Veeam Modern Data Protection

Aligning your business needs with Veeam Backup & Replication

Barry Coombs & Chris Snell

Modern Data Protection
Built for Virtualization

#1 VM Backup

Introduction

Backups and disaster recovery are far more important than most IT users give credit for. Events such as the September 11th attacks in New York, the Buncefield explosion in the UK or the 2011 earthquake and tsunami in Japan show that every organisation needs a robust and tested disaster recovery plan for every site. Whilst we may be happy in the thought that the likelihood of disasters of this size happening to our businesses is very small, the reality is your disaster could be brought on from something as small as a failing air conditioning unit to a virus outbreak. Disasters and data loss can occur at any time, for any site.

There are various quotes available on the Internet regarding the failure of an organisation to recover from a major disaster, the associated loss of IT infrastructure and data and how long that organisation could continue to trade. The truth is difficult to judge, but the best course of action is for any organisation to review its own procedures and ascertain how critical their data is.

Once the data has been valued, the dilemma facing every organisation is how much of their precious budget should be consumed by the IT organisation, and for the CIO/CTO, how much of his IT budget should be used to protect the company's data. Laws such as Sarbanes Oxley (known as SOX) require many companies to protect data to a certain degree regardless.

In the following paper, we are going to discuss how to identify the criticality of your systems and data to your business and examine the various technologies included in Veeam® Backup & Replication™, and how organisations can use these technologies to align their business needs to their data protection.

Background

Whilst working with IT we are used to focusing on the technology and how it aligns to the business needs; for example, thinking in virtual machines (VMs) and datastores rather than services and procedures. When planning our data protection strategies, we need to take a step back from the technology and understand the business requirements, and once we understand the business-led requirements, we are able to identify the correct areas of technology to meet these needs.

Identifying the business needs

Before looking at the technology, we need to review what IT is actually delivering to the business. Whilst we may like to think that we know the business's IT Services inside out, it may be the case that some of the most critical services are actually user deployed and rely on an access database on a local drive. A key step to identifying the services that your business consumes is to create a simple service catalogue.

This service catalogue should be your central point to help you identify the services in use within the organisation and their criticality to the business. From there, you will be able to expand the service catalogue to contain important information that will allow you to protect your VMs, and subsequently, the business services.

Building a service catalogue

There are many tools and methodologies in place for building an IT service catalogue, on a base level with some key headings and a spreadsheet we are able to capture key information to allow us to build our data protection strategy.

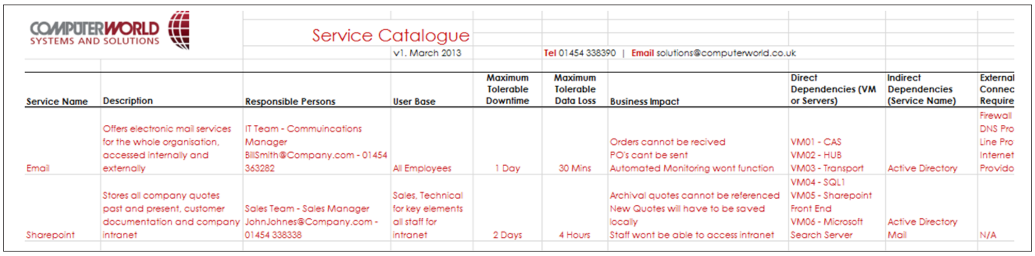

Service Catalogue

v1. March 2013

Tel 01454 338390 | **Email** solutions@computerworld.co.uk

Service Name	Description	Responsible Persons	User Base	Maximum Tolerable Downtime	Maximum Tolerable Data Loss	Business Impact	Direct Dependencies (VM or Servers)	Indirect Dependencies (Service Name)	External Connec Require
Email	Offers electronic mail services for the whole organisation, accessed internally and externally	IT Team - Communications Manager BillSmith@Company.com - 01454 363282	All Employees	1 Day	30 Mins	Orders cannot be recived PO's cant be sent Automated Monitoring wont function	VM01 - CAS VM02 - HUB VM03 - Transport	Active Directory	Firewall DNS Pro Line Pro Internet Provido
Sharepoint	Stores all company quotes past and present, customer documentation and company intranet	Sales Team - Sales Manager JohnJohnes@Company.com - 01454 338338	Sales, Technical for key elements all staff for intranet	2 Days	4 Hours	Archival quotes cannot be referenced New Quotes will have to be saved locally Staff wont be able to access intranet	VM04 - SQL1 VM05 - Sharepoint Front End VM06 - Microsoft Search Server	Active Directory Mail	N/A

