

Conversational Business Continuity and Disaster Recovery for Health Care



Sponsored by **veeAM**



Learn about:

- The value of business continuity and how to execute a plan
- Health care specific BCDR concerns and focus points

MINI
Edition

By **Wayne Dipchan**
(Senior Technical Infrastructure Architect, Technical Author/Speaker)

Sponsored by Veeam

Veeam® recognizes the new challenges companies across the globe face in enabling the Always-On Enterprise™, a business that must operate 24.7.365. 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 new solution that delivers high-speed recovery, data loss avoidance, verified recoverability, 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, while always supporting the current and future business goals of Veeam customers.

Founded in 2006, Veeam currently has 47,000 ProPartners and more than 242,000 customers worldwide. Veeam's global headquarters are located in Baar, Switzerland, and the company has offices throughout the world.

The Veeam logo consists of the word "VEEAM" in a bold, green, sans-serif font. The letters are spaced out, with the 'V' being significantly larger than the other letters.

To learn more, visit

www.veeam.com .

Conversational Business Continuity and Disaster Recovery for Health Care (Mini Edition)

by Wayne Dipchan

© 2017 Conversational Geek



Conversational**Geek**

Conversational Business Continuity and Disaster Recovery for Health Care (Mini Edition)

Published by Conversational Geek Inc.

www.conversationageek.com

All rights reserved. No part of this book shall be reproduced, stored in a retrieval system, or transmitted by any means, electronic, mechanical, photocopying, recording, or otherwise, without written permission from the publisher. No patent liability is assumed with respect to the use of the information contained herein. Although every precaution has been taken in the preparation of this book, the publisher and author assume no responsibility for errors or omissions. Nor is any liability assumed for damages resulting from the use of the information contained herein.

Trademarks

Conversational Geek, the Conversational Geek logo and J. the Geek are trademarks of Conversational Geek®. All terms mentioned in this book that are known to be trademarks or service marks have been appropriately capitalized. We cannot attest to the accuracy of this information. Use of a term in this book should not be regarded as affecting the validity of any trademark or service mark.

Warning and Disclaimer

Every effort has been made to make this book as complete and as accurate as possible, but no warranty or fitness is implied. The information provided is on an “as is” basis. The author and the publisher shall have neither liability nor responsibility to any person or entity with respect to any loss or damages arising from the information contained in this book or programs accompanying it.

Additional Information

For general information on our other products and services, or how to create a custom Conversational Geek book for your business or organization, please visit our website at ConversationalGeek.com

Publisher Acknowledgments

All of the folks responsible for the creation of this book:

Author:	Wayne Dipchan
Project Editor:	J Peter Bruzzese
Copy Editor:	John Rugh
Content Reviewer(s):	Karla Reina

The “Conversational” Method

We have two objectives when we create a “Conversational” book: First, to make sure it’s written in a conversational tone so that it’s fun and easy to read. Second, to make sure you, the reader, can immediately take what you read and include it into your own conversations (personal or business-focused) with confidence.


“Geek in the Mirror” Boxes

We infuse humor and insight into our books through both cartoons and light banter from the author. When you see one of these boxes it’s the author stepping outside the dialog to speak directly to you. It might be an anecdote; it might be a personal experience.



Within these boxes I can share just about anything on the subject at hand. Read ’em!

Business Continuity and Disaster Recovery for Health Care



How confident are you with your Business Continuity and Disaster Recovery (BCDR) plan?

If you had to push the button today, right now, could you have your workloads up and running within the agreed recovery time (RTO) and recovery point (RPO) of restoration? Do you have a clear understanding of what the agreed RTPOs are for any given application within your company? A rock solid,

BCDR plan is crucial for the survival of your business and can even mean the difference between life and death in a health care institution.

Understanding the direct link between technology and patient care should motivate you to ensure you have a BCDR plan you're confident in. Physicians, patients and staff rely on up-to-date patient information so they can make educated decisions and provide best patient care. This information flows from system to system, from the moment the patient is registered to discharge. All this data flowing from one system to another may be overwhelming when developing a solid BCDR plan.

We're all busy with day-to-day operations that keep Business Continuity and Disaster Recovery planning on our to-do list.



Some call that busy mentality
KTLO (Keeping the Lights On).
Trying to architect for the future
but too busy with KTLO.

Having a BCDR plan is similar to having a survival kit packed and ready for an emergency. You only need, or think about it, when a disaster strikes.

