same subnet used by the cloud gateways. This IP address should not be in the pool that was configured during the creation of the replication service — as the risk is that it could be assigned as an additional IP address to another NEA. Ideally, service providers should have two different pools, one loaded into the Cloud Connect configuration so that they can be assigned as public IP to tenants that need to publish their VMs over the internet, and another pool not loaded into Cloud Connect, but just used to configure the primary IP of each NEA.

As a recap, in our example our pool of public IPs has been divided like this:

scope	IP
gtw1	185.62.37.98
gtw2	185.62.37.99
gtw3	185.62.37.100
portal	185.62.37.101
Tenant pool	185.62.37.102 to 107
NEA pool	185.62.37.108 to 110
Disk	40 Gb

After the network extension appliance has been configured, the service provider can assign one or more public IPs to the tenant:

Edit Tenant1

Network Extension
Specify network settings to be used during failover.

Tenant

Bandwidth

Backup Resources

Replica Resources

Network Extension

Apply

Summary

Network extension appliances:

Name	Host	External network	IP address	Edit
NEA-Tenant1	vcsa.cloudco...	DPG-vcc_public	185.62.37.102	

< III >


☒ Allocate the following number of public IP addresses: 1

Public IP addresses are required for tenants to be able to perform full site failover. [Manage...](#)

< Previous Next > Finish Cancel

5.28: Allocate public IP addresses

When all settings are confirmed and applied, as you can see in figure 4.7, the first available IP address is automatically assigned to Tenant1, the Cloud Host is created, and the NEA for the tenant is deployed in the virtualized environment:



5.29: NEA deployed in the vSphere environment

The information shows us that there are two interfaces of the VM, as we configured this hardware plan to have 1 network with internet access. The external interface is connected to portgroup "DPG-vcc_public" (the shared VLAN where all the public IPs are published) while the internal interface is connected to portgroup **tenant.Tenant1.vlan112**.


This is a completely new port group that Cloud Connect created for this tenant, and the naming scheme is self-explaining. This is a portgroup created for tenants, specifically for Tenant1, and is using VLAN 112. If you check figure 4.6, VLAN 112 is the first ID of the pool assigned to network with internet access.

NOTE: The port group where the external interface is connected is considered, from security point of view, an external and untrusted area. Connections happening outside of the external interface are considered unprotected and unfiltered, unless a service provider is using additional security procedures to monitor and protect this network.

NEA deployment at the tenant side


Depending on different scenarios that may apply to a tenant, there are different moments when the corresponding Network Extension Appliance will be deployed at the tenant side:

- on a new customer consuming only replication resources, the NEA is deployed during the service provider setup wizard;
- on an existing customer already consuming backup resources, the NEA deployment is requested by Veeam Backup & Replication during the first partial failover attempt, otherwise there would be no tenant-side NEA to initiate the VPN tunnel:




5.30: No extension appliance warning at tenant side

Regardless of the scenario, Veeam Backup & Replication brings the customer to the corresponding section of the service provider setup wizard, where the tenant can configure his local Network Extension Appliance:



5.31: The network extension step during service provider setup wizard

The wizard tries to automatically choose the best network to connect the Network Extension Appliance too, but it's extremely important that the tenant edits the configuration of the NEA so that it is going to be connected to the same network of the VMs he wants to replicate towards the service provider:



The screenshot shows a dialog box titled "Network Extension Appliance Configuration" with a close button (X) in the top right corner. The dialog has a "Configuration" tab selected. Below the tab, there is a text label: "Specify host, resource pool, datastore and network settings for network extension appliance VM." The configuration fields are as follows:

- Host:** A text box containing "esxi03.skunkworks.local" and a "Choose..." button to its right.
- Resource pool:** A text box containing "Resources" and a "Choose..." button to its right.
- Datastore:** A text box containing "fas2020_vol1" and a "Choose..." button to its right.
- Network:** A text box containing "dvp-prodVM" and a "Choose..." button to its right.

Below these fields, there are two radio buttons for IP address configuration:

- ☐ Obtain an IP address automatically
- ☒ Use the following IP address:

When the second radio button is selected, a sub-dialog box appears with the following fields:


- IP address:** A text box containing "10 . 2 . 50 .200".
- Subnet mask:** A text box containing "255 . 255 . 255 . 0".
- Default gateway:** A text box containing "10 . 2 . 50 .254".

At the bottom of the main dialog, there are "OK" and "Cancel" buttons.

5.32: Configure networking for the tenant NEA

Note: The tenant NEA has only one network interface. If the tenant has received multiple networks from the service provider, one NEA will be deployed for each network. Only one NEA with multiple interfaces will be deployed at the service provider however

The NEA has deployed in the tenant vSphere environment and is ready to be powered on for any partial failover activity:




5.33: The tenant NEA as seen from local vSphere environment

Replication jobs

Now that the service provider has set up the cloud host for the tenant, it's time for the latter to start consuming his replication resources. In Veeam terms, a Cloud Host is the abstracted view of the multi-tenant environment offered by the service provider, and seen by the tenant as a remote virtualized host that can be used as a replication target.


In order to guarantee the most transparent user experience, Veeam Cloud Connect allows to replicate VMs towards the service provider by using the well-known replication jobs. Exactly like in a backup or backup copy job used to consume Veeam Cloud Connect backup, also in this case jobs are configured in the same exact manner, and only the target is different.

Once the customer has been assigned Replication resources in his subscription, the first thing he can do is rescan the services offered by the service provider:




5.34: Rescan the service provider

First, the tenant may notice that replication resources are now available by the fact that "Manage default gateways..." is enabled:



5.35: Manage default gateways is now enabled

By going into the Backup Infrastructure node, the tenant can see the Cloud Host listed under the available VMware resources, side by side with his local vSphere environment:



5.36: An overview of the Cloud Host

In the properties of the Cloud Host, the tenant can verify that the amount of resources are those requested upon subscribing the the Hardware Plan (10 GHz of cpu, 12 GB of memory, 500 GB storage and 1 network with internet access in our example), and he can verify that the service provider is using VMware vSphere 6.0.

Everything is ready for the first replication job towards Cloud Connect.

The new replication job is configured first by setting a name for the job itself:

New Replication Job

Name
Specify the name and description for this job, and provide information on your DR site.

Name

Name: Replication to VCC-R

Description:

Describe your DR site:

- ☐ Low connection bandwidth (enable replica seeding)
- ☒ Separate virtual networks (enable network remapping)
- ☐ Different IP addressing scheme (enable re-IP)

< Previous Next > Finish Cancel

5.37: Create a new replication job

It's important to select "Separate virtual networks (enable network mapping)." This option enables the network settings of the replication job, that will be important to correctly map tenant networks to the networks created inside Cloud Connect. Re-IP, on the other side, is not needed (and not available) in Cloud Connect.