Service name – This should be a friendly name that the service is known to the business as, at this level we shouldn't be thinking of application names necessarily.

Description – Provide a base description of the service and what it delivers to the business, As obvious as this may seem, remember that in a disaster situation it may not be you or the most appropriate persons recovering the systems.

Responsible persons – List key details for the persons responsible for administering and maintaining the service. This may well be someone within IT, but crucially could also include key users. If you have had a disaster and need to recover a service it is good to know which users may be able to assist you in fully testing the application prior to making it live to the business.

User base – Who are the consumers of the service?

Maximum Tolerable Downtime (MTD) – This figure isn't supposed to list ideologies around how quickly you would like to get the data online. Instead, it should list how long the business can cope without failing with the loss of this service.

Maximum Tolerable Data Loss (MTDL) – Again this isn't supposed to list the ideal scenario for how much data you can afford to lose. Instead, is though list how much data can the business tolerate to lose prior to having severe business implications.

Business impact – No matter how obvious the impact of losing the service would be, use this as an opportunity to delve a little deeper. For example, whilst loosing email services may mean email can't be sent and received, it could also mean that orders won't be received or that critical SAN alerting won't be received.

Direct Dependencies – Here is your opportunity to align the business services against your infrastructure. List the VMs and physical machines within your environment from which the service runs, such as the SQL front end and the IIS front end.

Indirect Dependencies – List the other services upon which this service is dependent, such as the active directory service or email services. By knowing this information, you are able to start formulating a plan for restart policies in case of failure.

Network Requirements – List the networking requirements that the service requires. This may be a certain firewall and an external link in the case of an email service, or, for an internal service this maybe a certain VLAN and router.

Storage Requirements – List the datastores or storage devices that this service is dependent upon.

Recovery Time Objective (RTO) – Once we have collected the information above, we are able to start setting the metrics that will allow us to set the standards of data protection for the business. The RTO sets a realistic expectation as to how long it is going to take IT to recover a given service. Keep in mind that your RTO should include the time needed to make the decision to failover in the first place.

Recovery Point Objective (RPO) – The RPO sets the expectation of how much data loss is acceptable during a disaster scenario.

The service catalogue should initially be completed by using existing knowledge within the IT department. Subsequently, further information should be gathered from key business stakeholders to ensure no element has been missed. At the initial stage, we would recommend leaving the RTO and RPO values, and once all the data has been collected, we recommend that agreed upon data protection tiers be put into place to set the RTO and RPO standards.

Data retention

Data retention requirements are an important aspect of any data protection strategy. These may be set out by a corporate policy or an industry regulatory body. Often we are able to move protected data to a cheaper storage medium the longer it will be retained. For example, data that requires short-term quick recovery will be replicated to a DR site, medium term will be stored on a backup server or NAS with cost-effective disk, and data that requires long-term retention will be moved to tape or the cloud.

You also need to think about how and where you data will be stored. Backup experts claim that to build a successful data protection and disaster recovery plan, you should follow the 3-2-1 rule:

- 3: You must have three copies of a backup file in different locations
- 2: You must use two different types of media to store copies of a backup file, for example, disk storage and tape
- 1: You must keep at least one copy of a backup file off site, for example, in the cloud or at the remote site

Differentiating between backup and replication

Whilst backup and replication are often used interchangeably, generally they are used for different purposes.

Backup is generally used for long-term data protection comprised of full VMs through to granular data such as e-mails. We generally will look to keep backups for an extended period of time to meet data retention needs. Recovery of backups is going to generally take longer than replication. Whilst Veeam is able to instantly restore VMs from a backup, a full scale recovery typically takes longer than booting up replicated VMs.