Aside from natural disasters, there are other types of events that can potentially bring your systems down, including hardware failure, cyber-attacks, and Operating System (OS) patches that conflict with the application, to name a few.

Let's say, you already have a BCDR plan. As time goes on, systems and technology change. Is your plan still valid? Your production workloads change over time. So should your BCDR plan.

In this book, we'll examine factors you need to consider when putting together a BCDR plan. We'll show you that developing a BCDR plan doesn't have to be a daunting task. You will soon be well on your way to a BCDR plan you can trust when disaster strikes.

What is Business Continuity?

The term “Business Continuity” refers to a continuation of your critical business applications during a disaster or outage.

With a perfectly implemented BCDR plan, your end users won't notice a disruption in their service.

Keep in mind, there are applications that have high availability and stay up no matter what. But for apps covered by a BCDR plan there is normally an outage as the disaster occurs. The plan then springs into action and minimizes the outage's duration.

The basic steps to achieve this goal are the same for companies of varying sizes and verticals. It's crucial that BCDR infrastructure be implemented at the point an application is deployed. If BCDR is baked into the project lifecycle and created proactively, going forward, all application deployments will adhere to your plan. Having a team dedicated to scalable, process-driven BCDR planning, testing, and improving is key to success.

Planning Steps for BCDR

First, put in place a governance committee to evaluate the criticality of each application.

Define SLA's for each level of your applications, based on their criticality. Often companies will define these SLA's in tiers. Each tier will have agreed upon Recovery Time Objectives (RTO) and Recovery Point Objectives (RPO).

RTO – The amount of time it takes to recover an application, starting from the time a disaster is declared to when users can log on and use the application.

RPO – The point in time that an application's data is restored counting backwards from the time of the disaster. This can be quantified by thinking about the amount of data in time that you are willing to lose. For example, a 15 minute RPO means that post recovery you will lose up to 15 minutes of data.

Once the tiers are defined, the governance committee should discuss each application in your environment and decide which tier they should be assigned to. The committee should include

management from both application and infrastructure departments. Any new application being introduced to the environments should be assigned a tier before being deployed.

Your tiers could be defined as follows:

- Tier 0 with RTO 0 minutes and RPO of 0 minutes
- Tier 1 with both RTO and RPO up to 15 minutes
- Tier 2 with RTO up to 4 hours and RPO up to 24 Hours
- Tier 3 with RTO up to 1 week and RPO up to 1 week
- Tier 4 Best Effort



In my experience getting key players from the application and infrastructure teams to agree has been the most difficult part of the plan.

Infrastructure services such as network connectivity, Active Directory, DNS, DHCP, etc. all need to be accounted for and assigned a tier. In many cases, the interface engine should also be considered as infrastructure and assigned an appropriate tier.

In a health care environment, it's important to consider the dependencies between applications when deciding which tier they should fall into. Typically, EMR systems have multiple inter-dependent applications. They are typically grouped as a suite of applications that would all fall into one tier.

Applications can be categorized into classifications such as:

- Mission Critical: Clinical, a direct impact on patient care
- Critical: Clinical, an indirect impact to patient care
- Essential: Possible financial impact but no impact on patient care
- Non-Essential: Business can run without these applications for some time without major disruption to end users. For example, archival or historical records.

The governance committee must be careful when deciding in which tier to place the applications. On the surface it feels like every application is critical

and can tolerate 0 downtime, however there are many factors to consider in order to achieve the most efficient BCDR plan.

Second, identify what technology needs to be in place to achieve your defined tier RPO's and RTO's.

This can be overwhelming as there are a plethora of options to choose from. The choice may rely on the skill sets you have in-house, or you may have had exposure to some trusted solutions in the past. These decisions must be balanced with a consideration of emerging technology that may offer better solutions.

For the purposes of this book, let's assume you have at least two data centers. These may be buildings owned by the company or rented rack space in a co-location data center. The cloud is also becoming more and more prevalent as a data center choice for companies. Even if you are using on-premises datacenters, you should look at tools that are cloud-ready. This will give you the option to leverage the benefits of the cloud in the future.

Specific regulatory requirements need to be met when dealing with patient information. As a professional in the health care field, you are likely very familiar with the Health Insurance Portability and Accountability Act (HIPAA). The standards set forth by HIPAA are to primarily protect sensitive patient data. These standards must be applied in many areas in technology:

- Physical Safeguards – Access to the datacenter racks or workstations
- Technical Safeguards – Account governance on the operating system or individual applications
- Technical Policies - Focus on the integrity of data. Is the data protected from being destroyed and/or altered? BCDR strategy and offsite backup are fundamental to insuring that any corruption or destruction of patient data can be quickly recovered.
- Network Security - Protect patient data in transmission from one system to another. Or while it is being replicated or backed up. Any software used to replicate patient information over a public network must use encryption.



