

Conversational Business Continuity and Disaster Recovery for Higher Education



Sponsored by **veeam**



Learn about:

- Higher Ed specific BCDR concerns, regulations and focus points
- The value of and what it takes to bring business continuity to Higher Education

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 49,000 ProPartners and more than 255,000 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 closely spaced and have a slightly rounded appearance.


To learn more, visit

www.veeam.com

Conversational Business Continuity and Disaster Recovery for Higher Education (Mini Edition)

by Wayne Dipchan

© 2017 Conversational Geek



Conversational Business Continuity and Disaster Recovery for Higher Education (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.

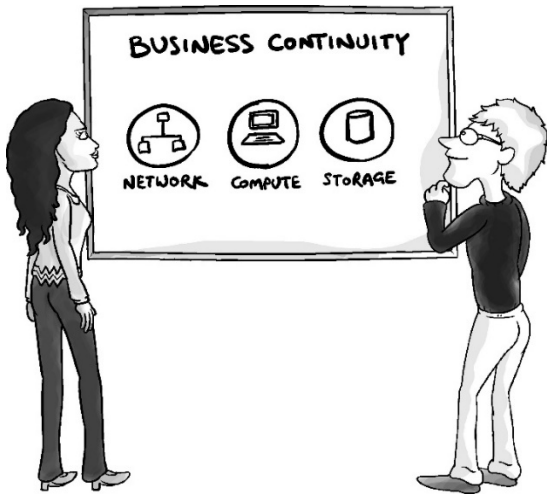
“Geek in the Mirror” Boxes

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 Higher Education



Keeping an education system's network environment running is, in some ways, more demanding than in any other type of organization. Most organizations need to simply meet the needs of their internal users, with a smaller subset of applications or services available to their external customer base. But education environments have,

by definition, the need to provide access to data and applications to faculty, staff, *and* students – each with unique needs. And this, of course, is anything but a traditional 9-to-5 type business; education environments are uniquely complex. Nearly every application – from email, to collaborative tools, to class schedules and grades, and everything in between - needs to be accessible around the clock, by just about any type of client device imaginable.

In some environments, it's almost like there are *no* tiers of applications – *everything is critical*.

All of this makes having a Business Continuity and Disaster Recovery (BCDR) plan in place so very crucial to the success of the educational institution.

Think about it – nearly any loss of service will have a major impact on a school's ability to function: registration, on-prem learning management systems, and the usual suspects like directory services, email, file and print services, etc. – an outage of any of these would bring a big part of the educational institution to a screeching halt.

So the question becomes *do you have a BCDR plan in place?* If you had to respond to a major loss of service (which we all, generally, refer to as the *disaster*) today, do you have a plan that defines the tiers of applications and services, the required recovery objectives (more on that later) for each, lists of system dependencies, and testing in place to ensure your plan is more than just a theory?

IT is pulled in so many different directions that stopping and putting a strategy and plan of any kind in place (for any strategic initiative, BCDR included) is tough. But the impact downtime can have on an educational system is too severe to risk.

The bottom line is, you need a BCDR Plan!



The bad news is your BCDR plan can't be focused on just *one* disaster, as a "disaster" can include the loss of data, systems, applications, connectivity, and locations – as well as any combination of these loss elements.

And then there are additional disaster-esque scenarios that can impact your environment. Ransomware infections today are designed to spread within a network, which can require the rebuilding or restoring of workstations and servers as part of the recovery plan. Cyber attacks involving compromised faculty or staff credentials can require reverting Active Directory back to an earlier state – which has further implications on the access to network resources by students, faculty, and staff. Incompatible patches, data breaches, malware, and phishing – just to name a few more – all can cause some degree of BCDR to come into play.

Lastly, remember that your BCDR plan isn't a static document; it's as ever-changing as the systems and technology your environment uses. As production workloads change over time, your business continuity plan needs to keep pace to ensure viability.

In this book, we'll examine factors you need to consider when putting together a BCDR plan in an educational environment. We'll also show how to proactively develop your plan, empowering IT for when disaster strikes.

Let's begin by level-setting the goal of *Business Continuity*.

What is Business Continuity?

While I'm guessing you have a better-than-rudimentary understanding of the term, let's use its most general definition, where *Business Continuity* refers to a continuation of your critical business applications during a disaster or outage, and/or changes to the business.