Replication allows us to protect complete VMs including their enclosed data. Typically we will run replication on a set schedule many times during a day or on a continuous schedule. With replicated VMs, we generally keep fewer recovery points as it can be more costly to store VMs online on live storage. If the need to bring the replicated VMs online occurred, we could typically complete this with ease in a few simple steps.

Defining data protection tiers for RTO and RPO

Once the service catalogue has been completed, suitable RTO and RPO values will need to be created. This will save undue complexity, and we have found creating protection tiers allows the requirement to be simplified and eases on-going management. MTD and MTDL values should be utilised that were collected within the service catalogue to help define the RTO and RPO values. These should not act as a goal, but rather ensure that these critical time points are met within a finalised strategy. Normally we are able to put together three to five tiers such as Gold, Silver and Bronze where we are able to allocate to all services to meet these needs.

For example, the tiers could be:

Gold Tier

DR: RTO = 2 hours; RPO = 4 hours; retention = 24 hours

Backup: RTO = 2 hours; RPO = 12 hours; retention = 7 years

Silver Tier

DR: RTO = 6 hours; RPO = 12 hours; retention = 12 Hours

Backup: RTO = 12 hours; RPO = 24 hours; retention = 6 months

Bronze Tier

DR: RTO = 24 hours; RPO = 24 hours; retention = 24 hours

Backup: RTO = 24 hours; RPO = 24 hours; retention = 3 months

Risk analysis

Hopefully there will be a corporate DR strategy in place that documents the full risk to the business, which will enable you to understand what you are protecting against within your data protection strategy. Unfortunately, and more often than not, a corporate DR strategy lives and dies with the IT department. As such, you may have to undertake your own risk analysis to understand what you need to protect against.

Some examples of risk are as follows:

- Hardware failure
- Software failure
- Fire
- Natural disaster
- Virus/cyber attack
- Malicious users

It may be also worth defining the priority of the risks that you are protecting against, as often when it comes down to making the budgetary decisions, you may need to make decisions regarding which elements will need to be dropped from the strategy.

DR plan & DR run book

Once we have fully defined our business needs and designed our chosen solution, it is imperative that a DR plan and DR run book are created. It may not always be possible for the person that has written and implemented the DR strategy to be the same person that will initiate a failover in the event of a DR situation.

These documents should contain a clearly defined plan that details who should take what actions and in what order. By automating these processes (e.g. IP changes), the run book can be simplified further, which will be of a huge advantage during a DR event. The DR plan should clearly define how a DR is initiated and whom is able to declare a DR event. It should further outline timeframes and processes that will need to be called upon. The run book should breakdown the technical steps into a simplified process to allow whoever is necessary or available to conduct the failover.

Storing these key documents will need to be carefully considered, as storing a copy on the production file server may not be the most intelligent action if you are trying to recover that very server. Consider storing copies at each site and maybe a printed copy with a legal representative.

Test, test and test again

We are firm believers that your DR / data protection plan is only as good as your last test. Creating a firm and robust plan, including a schedule, to test your backups and replication is the key to success. Not only will this allow you to prove the technology, but it will also firm up the DR plan.

Once you have identified, defined and documented your business needs, it is time to look at how you deliver the requirements through technology. In more traditional IT infrastructures, often the technical capabilities force a shape on the final solution that isn't ideal and due to either cost constraints or simple lack of functionality. The final solution does not fully meet the carefully defined and agreed upon plans and processes.

Virtualisation has had a major impact in the areas of what can be achieved in DR, data protection and operational flexibility. DR no longer needs exactly the same IT infrastructure to be replicated on a second site. When protecting applications, not only can you protect the application data, but also the application and the hardware configuration required to run the application all in one backup pass. This flexibility means you can define your business requirements and deliver them, rather look at what your infrastructure can do and then define your business requirements accordingly.

Whilst virtual infrastructures give the flexibility to deliver more, you still need the right tools to run the processes, manage the data and perform recovery functions in a timely, controlled way when required. By using the correct tool, you can also extend the capability of your virtual infrastructure so that it far exceeds anything possible with a traditional physical infrastructure.