Next step, a tenant selects the virtual machines that he wants to replicate towards Veeam Cloud Connect:

Edit Replication Job [Replication to VCC-R]

Virtual Machines
Select one or more VMs to replicate. Use exclusion settings to exclude specific VMs and virtual disks from replication.

Virtual machines to replicate:

Name	Type	Size
test-2012	Virtual Machine	40,0 GB
lamp	Virtual Machine	1,8 GB


Add...
Remove
Exclusions...
Source...
Up
Down
Recalculate

Total size: 41,8 GB

< Previous Next > Finish Cancel


5.38: Select one or more VMs to replicate

In this example, we are replicating a Windows 2012 VM ("test-2012") and a Linux VM ("lamp"). Next, we select the destination:




5.39: Specify the destination of the replication job

This is the only difference between a local replication, and the one towards Veeam Cloud Connect. Select the cloud host as a target and then choose the cloud host published by the service provider:



5.40: Select the Cloud Host


Then, it's time for the network mapping:



5.41: Network mapping


Here, a tenant is going to see the source network of any replicated VM and be able to map each network to a network created by the service provider. In this simple case, the tenant has only one network with internet access, so every VM is going to be mapped to the single network made available in the hardware plan by the service provider. In more complex environment, there will be multiple source and target networks to be coupled.

The rest of the job configuration follows the same steps of any replication job. We are only showing the configuration of the WAN Acceleration here:



5.42: Enable replication through WAN Accelerators


Once the job is saved and scheduled, the VMs are replicated to Veeam Cloud Connect, and the service provider can see them in the vSphere environment, under the resource pool created for the specified tenant:



5.43: The virtual machines in tenant's resource pool

Virtual machines are also named with the tenant's name as a prefix, so no name conflict or confusion can be generated inside the shared vSphere environment.

In the tenant environment, the replicated virtual machines show up in "Ready" state in Veeam Backup & Replication. They are ready to be used for failover activities:




5.44: Virtual machine replicas in ready state

Failover plans


Once all the protected virtual machines have at least one restore point stored into Veeam Cloud Connect, a tenant can use the failover capabilities he subscribed.

In order to complete a full failover, a failover plan must be configured. A failover plan is a group of virtual machines that Veeam Backup & Replication has to manage as a single entity, following the boot order and delays configured in the plan itself. When it comes to Veeam Cloud Connect, a "cloud" failover plan has additional options.




5.45: Create a cloud VMware failover plan

After giving the new failover plan a unique name (unique for the tenant, as multiple tenants at the service provider can have the same name for their failover plans without any problem), the tenant selects the virtual machine he wants to add to this plan:




5.46: Select replica VMs

Only virtual machines replicated to Veeam Cloud Connect are shown in the selection screen, and only those with at least one complete restore point available can be selected. Replica VMs are added to the cloud failover plan, where boot order and delay can be configured:



5.47: Replica VMs added to Failover Plan

In the next step, the tenant configures the different default gateways for each available network:




5.48: Specify default gateways for cloud networks

When a full failover is initiated, the network extension appliance enables the different internal networks and set the configured default gateway's IP address as its own address. The result is that any virtual machine connected to a cloud network can reach its original gateway without changing its original IP address. The NEA is simulating the original gateway of the tenant's production network.

The next step is the configuration of Public IP Addresses. The network extension appliance has one external interface, published over the internet. In addition to its primary IP address, a service provider can assign additional IP addresses to be used by tenants to publish services running on the failed over VMs.

First, a tenant enables the option "Assign public IP addresses to use during full site failover," then by using the "Add" button, he creates publishing rules like in a firewall:




The screenshot shows a dialog box titled "Public IP Address Mapping Rule" with a red close button in the top right corner. The dialog contains the following fields and controls:

- Replica VM:** A text box containing "test-2012" and an "Add VM" button to its right.
- Public IP address:** A dropdown menu showing "185.62.37.102".
- Port:** A spinner box showing "3389".
- Internal IP address of replica VM:** A dropdown menu showing "10.2.50.110".
- Port:** A spinner box showing "3389".
- Description:** A text box containing "Access test-2012 via RDP".
- Buttons:** "OK" and "Cancel" buttons at the bottom right.

5.49: Add a public IP address mapping rule for a Windows VM

For a Windows VM, the internal IP address of the replica VM is automatically recognized. However, even non-windows VMs can be published by manually writing their IP address in the box (lamp is a Linux VM):



The screenshot shows a dialog box titled "Public IP Address Mapping Rule" with a red close button in the top right corner. The dialog contains the following fields and controls:

- Replica VM:** A text box containing "lamp" and an "Add VM" button to its right.
- Public IP address:** A dropdown menu showing "185.62.37.102".
- Port:** A spinner box showing "80".
- Internal IP address of replica VM:** A dropdown menu showing "10.2.50.54".
- Port:** A spinner box showing "80".
- Description:** A text box containing "Webserver running on lamp".
- Buttons:** "OK" and "Cancel" buttons at the bottom right.

5.50: Add a public IP address mapping rule for a Linux VM


Once the failover plan is configured, it is stored directly with the service provider, and what the tenant sees locally is read in real time from the service provider. This is done because, upon a failure, a tenant may lose his entire environment, thus also losing the Veeam Backup & Replication installation and its configuration like the failover plans. By storing the failover plan directly at the service provider, it can be initiated by the service provider or by the tenant (using the cloud portal) without the need for a local Veeam installation at the tenant.

Partial failover

A partial failover is the scenario where a tenant still has his infrastructure up and running, and only one or more virtual machines are having issues. In this situation, nobody wants to run a complete site failover to solve issues of a few VMs. Partial failover allows you to start the replica VM of one or more VMs at the service provider side, and let all the other VMs run at the tenant side.

To make this possible, the technology integrated into the Network Extension Appliance extends — hence the name — any customer network to the service provider site, so that production VMs can communicate with replicas without any change in the IP addressing.

This happens because NEA creates a Layer 2 VPN tunnel for each network involved that transparently extends the tenant network to the corresponding service provider network.



5.51: Veeam Cloud Connect partial failover


The cloud gateway at the service provider is responsible for interconnecting the two NEAs, at the tenant and at the service provider. Thanks to this interconnection, OpenVPN Client running at the tenant can initiate a VPN tunnel towards the OpenVPN Server running in the service provider tenant. The final result is that a Layer2 tunnel is created between the two networks, and thanks to a Proxy-ARP solution running in both the appliances, packets can travel inside the tunnel and VMs can communicate with each other, regardless of which site they are powered on.

NOTE: Virtual machines running at the service provider can reach the internet by using the internet connection of the tenant. Any packet created at the service provider, and with a destination other than its own subnet, is forwarded to the default gateway, which is running at the tenant side.

Partial failover operation


To initiate a partial failover, the tenant selects the virtual machine he wants to failover from the ready replicas.

Note: Veeam doesn't verify if the original virtual machine is still running, thus possible IP address conflicts may occur if the tenant doesn't verify this information prior to starting the partial failover.