If you are thinking about extending into the cloud, do not assume that all cloud providers meet the regulatory requirements for health care. Make sure they provide documentation on what standards they meet. Also ask what happens if you want to take data out of the cloud. Note ingestion/exgestion fees.

Let's apply this information to the tiers we defined earlier. This will help you see the strategy in practical terms to relate the information to your environment.

Tier 0

How can we achieve an RTO and RPO of 0 minutes? Basically, this will have to be an active/active configuration. Workloads providing a tier 0 service will need to be up online and servicing requests in all data centers. For the most part, infrastructure services will fall into this tier and they usually have

built in mechanisms to assure this active/active configuration.

Web applications can also be load balanced between data centers by using appliances such as F5 or NetScaler. A thick client application could be virtualized, then the publication of the application between all data centers could be load balanced using technologies such as Citrix and App-V. The data that supports these applications will need to be kept in synch behind the scenes possibly in SQL or Oracle DB and flat image files.

Tier 1

RTO and RPO of 15 minutes. Here is where the critical clinical applications will go as these applications have a direct impact on patient care. These are likely the applications used in the emergency department or the overall Electronic Medical Record (EMR) application.

These workloads will need to be replicated from the primary data center to any secondary data centers at intervals of 15 minutes. You are essentially making an exact copy of the workload in the

secondary data center and updating that copy with any changes every 15 minutes. Here is where you will need to be judicious in determining that an application is a tier 1 app. The connection between data centers has limited bandwidth (bandwidth size will depend on your type of connection). This connection may also be used for production traffic. Therefore, there is a need to consider the amount of traffic being replicated and how often that replication occurs. When deciding what tool to use, you should look at the Wide Area Network (WAN) replication optimization features (you may also consider throttling bandwidth used for replication on the network) and the ability to replicate only changed blocks.

Something to keep in mind is when you first set up replication, there is a real potential to saturate the bandwidth on the link between your data centers as you will be replicating the whole VM. Depending on the amount and size of the VMs, the bandwidth utilization may affect your production traffic on the link. It's recommended that you seed secondary datacenters with your VMs before turning on replication.

This will accomplish your RPO of 15 minutes or less, but what about the RTO of 15 minutes or less?

To achieve this, a well-orchestrated workflow to implement the BCDR strategy is needed. Most software tools built to perform replication include or offer an orchestration toolset that will allow you to manage, monitor, and troubleshoot replication. The toolset will allow you to easily failover and failback workloads from the primary datacenter to the secondary datacenters within the 15-minute RTO tolerance. Workloads can also be grouped by dependency, thus making sure you bring up services in the correct order.

If there are two physical data centers, the underlying hypervisor infrastructure will need to be present and running at both datacenters. There is also the option to replicate workloads into the cloud and use a disaster recovery as a service (DRaaS) offering. The cloud option alleviates the need to have the hypervisor layer sitting and waiting for a disaster so it may be a more efficient solution.



An efficient use of two physical datacenters can be reached by running Development, Test, and Quality Assurance workloads in the secondary data center.

These workloads can be shut down when a disaster is declared and thus free up the resources on the hosts to run the Production workloads. The shutdown of the non-essential workloads can be automated with the orchestration toolset.

This will keep asset utilization higher by preventing the waste of resources sitting and waiting for a disaster.

At this point, we need to delve into the network-addressing side of this solution. Usually when there are two separate data centers connected by a WAN link, there are also differing layer two subnets or VLANs in each datacenter.

This poses a problem when bringing up replicated workloads that have the IP configuration from the source datacenter embedded in the image. If the VLAN the VM resided on in the primary datacenter does not exist in the secondary datacenter, the VM will not be able to communicate with anything on the network. There are two ways to solve this issue:

1. Change the IP address information on the VM to match the VLAN in the secondary datacenter. That sounds easy enough and it would be if there are a manageable number of VMs. If there are hundreds or thousands of VMs we start to see the difficulty of manually changing the IP address of each VM. And unfortunately, you will not meet a RTO of 15 minutes. Thankfully, this task can be automated by using predefined IP address information for each VM with the secondary data center's IP scheme. An orchestration tool can be used to assign the IPs to the replicated VMs so when they are brought up in the secondary datacenter the IP is

automatically updated. This solution is ideal if there is no dependency on the IP of a workload.