One such product that really extends the virtual infrastructure is Veeam Backup & Replication. In this product set, there are a number of features such as SureBackup® and Virtual Lab that will allow you to easily test backups and replicas with ease on a regular basis and help you deliver your agreed upon levels of service to the business. In the next section, we will look at some of the key features that Veeam provides and how they tie into the delivery of RPOs, RTOs, SLAs and more.

How Veeam Backup & Replication can help you meet your business needs

Veeam Backup & Replication is a robust data protection platform for your virtualised environment. Veeam Backup & Replication is agentless, meaning that there is no troublesome software to install for specific use cases, allowing all the recovery options to work from simple image-level backups.

Veeam offers organisations the following types of technology to protect their servers:

Disk

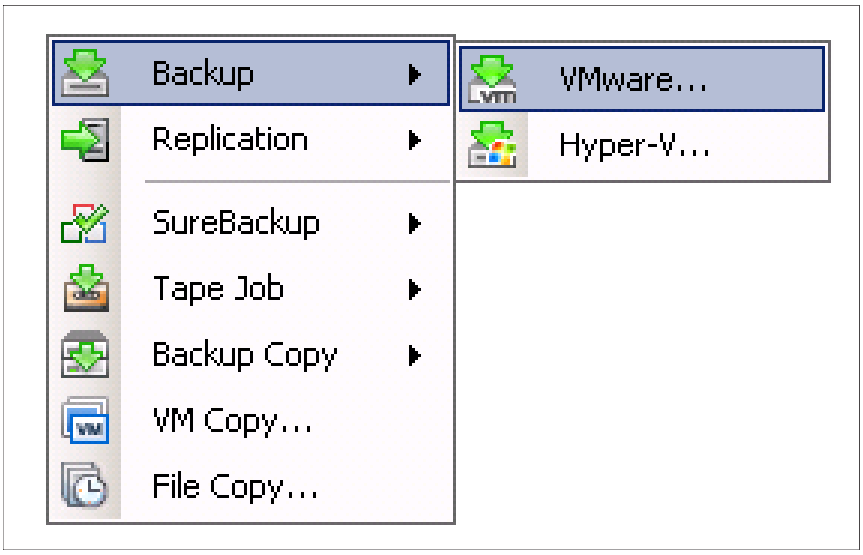

Backup to disk has many technical advantages, with the primary advantages being speed and simplicity. Backup to Disk is generally considered faster than using a direct to tape process, and quicker backups mean shorter backup windows—or the possibility of protecting more data in the same window.

Along with improved speed, backup to disk offers the ability to deduplicate and compress data on the fly. Inline, source-side deduplication offers faster backups as data is only moved once, so there is less impact on the network infrastructure too. Smaller backup files, aided by compression, are also a result.

The reliability of disk systems over tape is also a consideration. Most IT workers will have heard of an instance when a tape recovery was halted due to corrupt data.

Tape

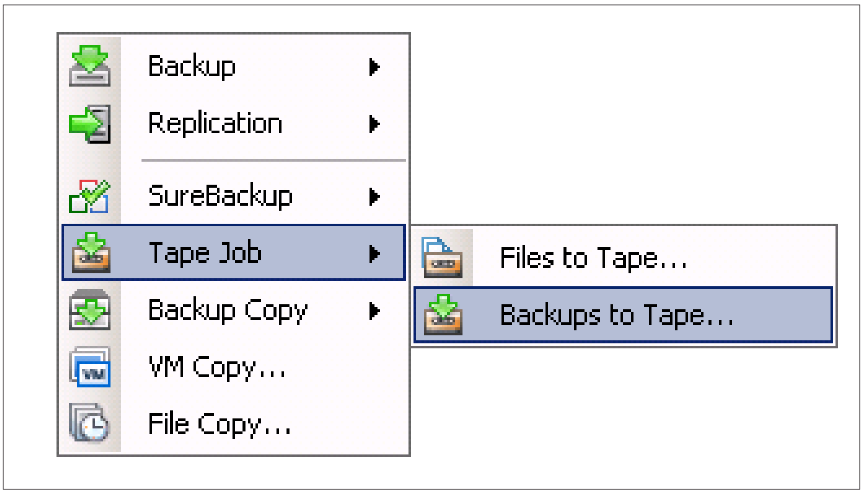

Tape has been ubiquitous as a backup medium for many years, with its lower cost being a major factor. However, tape backup technology has not been without its problems.