5.52: Start the failover of a single VM

In the wizard, the tenant can add additional VMs to the partial failover, and he can select the restore point he wants to use for each of them:



5.53: Select the restore point to be used for the failover


The wizard is finished, and after a few seconds the operation is completed:



5.54: Partial failover is completed successfully


What has happened behind this screen? A few things.

On both sides, NEAs are started so that the VPN tunnel and the Proxy-ARP components are up and running. This is the NEA at the tenant side:




5.55: NEA at tenant side is started

On the service provider side, both the NEA and the requested replica VM are started. The replica VM has the same configurations and IP address as its original copy:



5.56: test-2012 VM is started at the service provider side

The service provider can also verify in the Veeam Backup & Replication console that the failover has been started by Tenant 1:



5.57: test-2012 is in failover state at the service provider

And, he can also see the two tasks originated by the failover:

Type in an object name to search for							
JOB NAME	SESSION TYPE	STATUS	START TIME ↑	END TIME	TENANT	DATA SENT	DATA RECEIVED
VPN tunnel	VPN Tunnel	Active	5/5/2016 12:10 PM		Tenant1	184.9 KB	779.2 KB
test-2012	Cloud Failover	Success	5/5/2016 12:09 PM	5/5/2016 12:11 PM	Tenant1	0.0 KB	0.0 KB
AD in the Cloud	Cloud Backup Copy	Success	5/5/2016 8:15 AM	5/5/2016 8:22 AM	Tenant3	50.0 KB	149.5 KB
AD in the Cloud	Cloud Backup Copy	Success	5/5/2016 8:00 AM	5/5/2016 8:14 AM	Tenant3	2.0 MB	331.7 MB
DK-CloudConnect-3	Cloud Backup Copy	Failed	5/5/2016 7:01 AM	5/5/2016 9:02 AM	Tenant4	0.0 KB	0.0 KB
DK-CloudConnect-3	Cloud Backup Copy	Success	5/5/2016 7:01 AM	5/5/2016 7:01 AM	Tenant4	0.0 KB	0.0 KB
DK-CloudConnect-2	Cloud Backup Copy	Failed	5/5/2016 5:01 AM	5/5/2016 7:02 AM	Tenant4	0.0 KB	0.0 KB
DK-CloudConnect-2	Cloud Backup Copy	Success	5/5/2016 5:00 AM	5/5/2016 5:01 AM	Tenant4	0.0 KB	0.0 KB
DK-CloudConnect-1	Cloud Backup Copy	Failed	5/5/2016 4:00 AM	5/5/2016 6:01 AM	Tenant4	0.0 KB	0.0 KB
DK-CloudConnect-1	Cloud Backup Copy	Success	5/5/2016 4:00 AM	5/5/2016 4:00 AM	Tenant4	0.0 KB	0.0 KB
DK-CCJOB	Cloud Backup Copy	Success	5/5/2016 12:53 AM	5/5/2016 12:59 AM	Tenant4	353.1 KB	1.4 MB
DK-CCJOB	Cloud Backup Copy	Success	5/5/2016 12:45 AM	5/5/2016 12:53 AM	Tenant4	479.4 KB	118.5 MB
DK-CCJOB	Cloud Backup Copy	Success	5/5/2016 12:44 AM	5/5/2016 12:45 AM	Tenant4	22.7 KB	5.0 MB
DK-CCJOB	Cloud Backup Copy	Success	5/5/2016 12:41 AM	5/5/2016 12:44 AM	Tenant4	392.9 KB	15.9 MB
DK-CCJOB	Cloud Backup Copy	Success	5/5/2016 12:01 AM	5/5/2016 12:41 AM	Tenant4	340.1 MB	23.8 MB
DK-CCJOB	Cloud Backup Copy	Success	5/5/2016 12:00 AM	5/5/2016 12:00 AM	Tenant4	0.0 KB	0.0 KB
Replication to VCC-R	Cloud Replica	Success	5/4/2016 9:31 PM	5/4/2016 10:44 PM	Tenant1	31.3 KB	579.4 MB

5.58: The task list at the service provider

There is a completed "Cloud Failover" task, related to the power on of the replica VM, and a "VPN Tunnel" in active state, as the failover is still in the process.

The final result for the tenant is that any connection towards the failed over VM happens as usual:

```
Administrator: Command Prompt

C:\Windows\system32>ping test-2012

Pinging test-2012.skunkworks.local [10.2.50.110] with 32 bytes of data:
Reply from 10.2.50.110: bytes=32 time<1ms TTL=128
Reply from 10.2.50.110: bytes=32 time<1ms TTL=128
Reply from 10.2.50.110: bytes=32 time<1ms TTL=128
Reply from 10.2.50.110: bytes=32 time=1ms TTL=128

Ping statistics for 10.2.50.110:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
    Approximate round trip times in milli-seconds:
        Minimum = 0ms, Maximum = 1ms, Average = 0ms

C:\Windows\system32>ping test-2012

Pinging test-2012.skunkworks.local [10.2.50.110] with 32 bytes of data:
Reply from 10.2.50.110: bytes=32 time=84ms TTL=126
Reply from 10.2.50.110: bytes=32 time=45ms TTL=126
Reply from 10.2.50.110: bytes=32 time=38ms TTL=126
Reply from 10.2.50.110: bytes=32 time=38ms TTL=126


Ping statistics for 10.2.50.110:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
    Approximate round trip times in milli-seconds:
        Minimum = 38ms, Maximum = 84ms, Average = 51ms

C:\Windows\system32>_
```

5.59: The connection to the original VM and its replica

We can see two ping operations here: The first one against the original VM has a time below 1 ms and a TTL of 128, signs that the ping was connecting to a local VM. The second test has higher latency and a TTL of 126, a clear sign that the connection is still over a Layer2 network — both machines are in the same subnet — but the link is towards a remote location.

The partial failover is correctly working, and can be kept up and running as long as the tenant needs it. Once the failover is not needed anymore, the tenant can choose among different options:



5.60: Options for a failed over VM

Undo failover stops the replica VM at the service provider side, and any change applied to that VM is lost. If the tenant has made some changes to the replica VM and wants that version to be the one to be used from now on, he can use the failback option to replicate the replica VM back into his production site.

Full Failover


A full failover is the scenario where a tenant only relies on the copy of his infrastructure stored at the service provider to run his virtual machines, as his local infrastructure is not up and running. The two differences with the Partial Failover can be listed here, from the point of view of Veeam Cloud Connect:

- only the NEA at the service provider side is involved, as the full failover assumes there is no infrastructure component left at the tenant side;
- the full failover has to be started using a cloud failover plan. Any failover of a single VM starts a partial failover.

In a full failover situation, the Network Extension Appliance doesn't act anymore as a VPN extension, but its two roles are to become the new default gateway for every network created in the hardware plan, and to publish services running in failed over VMs to internet, like a firewall.