2. If there is a need to keep the IP the same in the secondary data center you want to confirm whether any application is configured to communicate with the IP or DNS name. This is where a Software Defined Network (SDN) comes in. Companies such as Cisco and VMWare have developed a way to virtualize the network. The software defined network is one slice of the overall idea of moving toward the Software Defined Data Center (SDDC). SDN makes it possible to stretch layer two VLANs across WAN links so the source VLAN also lives in the secondary data center, eliminating worry about IP address issues or updating IP info on boot up.

Other Tiers

Other tiers will, for the most part, follow the same structure as explained for tier 1 applications. The difference comes with the replication intervals.

Tier 2 with a RPO of 24 hours will require a daily replication usually during non-business hours, or in the case of a hospital, whenever there are the least number of users on the systems.

Tier 3 with a RPO of one week will require weekly replication of changed blocks also during the time when there is the least number of users.

Tier 4 Best effort may not need to be recovered depending on the duration of the outage.

Finally, document and socialize an agreed upon playbook for BCDR implementation.

When disaster strikes, trying to determine who owns which process is not what you want to be doing. The BCDR team should develop a clearly documented step-by-step procedural guide to implement the BCDR plan. This document should detail what order services should be brought online, what team is responsible for bringing each service online, and how you are going to communicate that a service is online and ready to move to the next process. All teams involved should know ahead of time what has been assigned to them.



Make sure to keep a physical copy of this document as you may not have access to the soft copy during a disaster. You do not want your BCDR plan to fail before it begins, because you cannot access to electronic file on an offline application.

Granularity of Recovery

So far we have focused on a complete data center outage and the need to bring up all the applications in a secondary data center. A more likely scenario may be just one application failing for some reason and the need to bring just that one system or service up in the secondary data center. Orchestration tools make this possible. You can select the workloads you want to fail over and group workloads to fail over as one unit. Of course, you would need to keep in mind your IP strategy and make sure the service is able to communicate with other upstream or downstream services.

How to Gain Confidence

Testing is a crucial step of ensuring a good BCDR plan.

The amount of planning and staff involved in the BCDR testing depends on the size of the infrastructure. Some companies can turn the connectivity off to the primary data center and have all infrastructure and application teams sign off on their part of the recovery. This will highlight any inefficiencies with the plan or any unexpected results that can be remediated before the next test. Of course, this will need to take place outside of regular business hours. This presents a real challenge to a health care facility where regular business hours are 24 hours a day, 7 days a week. Typically, there are some departments that have more standard 9am – 5pm Monday to Friday hours. For these departments, you can schedule testing for their applications on evenings or weekends.

But what about the other 24/7 departments—which are typically most critical? For those, look to your orchestration toolset. Many of these tools provide features that enable a full BCDR test of a workload

or multiple workloads while keeping the production workload running. This is accomplished by bringing up the replicated workloads in the secondary data center within an isolated network. The isolated network will prevent duplicate name and IP address conflicts on the network. This feature allows the application team to connect to the application from within the isolated network, test the application and eventually sign off on the success of the testing. Some orchestration tools allow you to test the infrastructure piece of a fail over with the push of one button. Reports can be generated and sent to management for confirmation of error-free testing. Here again, of course, any issues encountered need to be documented and remediated and then tested again in the next BCDR testing cycle.

Key Takeaways

With the move toward SDDC and the orchestration tools available today, BCDR planning, implementation, testing, and training does not have to be a daunting task. With the right governance in place, coupled with the technology available, you can have confidence that downtime will be completely avoided should disaster strike.

NOTES



VEEAM


IT'S HERE

NEW Veeam Availability Suite 9.5

AVAILABILITY for the Always-On Enterprise

go.veeam.com/v9-5

With continuity of services being a key requirement for most health care-focused businesses, it's essential that a BCDR plan be well-thought-out and easy to execute. In this book, learn how to ensure downtime is avoided in a health care setting should a disaster strike.



About Wayne Dipchan

Wayne Dipchan (MCSE/MCDBA) has nearly 15 years of enterprise IT experience, ranging from educational organizations, to private investment banking, to health care. He is a published author and technical speaker.



ConversationalGeek®

Visit conversationalgeek.com for more books on topics geeks love.