Backup to tape can happen in one of two ways—either direct to tape (D2T), or backup to disk, and then on to tape (D2D2T). Here we will concentrate on the latter, especially as it helps to achieve the aim of having three backup copies on at least two different types of media.

The first point to consider with a D2D2T process is how long to wait until moving the data on to tape.

If data is not being copied off site via another process, then the answer is that the backup data should be moved to tape as soon as possible, and just as importantly, the tape should then be moved off site just as quickly.

Tape can now reasonably be considered as an archive technology, so is it necessary to archive after every backup? It is common to see Veeam customers keep four weeks of backup within their primary disk backup repository, and only copy the data to tape once per week before moving the tapes off site immediately. As long as every backup is copied off site within a short timeframe

This also means that customers can choose between normal incremental backups, or benefit from using reverse incremental technology, which normally isn't suitable for D2D2T as it always creates a full backup set for every backup.

Backup Copy

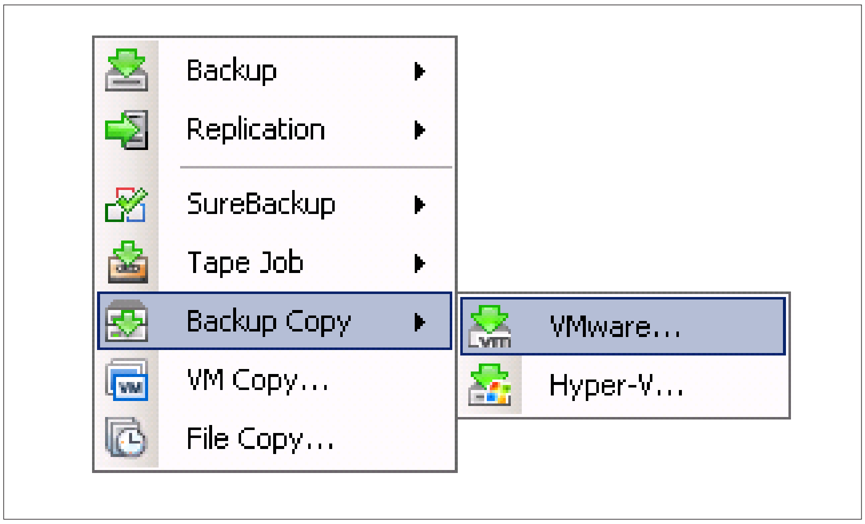

Perhaps the simplest solution to taking data off site is to be able to copy it automatically. This of course requires that the network bandwidth between sites is capable of copying each backup set across, which in turn requires that the backup changes are small enough to be copied.

Veeam Backup & Replication assists with this process by using built-in WAN acceleration to minimise the volume of data sent. This is achieved by using caches located at the source repository and target repository, along with variable block size algorithms to compare and remove redundant blocks before transferring VM data, thus significantly reducing the amount of traffic going over the network.

An organisation can choose which VMs have their backup data copied from one site to another. This allows only the data for the most critical machines to be copied, reducing the load on business critical networks.