There are three different ways to start a full failover:

- the tenant selects the cloud failover plan from his Veeam installation and starts it. This is an unlikely situation as, again, a full failover is usually required because the tenant has lost his infrastructure;
- the service provider, upon a request by the tenant, selects the corresponding cloud failover plan and starts it on behalf of the tenant. This is a more likely scenario, especially for those tenants not confident in using the self-service capabilities of Veeam Cloud Connect:




5.61: A service provider can start a cloud failover plan on behalf of a tenant

- The tenant can connect to the Veeam Cloud Connect Portal using the credentials received from the service provider upon subscribing to the service, and he can start his failover plan without involving the service provider at all.


Start the failover plan using the Cloud Connect Portal

The tenant starts the procedure by connecting to the Veeam Cloud Connect Portal and logging in with the received credentials:




5.62: The login screen of the Veeam Cloud Connect Portal

After the tenant has successfully log in into the portal, he can see his failover plan(s), safely stored at the service provider:



5.63: The list of available cloud failover plans

The tenant can now select one of the available cloud failover plans by using the "Start" button on the top left to start it. The portal asks to the tenant which point in time the tenant wants to use as a restore point:



5.64: Select the point in time to be used as a restore point

Immediately after the selection, the failover plan is executed. After a few seconds, the plan is successfully completed:

The screenshot shows the Veeam Cloud Connect Portal interface. The left sidebar has a 'FAILOVER PLANS' section with a 'SESSIONS HISTORY' link. The main content area displays a table of failover plans. The 'Cloud failover plan' is selected, showing its status as 'Completed' and a detailed log of the execution process.

NAME	STATUS	CREATED ↓	FINISHED
lamp	✓	5/5/2016 01:31:25 pm	5/5/2016 01:32:42 pm
test-2012	✓	5/5/2016 01:31:25 pm	5/5/2016 01:32:12 pm
✓ Cloud failover plan	✓	5/5/2016 01:31:23 pm	5/5/2016 01:32:43 pm

Job started at 5/5/2016 1:31:23 PM
 Failover plans view can be refreshed manually by pressing F5
 Building VMs list
 Setting up network extension for tenant Tenant1 with routing between networks disabled
 Processing VM: test-2012
 Waiting 60 sec before the next VM
 Processing VM: lamp
 Failover plan executed, 2 VMs processed. Successes: 2, Warnings: 0, Errors: 0.
 Job finished at 5/5/2016 1:32:43 PM

5.64: The cloud failover plan is executed successfully via Veeam Cloud Connect Portal

The tenant can open the details of the two failed over VMs and check the IP publishing rules:

The screenshot shows the Veeam Cloud Connect Portal interface with the 'SESSIONS HISTORY' link selected. The main content area displays a table of failover plans. The 'test-2012' and 'Cloud failover plan' rows are expanded, showing detailed logs of the failover process for each VM.

NAME	STATUS	CREATED ↓	FINISHED
lamp	✓	5/5/2016 01:31:25 pm	5/5/2016 01:32:42 pm
test-2012	✓	5/5/2016 01:31:25 pm	5/5/2016 01:32:12 pm
Cloud failover plan	✓	5/5/2016 01:31:23 pm	5/5/2016 01:32:43 pm

Validating VM
 Performing failover for VM lamp to state as of less than a day ago (9:49 PM Wednesday 5/4/2016)
 Reverting VM to the restore point snapshot
 Powering on VM
 Applying full site failover settings
 Network 10.2.50.0/24 (HWP-tenant1 network 1 (internet)) has been already initialized with internet access enabled
 Enabling endpoint 185.62.37.102:80 to access service 10.2.50.54:80 (Webserver running on lamp)
 VM lamp has been failed over to the state as of less than a day ago (9:49 PM Wednesday 5/4/2016) successfully


Validating VM
 Performing failover for VM test-2012 to state as of less than a day ago (9:32 PM Wednesday 5/4/2016)
 Reverting VM to the restore point snapshot
 Powering on VM
 Applying full site failover settings
 Initializing network 10.2.50.0/24 (HWP-tenant1 network 1 (internet)) with internet access enabled
 Enabling endpoint 185.62.37.102:3389 to access service 10.2.50.110:3389 (Access test-2012 via RDP)
 VM test-2012 has been failed over to the state as of less than a day ago (9:32 PM Wednesday 5/4/2016) successfully

5.65: The details of both failed over VMs

Here, you can see the two different publishing rules that were previously configured during the creation of the failover plan:


185.62.37.102:80 -> 10.2.50.54:80 185.62.37.102:3389 -> 10.2.50.54:3389

The first rule publishes the webserver running on the VM that is called lamp:



5.66: Webserver is published via NEA Public IP Address

In the same way, the windows VM is now reachable via RDP protocol via the NEA appliance and its public IP address:



5.67: RDP is published via NEA Public IP Address

The failover plan has been correctly executed, and all the needed publishing rules have been applied successfully.

Monitoring Cloud Connect with Veeam ONE

Veeam® ONE™ is a powerful monitoring, reporting and capacity planning tool for the Veeam backup infrastructure and VMware vSphere and Microsoft Hyper-V virtual environments. It helps enable Availability for the Always-On Enterprise™ by providing complete visibility of the IT environment to detect issues before they have operational impact.

Veeam ONE is available both as a paid product and in a free edition. Veeam ONE Free Edition can help service providers track each tenant's valuable Veeam Cloud Connect usage data. This free feature provides three Veeam Cloud Connect infrastructure reports, including reports showing:

- Quota usage over the past period
- Date estimates when cloud repositories will run out of available storage capacity
- Information to help avoid overprovisioning backup repositories

The feature also gives real-time monitoring of Veeam Cloud Connect component health, with predefined alarms and built-in summary dashboards in Veeam ONE Monitor Client.


What you can do with Veeam ONE Free Edition

A service provider is going to deploy and use Veeam ONE Free Edition to start monitoring his Veeam Cloud Connect infrastructure.

The first step towards this goal is to install Veeam ONE. For this activity, a dedicated server is preferred:

ONE	
server name	one.cloudconnect.local
IP Address	10.10.51.42
Operating System	Windows Server 2012 R2
Installed components	Veeam ONE
vCPU	2
RAM	2 Gb
Disk	40 Gb

On this server, the service provider installs Veeam ONE. There is no dedicated installation media for the Free Edition. Obtain a Free Edition installation by selecting this edition during the common setup:



6.1: Install Veeam ONE in Free Edition mode


Veeam ONE Free Edition is missing multiple features compared to the full edition. But, Veeam decided to let service providers monitor their Veeam Cloud Connect infrastructures with the Free Edition. Specifically:

- Performance monitoring and alerting for Veeam Cloud Connect infrastructure: Displays Veeam Cloud Connect jobs including their latest state and performance statistics
- Veeam Cloud Connect reporting: This report pack provides information about Veeam Cloud Connect infrastructure, including user quota usage, capacity planning for cloud repositories and configuration data for cloud gateways and repositories. The reports include Cloud Connect Inventory, Overprovisioned Backup Repositories, Veeam Cloud Connect Replication Provisioning and Veeam Cloud Connect User Report.