Virtual machine replication

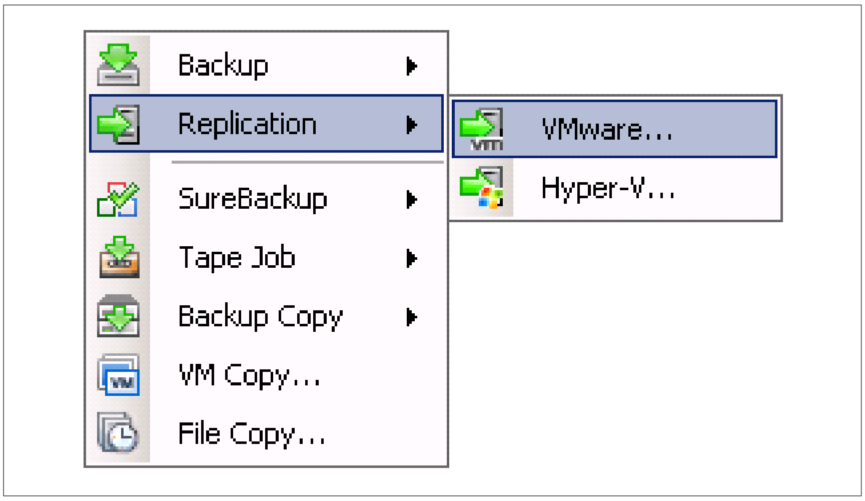

It is also possible to replicate entire VMs or parts of a virtual infrastructure. Replicated VMs offer an organisation the opportunity to have a ready-configured, warm virtual environment in an offsite location with an associated short recovery time. With Veeam Backup & Replication, we are able to replicate VMs from an organisation's VMware or Hyper-V environment onto another at remote site through the Veeam Backup & Replication proxies that are configured at each site, offering compression of the data between the sites. When recovery is necessary, Veeam Backup & Replication is able to customise the VM's IP addresses on a per-VM basis to allow for changes in the network addresses at the DR site.

Schedules

The technologies mentioned above can be scheduled to meet a number of different requirements with the backup to disk and replication technologies, allowing you to offer near-continuous data protection (near-CDP) where required or scheduling an archive to tape on a weekly or monthly basis.

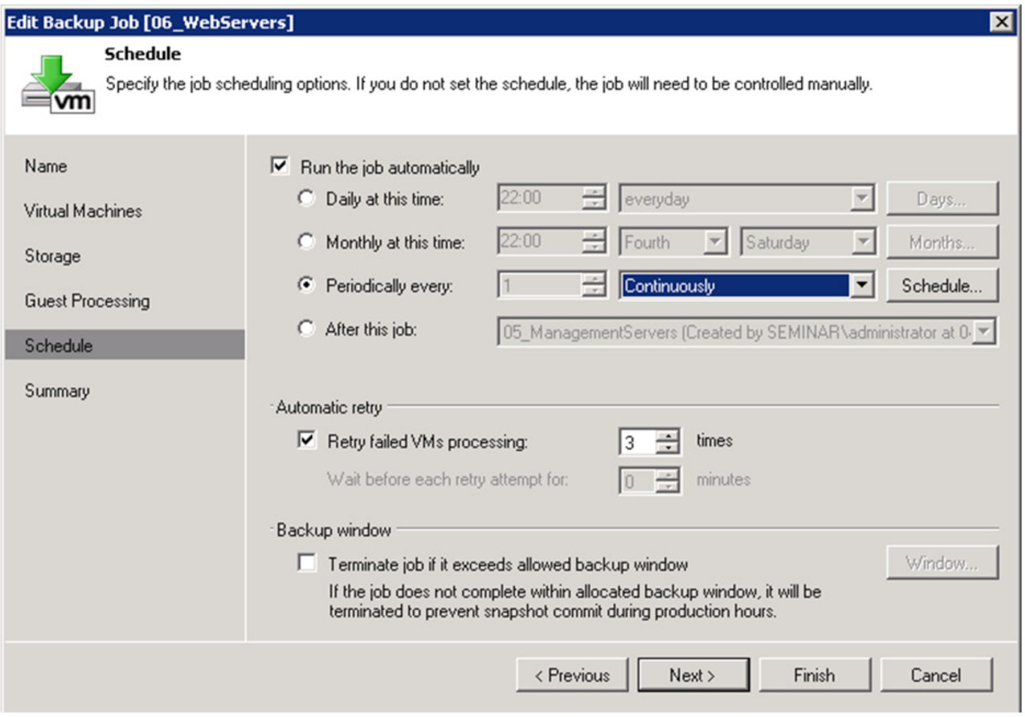

Edit Backup Job [06_WebServers]

Schedule

Specify the job scheduling options. If you do not set the schedule, the job will need to be controlled manually.