Monitoring, alarms and reporting


Monitoring capabilities in Veeam ONE (not limited to the Veeam ONE Free Edition) are limited as of now to backup resources. Replication resources cannot be monitored today.

In the monitoring panel, a service provider can open the Data Protection View, and different information will be available here:




6.2: Data Protection View

From here, a service provider can select one of the Cloud Connect components and see additional details in real time. For example, WAN accelerators usage and data savings:




6.3: WAN Accelerators statistics

Other information can be obtained in regards to Cloud Repositories:



6.4: Cloud Repositories statistics

And also Cloud Gateways:



6.5: Cloud Gateways statistics

Reporting

Real time monitoring is of great help for an operations team that has to control the behavior of a Veeam Cloud Connect environment, but it may lack some historical information that could help analyze and predict trends, like storage consumption.

Veeam ONE Free Edition also offers some pre-defined reports that service provider can run for this reason:

The screenshot shows the Veeam Cloud Connect interface. On the left is a sidebar with a search bar and a list of folders including 'All folders', 'My reports', 'VMware infrastructure assessment', 'Veeam Cloud Connect', 'Veeam Backup & Replication', 'VMware overview', 'VMware monitoring', 'VMware optimization', 'VMware capacity planning', 'VMware configuration tracking', 'Custom reports', 'Offline reports', 'All deployment projects', and 'User reports'. The main area displays a list of pre-defined reports under the 'VEEAM CLOUD CONNECT' header. The reports are:

Name	Folder
Cloud Connect Inventory This report provides general inventory information about your Cloud Connect infrastructure.	Veeam Cloud Connect
Cloud Connect Replication Provisioning This report helps to identify potential impact of excessive over-provisioning of hosts and datastores using for cloud replication.	Veeam Cloud Connect
Over-provisioned Backup Repositories This report helps to identify potential impact of excessive over-provisioning of repositories in your backup environment.	Veeam Cloud Connect
Veeam Cloud Connect User Report This report provides data about Cloud Connect users, their quota usage over the past period, also it estimates the date when cloud repositories will run out of available storage capacity.	Veeam Cloud Connect

6.6: Pre-defined reports for Veeam Cloud Connect


With these reports, service providers can obtain information about the status of their Veeam Cloud Connect environment without having to spend time to create and customize dedicated reports.

NOTE: Reports cannot be scheduled in the Free Edition. Service providers need to upgrade to the Full Edition to have scheduled reporting.

Alarms

Finally, Veeam ONE has a complete set of pre-defined alarms to monitor and be notified about any possible issue the infrastructure can have, both the virtualized environment and the Veeam infrastructure itself.

First, the service provider configures how he wants to be notified about alarms in the general options:



The screenshot shows the 'Server Settings' dialog box with the 'SMTP Settings' tab selected. The dialog has a blue title bar and a tabbed interface. The 'SMTP Settings' tab is active, showing fields for SMTP server, port, from address, login, password, and checkboxes for SSL security and authentication. There is also a 'Send test email...' button. Below these fields is the 'Email format' section with radio buttons for 'HTML' (selected) and 'Plain Text (KB Articles will not be included)'. At the bottom right are 'OK' and 'Cancel' buttons.

Server Settings

SMTP Settings | Notification Policy | SNMP | Monitored Hosts | Monitored Datastores | Monitored VMs | Other Settings

SMTP server settings:

SMTP server: Port:

From:

☐ Enable SSL security

☐ Use authentication

Login:

Password:

Email format:

Send email notifications in this format:


☒ HTML ☐ Plain Text (KB Articles will not be included)

6.7: Veeam ONE server settings

Service providers can configure email settings and use SNMP as the preferred method to receive notifications here.

Once the notification options are configured, different alarms related to Veeam Cloud Connect are available. Here is one example: A service provider wants to be notified if a cloud gateway is not available — as this may cause some issues to the incoming connections — and additional load to the remaining cloud gateways.

There is a pre-defined alarm for this condition. A service provider can access this alarm by selecting Cloud Gateways node in the monitoring pane, then the alarm tab, then the "Defined alarms" option. In the following box, the service provider sees the alarm already created and enabled by default:




6.8: Pre-defined Cloud Gateways alarms

By editing the alarm, a service provider can configure the rules that trigger the alarm (like how many minutes the cloud gateway has to be unreachable to be considered offline), the actions (send an email, send an SNMP trap, run a script) and a suppression schedule (for recurring maintenance windows, etc.).

APPENDIX A: SSL Certificates generation


Veeam Backup & Replication™ gives Service Providers the ability to generate and use a self-signed certificate during the initial configuration of Veeam® Cloud Connect. This is a quick and easy method to complete the deployment and to test it, but gives lower security to customers, since they cannot verify the certificate, and thus proving the authenticity of the Service Provider.

When a user connects to a Cloud Connect environment and uses a self-signed certificate, this is the result:



A.1: warning when using self-signed certificates

The reason for the warning is that the self-signed certificate is not signed by any of the recognized Certification Authorities:



A.2: A self-signed cert generates a trust warning

In order to properly protect Cloud Connect and give their customer comfort, the Service Provider should use a proper and generally recognized certificate, issued by one of the Certification Authorities recognized by operating systems.

Create the Certificate Signing Request (CSR)

In public key infrastructure (PKI) systems, a **certificate signing request** (also displayed as **CSR** or **certification request**) is a message sent from an applicant (the Service Provider running Cloud Connect in our case) to a Certificate Authority in order to apply for a digital identity certificate. The most common format for CSRs is the PKCS #10 specification.

Before creating a CSR, the applicant first generate a key pair, keeping the private key secret. The CSR contains information identifying the applicant (such as a distinguished name in the case of an X.509 certificate) which must be signed using the applicant's private key. The CSR also contains the public key chosen by the applicant. The CSR may be accompanied by other credentials or proofs of identity required by the certificate authority, and the certificate authority may contact the applicant for further information.

The first and most important operation a Service Provider should do is to decide the public fully qualified domain name that Cloud Gateways will use to be contacted by users. This name should match the one used in DNS and the one used in the CSR. In this guide, the public domain of the Cloud Connect service is **virtualtothecore.com**, and the fqdn (fully qualified domain name) is:

cc.virtualtothecore.com

In order to create the CSR, on the Windows Server running Veeam Backup & Replication (vbr. cloudconnect.local in this guide) a Service Provider needs first to create with a text editor an .inf file. This file (it can be called request.inf) should contain a text like this:

```
;----- request.inf -----

[Version] Signature="$Windows NT$"

[NewRequest]

Subject = "CN= FQDN, OU=Organizational_Unit_Name, O=Organization_
Name, L=City_Name, S= State_Name, C=Country_Name" ; replace
attributes in this line

KeySpec = 1

KeyLength = 2048 Exportable = TRUE FriendlyName = "cc" MachineKeySet
= TRUE SMIME = False PrivateKeyArchive = FALSE UserProtected = FALSE
UseExistingKeySet = FALSE

ProviderName = "Microsoft RSA SChannel Cryptographic Provider"
ProviderType = 12

RequestType = PKCS10 KeyUsage = 0xa0

[EnhancedKeyUsageExtension]

OID=1.3.6.1.5.5.7.3.1 ; this is for Server Authentication

[RequestAttributes]

; SAN="dns=cc.virtualtothecore.com"
```

The text parts related to the **virtualtothecore.com** example are to be changed with the specific values of the Service Provider. To obtain a valid certificate from a Certificate Authority, a proper domain name should be used. Thus, I've used for this procedure my blog domain name virtualtothecore.com, and so the FQDN is **cc.virtualtothecore.com**. You will also have to write your own information in the "Subject" line.

Note that, if you want to generate a request for a wildcard certificate, the CN portion of the subject must start with the * symbol.

After the configuration file has been edited, it can be saved in a useful location like a dedicated folder c:\certificates. Then, the Service Provider has to open a command prompt with Administrator rights (right click and select "Run as Administrator"), move into c:\certificates and use this command:

certreq -new request.inf certreq.txt

If you open the created **certreq.txt** file, its content is like this:

```
-----BEGIN NEW CERTIFICATE REQUEST-----
MIID9jCCAt4CAQAwTElMAkGA1UEBhMCSVQxETAPBgNVBAGMCExvbWJhcmR5MQ8
wDQYDVQQHDAZ
WYXJlc2UxEzARBgNVBAoMClNrdW5rd29ya3MxCzAJBgNVBAsMAklUMSAwHgYDVQQD
DBdjYy52aX J0dWFSdG90aGVjb3JlLmNvbTCCASIwDQYJKoZIhvcNAQEBBQADggEPADCC
CAQoCggEBA JUBkduH0 QfJbnt2ryIjdn5z8euMM4zHyd4C
FB d2eCXAnfaskOc3F9eW9zP1KMk0Z/8K9GfezZDkMcbno5h
nIkuwBcLoHJUeiWQDmlaDutxvgvo1RO2TEQJes5CBKB7vrEakRCco3Cq26rXEPaRx
1MjdmCOVyk 2weF9TJNUIFr1Tadw/
NWCLqWUw4ZGBsDJL0lftuQe0VmxJciZC1EZQXppsXSanSdaIZECJzHUS u0wA5nZL9pl
tv03593Kqr+qYkbocRj+T2hixA7n+Y8Bi5pO6pDOs/UdCQodteb0qCcLUCXBtQoi
mEL7uwtAPQ07RfiTX9EIeeIxx0+FHD6T7UCAwEAAACCATowGgYKKwYBBAGCNw0CAzEMF
go2LjIu OTIwMC4yMFMGCSqGSIb3DQEJJDjFGMEQwDgYDVR0PAQH/
BAQDAgWgMBMGA1UdJQQMMAoGCCsGAQU
FBwMBMB0GA1UdDgQWBATEaoWXriXLI1DePK17Mxh2s8ryRzBTBgkrBgEEAYI3FR
QxRjBEAgEJDB
Z2YnIuY2xvdWRjb25uZWNOlmxvY2FsDBpDTE9VRENPTk5FQ1RcYWRtaW5pc3R
yYXRvcgwLY2Vyd
HJlcS5leGUwgcYKKwYBBAGCNw0CAjFkMGICAQEeWgBNAGkAYwByAG8AcwBvAGYAdAA
gAFIAUwBB ACAAUwBDAGgAYQBuAG4AZQBsACAAQwByAHkACAB0AG8AZwByAGEAcABoAG
kAYwAgAFA AcgBvAHY
AaQBkAGUAcgMBADANBgkqhkiG9w0BAQUFAAOCAQEAAQaUqU2Y97wH3JhgiDvn85HEZq
+60a4WqgX XHiriIG1FnJwuzdG3k+m185N+smSX/
VlIXT9fITak034muIRpqwNJR7fz4gPaLnmNowa3Don1la 8TihI47Pezl8h76ig
04hFfSOUH7Z4Atq+2XZ55lj/mRksq2oVZUeEzHCf0V7MSQD6M3Yf/WLJGL ZG/
kDexwDz2I5W9q6vu2OwmD0eA2mHW1RjycqBJktyaZ7Hy6BF1T1F3AVyJYpTVMT/
IbDAz MYZQ 4U1/
bsKD5ZHkY2WhrRkD4D2UQpFShPdlaCYf3OP9F9FbLY4mZ7yKaQxrZWaKqRzKEa
EMPng8IKt DYJRCVAw==
-----END NEW CERTIFICATE REQUEST-----
```

Obtain a signed certificate

With the Certificate Request correctly created, it's time to obtain a signed certificate from a Certificate Authority. There are several online services where service providers can get a certificate, and some of them also offer free certificates with time limits that are useful for testing SSL connections.

The involved steps vary depending on the selected Certificate Authority, but it usually involves a validation of the CSR, a check against the registered domain via WHOIS protocol to collect the registrant email address and a verification sent to this email to validate the authenticity of the request.

Just as an example, this is the CSR verification done by the Certification Authority I've used:

Verify Server URL

The CSR you generated is designed to work with the following URL:

https://*.virtualtothecore.com

If this is not the correct URL (computed from the Common name in the CSR), or if any of the CSR Information below is incorrect, then please generate a new CSR and click the Replace CSR button.

Replace CSR

CSR Information

Common Name: *.virtualtothecore.com

Organization: Veeam Software

Org. Unit: Veeam Cloud Connect

Locality: Baar

State: ZG

Country: CH

Encryption Type: RSA

*** Hashing Algorithm:** SHA-256 with RSA or DSA and SHA-1 root

*For best RSA browser compatibility, choose the SHA-256 with RSA or DSA and SHA-1 root option.

Certificate Transparency: ☐ Do not make my certificate information public. Check the box if you do not want to publicly share your common name and SANs.

Learn more about [Certificate Transparency](#)

Note: The value for the Common Name must exactly match the name of the server you plan to secure.

A.3: CSR Verification

-----BEGIN CERTIFICATE-----

-----END CERTIFICATE-----

137

Install the Signed Certificate

Back in the Veeam Backup & Replication server, the service provider has to create a text file in c:\certificates and call it cert.cer. Then, open it with a text editor and paste in it the certificate text was received from the Certification Authority.


Then, open again a high privileges command prompt, go into the c:\certificates directory, and run this command:

certreq -accept cert.cer

Once the command is executed, the certificate is stored in the local Certificate Store of the Veeam Backup & Replication server.