☑ Run the job automatically

○ Daily at this time: 22:00 everyday Days...

○ Monthly at this time: 22:00 Fourth Saturday Months...

● Periodically every: 1 Continuously Schedule...

○ After this job: 05_ManagementServers (Created by SEMINAR\administrator at 0

Automatic retry

☑ Retry failed VMs processing: 3 times

Wait before each retry attempt for: 0 minutes

Backup window

☐ Terminate job if it exceeds allowed backup window Window...

If the job does not complete within allocated backup window, it will be terminated to prevent snapshot commit during production hours.

Recovery options

Veeam Backup & Replication only requires one image to be taken for each backup or replica, but it offers a variety of recovery mechanisms from that single image. Different business needs require not just different backup and archiving methods, but also different speeds and methods for recovery.

The most commonly required speed for recovery is "immediately." Simple-to-use 1-Click Restore is available for both VMs and file or folders, along with Veeam's patented Instant VM Recovery™ for rapid and simple recovery of a full VM in emergencies.

Tools for easy, granular recovery for application items are also available via Veeam Explorer™ tools (for Exchange, SharePoint and Storage Snapshots) and wizards for the U-AIR® (Universal Application-Item Recovery) provide functionality for SQL and Active Directory.

SureBackup provides an automated system for testing recovery plans by ensuring that any business need for disaster recovery plans are tested in an automated fashion.

Aligning Veeam Backup & Replication with your data protection strategy

With the backup methods mentioned above, Veeam offers many options to help you achieve your defined business goals. You are able to offer different data protection and retention strategies by utilising the mix of the technologies Veeam Backup & Replication provides.

Organisations can protect their most important VMs by backing them up multiple times a day and replicating the VMs on a defined schedule to a DR site in a different location, and archiving the backups off site utilising built-in WAN acceleration and Backup Copy jobs to achieve near-CDP. With less important VMs, organisations may decide to back up the VMs once per day and archive to tape once per week.

By reviewing your business goals and service catalogues alongside Veeam Backup and Replication v7's technologies, it should be easy to create a robust data protection plan no matter the size of your business.

About the Author

Barry Coombs is a Senior Technical Architect for Computerworld Systems LTD a virtualisation focused value-added reseller by day and an avid blogger by night—following everything to do with the virtualisation, storage and cloud industries.

His responsibilities range from identifying new technologies and architecting solutions for customers to speaking and hosting customer-focused events surrounding virtualisation, storage and cloud computing.

Barry was awarded VMware's vExpert award for contributions to the VMware Community in 2010 through 2013.

You can follow Barry on his blog www.virtualisedreality.com or on twitter @VirtualisedReal.

About the Author

Chris Snell has worked in technical roles worked for backup and disaster recovery software companies for over a decade and brings a wealth of knowledge of the industry. Chris is one of Veeam's longest standing Systems Engineers and is well placed to keep you up-to-date with their latest technology.

About Veeam Software

Veeam® is Modern Data Protection™ - providing powerful, easy-to-use and affordable solutions that are Built for Virtualization™ and the Cloud. Veeam Backup & Replication™ delivers VMware backup, Hyper-V backup, recovery and replication. This #1 VM Backup™ solution helps organizations meet RPOs and RTOs, save time, eliminate risks and dramatically reduce capital and operational costs. Veeam Backup Management Suite™ provides all the benefits and features of Veeam Backup & Replication along with advanced monitoring, reporting and capacity planning for the backup infrastructure. Veeam Management Pack™ (MP) extends enterprise monitoring to VMware through Microsoft System Center and also offers monitoring and reporting for the Veeam Backup & Replication infrastructure. Veeam also provides free tools for the virtualization community.

Founded in 2006, Veeam is privately-owned and has been profitable since 2009. Veeam currently has over 20,000 ProPartners and more than 80,000 customers worldwide. Veeam’s global headquarters are located in Baar, Switzerland and has offices throughout the world. To learn more, visit http://www.veeam.com.