In the Cloud Connect section of the Veeam Console, service provider can now select "Manage Certificates" and use the new certificate. First, choose "Select certificate from Certificate Store".

In the following screen, "Pick Certificate," the imported certificate is listed together with the pre-created and self-signed certificates. Select the bought certificate (a wildcard certificate in this example):



A.4: Pick the new wildcard certificate


Before completing the wizard, you can see a summary of the certificate parameters. Among them, you can see the Thumbprint of the certificate; this can be sent to customers for additional verification.



A.5: Certificate Summary


The certificate is now ready to be used for SSL cyphered connections.

NOTE: To manage certificates, service providers can use the Certificates MMC (Microsoft management console), a graphical interface to interact with the Certificate Store. When configured, it only requires you to select "Computer account" and then "local computer".



A.6: Certificates MMC


If a service provider opens the certificate to see additional details, this is what he will see:



A.7: Certificate details

The certificate is issued to ***.virtualtothecore.com** as requested, it's valid, and the Certification Authority ("Issued by") is recognized; this means Windows is able to recognize the Certificate Authority that signed the certificate as valid.

Connections to the Cloud Gateways can now be completed without any warning:



A.8: Certificate is successfully validated


APPENDIX B: How encryption works

Veeam® Cloud Connect offers complete encryption for data at rest thanks to Veeam Backup & Replication™ encryption capabilities, but also for data in flight, so that any information exchanged between a tenant and a service provider is protected while traveling over an insecure channel like the internet.

Security savvy people would probably like to know more details about how network encryption is implemented in Veeam Cloud Connect. This appendix is designed to give you all the details you need.

SSL and AES

Veeam Cloud Connect uses two different encryption technologies: SSL (Secure Socket Layer) and AES (Advanced Encryption Standard).



B.1: SSL and AES are the two encryption algorithms used in Veeam Cloud Connect

SSL is a generic term related to a family of communication encryption technologies. In more detail, Veeam Cloud Connect uses the latest TLS (Transport Layer Security) protocol and never falls back to older and insecure versions of SSL.

AES is also used together with SSL. It's a symmetric encryption algorithm, and the de-facto standard for advanced data encryption.

During the initial connection between a customer and a service provider, the communication channel over a single TCP port (6180 by default) is protected with SSL. By verifying the SSL certificate published by the service provider, and comparing it to the hostname the provider is using for publishing Veeam Cloud Connect itself, the customer is assured that he is effectively connecting to the chosen service provider.

Providers can also create and send customers the fingerprint of their own SSL certificate for additional security.

SSL is an asymmetric encryption algorithm. There is a private key that is safely stored with the service provider, and a public key that is published over the internet and retrieved by users. This technology is well suited for protecting communications over unsecured channels like the internet, especially for the initial handshake. But, at the same time, it's not suitable to exchange large amount of data. As an asymmetric algorithm, its performance during decryption is far worse than during the encryption. Also, it guarantees the authenticity of the server publishing the key — in this case the Cloud Connect environment — but not the user that is sending data to it.

While all control and configuration commands use the SSL/TLS tunnel, the exchange of data between a customer and a service provider — for example during a backup operation — uses the other Veeam Backup & Replication encryption based on AES-256.

Key exchange


One of the security issues of using a symmetric key over an insecure channel, like the internet, is the fact that anyone who holds a copy of the key can decrypt the cyphered data. So, it's paramount to implement a secure mechanism to safely exchange the AES keys between the customer and the service provider.

Let's see how this happens in Veeam Cloud Connect.

There are two Veeam Backup & Replication installations involved in Cloud Connect. One at the service provider, publishing the service, and one at the customer side, consuming the service.


Every activity is initiated by the customer side. When a new operation towards Cloud Connect needs to be started, the customer side sends a control command over the SSL tunnel.

The Cloud Connect installation at the service provider responds to the "Start" command, and it creates the encryption Key A. This key uses AES-256.




B.2: Cloud Connect Service creates AES key A

Using the SSL protected tunnel, Key A is securely and directly passed to the job manager of the customer and the target data mover at the service provider via a local network communication.




B.3: Key A is passed to customer and service provider data movers

The customer job manager then creates their own encryption key, Key B, to be used for encrypting traffic between the customer and the service provider data movers.



B.4: Customer VBR installation creates Key B

A new tunnel is initiated by the customer job manager using Key A. Thanks to this, only the service provider side is able to decrypt data coming from the customer side. Using this tunnel, Key B is delivered in a secure way to the data mover at the service provider side. At the same time, Key B is also delivered locally to a customer data mover using the local network.



B.5: Key B is delivered to customer and service provider datamovers

From here on, data transfers of the backup job payload is encrypted using Key B. In this way data is encrypted in flight by an encryption key that is created by the customer and not by the service provider.

Thanks to the use of Key B, customer data is safely sent to the service provider. Any attempt to intercept and modify the encrypted traffic raises a security warning as the key is only owned by the user and his service provider. This guarantees to customers the avoidance of possible middleman attacks.

When additional data transfers need to be done during a new session between the customer and the service provider, the entire process is repeated from the beginning and new keys, Key A and Key B, are generated again.


Finally, when WAN acceleration is used, the process is exactly the same, and Key B is also passed to both WAN accelerators at source and target locations. In this way, WAN acceleration is able to decrypt on-the-fly data blocks encrypted by the users.

This solution allows the protection of Cloud Connect user communications, as it's leveraged regardless of whether or not the user is also encrypting their backups. For backup file encryption an additional set of AES keys are created to encrypt backup files. Those keys are not related to the communication keys.

APPENDIX C: Customize your Veeam Cloud Connect Portal

A service provider can guarantee that tenants have easy access to their cloud failover plans by deploying the Veeam® Cloud Connect Portal. Tenants can run cloud failover plans to switch to VM replicas in the cloud disaster recovery (DR) site in an easy and secure way.

By default, Veeam Cloud Connect Portal looks like this:



C.1: Veeam Cloud Connect Portal login screen

Since many tenants will access the Veeam Cloud Connect Portal on a regular basis, it represents a critical component of the service provider's branding strategy. It communicates who the service provider is, and what they offer, through its organization and appearance.

Service providers can give Veeam Cloud Connect Portal that familiar look and feel by using the company's color scheme and name. This is very easy to achieve because Veeam provides both simple and quick options to brand the Veeam Cloud Connect Portal.

Company name

A company's name plays a crucial role not only in a brand's growth, but also in the customer's perception of it. This makes the company name one of the first impressions the company has on many people.

Most of the visible information in Veeam Cloud Connect Portal is editable via app.js file, and the company name is no exception.

In order to replace Veeam® with the name of your company, you should find the following files, the login page:

```
[PortalDir]Scripts\build\production\LoginApp\app.js
```

and the application page:

```
[PortalDir]Scripts\build\production\veeamCloud\app.js
```

Locate the string including Veeam© in both files:

```
componentCls:"app-header-title",bind:{html:'<div class="sup">Veeam
<sup>®</sup>
</div>
```

Replace it with the appropriate name (Wonderland Corp. in this example):

```
componentCls:"app-header-title",bind:{html:'<div
class="sup">Wonderland Corp.<sup>®</sup>
</sup></div>
```

You'll see something similar to the following:



Type in your credentials:

User	Enter your Cloud Connect user name
Password	Enter your password
☐ Remember me	
Login	

C.2: Personalized company name in Veeam Cloud Connect Portal

The other missing part is the portal name. That is also configurable via text editor. This time, we're going to replace the line containing the application name:

```
appName: " Cloud Connect Portal "
```

Enter your new application name:

```
appName: " Awesome Portal "
```


The application title still preserves its previous name. So, we have to open web config file:

```
[PortalDir]\Web.Config
```

Change the title name:

```
<add key="AppName" value="Awesome Portal"/>
```

Now, everything is set up how we want, so it's time for us to proceed with color scheme customization.



Type in your credentials:

User:

Password:

☐ Remember me

Color

Branding and color go hand in hand because color provides a direct method for communicating meaning and messages without words.

The Veeam Cloud Connect Portal color scheme is present in three CSS files:

```
[PortalDir]Scripts\build\production>LoginApp\resources>LoginApp-all.css
```

```
[PortalDir]Scripts\build\production\VeeamCloud\resources\VeeamCloud-all_01.css
```

```
[PortalDir]Scripts\build\production\VeeamCloud\resources\VeeamCloud-all_02.css
```

By default, Veeam uses green (#54b948). So, you need to replace it across the mentioned documents with a hex code for the color you want. For instance, the following example uses blue (#4c05ff):



Most of information regarding styles, names and colors of the Veeam Cloud Connect Portal are present inside the js or css files. Having a basic understanding of JavaScript and CSS should help you to modify it even further.

APPENDIX D: Advanced Registry settings

Veeam® Cloud Connect has many features and options that are available in the Graphical Interface. Service providers can configure and tune the software via the interface, or by using PowerShell or restful API as needed.

But, for even more special configurations, additional registry keys are available in the software.

WARNING: *Default parameters are configured in the software because those are considered the best value for any given option, according to Veeam's tests and field experience. Please change these values carefully, and always take your time to evaluate the effect any single change may have. If not sure, please involve Veeam Support to assist you.*

Also note, unless stated differently, all registry keys need to be created in:

HKLM\SOFTWARE\Veeam\Veeam Backup and Replication

Finally, the Veeam Cloud Connect Service frequently rescans the registry, so it is not usually necessary to restart it to apply registry changes.

General

CloudIgnoreInaccessibleKey:DWORD = 1

When assigning a certificate to Cloud Connect Service, it tries to get the private key of the certificate, in order to check if it's accessible. If not, an error will show up in the GUI preventing the wizard from proceeding. This is a proactive way to handle situations when the certificate is imported — for example — under a wrong user account or due to another reason leading Cloud Service to be unable to access the private key. Without this verification, the provider will pass all the wizard steps, but all the tenants' jobs will fail. For newer cryptographic providers — such as Microsoft Software Key Storage Provider — there might be a situation when the said pre-check fails permanently, but after that, all the actual usage of the certificate via API goes well. So, we still don't want to disable the check for all providers permanently, but those having "modern" certificates do need to set this registry key as a workaround.

CloudConnectEnhancedSecurityMode:DWORD = 1

This key has to be set on the tenant's side. When enabled, a strict match of cloud gateway FQDN against the Cloud Service certificate is required for every cloud connection in order for jobs not to fail. Otherwise, (default value is 0) the certificate is checked against the provider's FQDN upon initial connection only, and any mismatched names of the certificate and the Gateway Name/IP are ignored. This setting is needed for those providers who have a certificate issued by trusted CA, but have specified gateway IP instead of FQDN in the Gateway wizard. In this case, the tenant's Veeam Backup & Replication cannot validate a certificate against a gateway name, as it only knows the gateway IP to connect to (not FQDN).

CloudReplicaNoStaticIpSDetectedWarning:DWORD = 1

It removes the warning when a Linux VM is replicated, stating that the IP address is not identified: "Static IP address not found".

`CloudConnectionTimeoutSeconds:DWORD = 15`

It only works on V9.0.0.1491+, do not apply on any older version. Key to extend SSL Connection attempt, if connection attempt times out after 15 seconds. This is a Tenant Side key.

`DisableVpnServerFirewall:DWORD = 0`

Set value to 1 to disable Firewall on Service Provider NEA. By default, only IP addresses fetched from CloudGateway servers are allowed to connect to the NEA during a partial failover.

`CloudConnectReportTime: REG_SZ="HH:MM"`

Specifies the local time (in HH:MM 24h format) when the daily Cloud Connect e-mail report is to be sent.

`DisableSuccessCloudConnectReport:DWORD=1`

Disables sending the daily Cloud Connect e-mail report if all tenants' jobs have a result of "success." The report will be sent only if at least one error or warning is present.

Anti DDoS prevention

The cloud gateways have dedicated configurations to prevent DDoS (Distributed Denial of Service) attacks.

`PeerCloudConnectionsLimit:DWORD = 64`

Allowance for number of tenant connections to a gateway. The key goes on gateway servers only. The default value was 16 (v8), then increased to 64 (v9).

`MaxSimultaneousCloudConnections:DWORD = 1024`

Sets the number of concurrent streams to a gateway (regardless of tenant count). The key goes on the gateway servers only. The default value was 256 (v8), then increased to 1024 (v9).

These keys should be specified on a cloud gateway in:

`HKEY_LOCAL_MACHINE\SOFTWARE\Wow6432Node\Veeam\Veeam Gate Service`

Encryption

`EncryptedTenantBackupsOnly:DWORD = 1`

This key forced every incoming tenant to send only encrypted backups to Veeam Cloud Connect backup.

Note: This option is enforced for every tenant; it's not possible to set this option per single tenant.

About the Author



Luca Dell'Oca (vExpert, VCAP-DCD, CISSP) is EMEA Cloud Architect for Veeam Software based in Italy. Luca is a popular blogger and an active member of the virtualization community. Luca's career started in information security before focusing on virtualization. His main areas of expertise are VMware and storage design, with a deep focus on Cloud Service Providers and Large Enterprises.

Follow Luca on Twitter [@dellock6](#) or [@Veeam](#)

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 less than 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 39,000 ProPartners and more than 193,000 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>.

***FAST,
SECURE***


***BACKUP
TO THE CLOUD***

Veeam Cloud Connect provides a fully integrated, fast and secure way to backup, replicate and restore from the cloud

Customer On-Premises Infrastructure



Service Provider Infrastructure



SSL/TLS
Replication
WAN Acceleration
(optional)