But defining Business Continuity in an *education setting* is a bit trickier, as you really need to figure out whether your specific operations can be categorized as *one or two* distinct sets of network operations. There are aspects of your network that are clearly devoted to the daily operations of the education of students, while other systems are all about the back-end processes that keep students in seats year after year. Some of you may use separate Active Directory instances, some may have everything staff-related on-prem with everything student/faculty-related in the cloud. It's not exactly the same in each case.

So, as we discuss the steps necessary for BCDR, keep in mind that your environment is unique and may require you to devise two separate plans.

Whether you consider everything as one environment or not, the best-case scenario of a well-planned, tested and implemented BCDR plan is for your end users to never see a disruption in their service. But, in reality, the goal is to minimize both downtime and data loss during the downtime duration – as close to zero as possible.

Achieving this goal is easy in some cases, as there are applications that have high availability and stay up no matter what. But for other applications, reaching the goal of zero disruption will take a bit more work, as disasters normally equate to an outage.

The basic steps to achieve this goal are the same for educational institutions 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 your success.

Step 1: Define BCDR Elements and Objectives

Before you ever plan, there's a bit of investigation and discovery that needs to take place to define the specific goals you have for the plan. You don't just start out planning "we'll recover system X first and then application Y" – you need to define when, how and what you wish to recover.

Your objectives should be comprised of a number of elements that will be associated together later in the process. These include:

Application Criticality

A governance committee – composed of members of IT's application and infrastructure teams, line of business owners, department heads, and even end-users – needs to consider and establish the criticality of each of your applications. Generally, organizations simply refer to “application tiers” using subjective terms where one application is more important than another. But your BCDR plan needs to have far more objective definitions that can't be misconstrued.

Because we're ultimately going to use defined levels of criticality as part of a recovery effort, the best standard is to define a service level agreements (SLA) for each application and group them into tiers. Defining an SLA is as simple as asking “how long can we afford to have this application down?” and “how much of this application's data can we afford to lose?”

Applications can be categorized into criticality classifications such as

- **Mission Critical:** Business operations come to a complete stop if unavailable.
- **Critical:** Business operations are impacted but not completely down.
- **Essential:** Possible financial impact but no impact on business operations.
- **Non-Essential:** Business can run without these applications for some time without major disruption to end users. For example, archival or historical records.

To help you categorize your applications , think of the mission essential functions (MEFs) of your organization and the applications that keep them going. Map those applications and data to operations that may cause the following: disruption of research, departure of faculty and students, well-being of students, loss of revenue, legal harm, impact on business partners or other units.

Recovery Objectives

These classifications can then be further honed down by defining specific Recovery Time Objectives

(RTOs) and Recovery Point Objectives (RPOs) for each application.

An RTO is the acceptable amount of time it takes to recover a given data set, system, or application, starting from the time a disaster is declared to when normal access is restored.

An RPO is the point in time to which the data, system, or application is restored, counting backwards from the time of the disaster.

To put these into perspective, a critical infrastructure service like Active Directory may have RTO and RPO measures in single-digit minutes (or even zero minutes), whereas the files used by the marketing department, fundraising or training may have an RTO and RPO of hours or days. Basically, the allowable downtime for each application is linked to how critical is the business unit, which must put areas such as environmental or public surety on top of your mind when it comes to recovery.

Recovery Tiers

Once the objectives are defined for all applications, you will begin to see how similarly critical

applications have similar or the same recovery objectives. The governance committee should group like applications into tiers. 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

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 when you are knee-deep in a full disaster with nothing running, the most efficient BCDR plan accounts for a systematic recovery of needed services, on-by-one.



While infrastructure services such as network connectivity, Active Directory, DNS, DHCP, etc. all need to be accounted for and assigned a tier (likely *tier 0*), because these are foundational elements on which every other application relies, you should consider having these services *already replicated to a recovery environment*. This will allow failover and recovery of impacted data and applications to be quicker and more efficient when a disaster occurs..

Dependencies

Many applications have interdependencies. Take the use case of an app designed to facilitate student registration for courses. There's a front-end web application, a back end LMS, potentially the use of Active Directory for authentication. It's a complex mix of services that, should any one of them not be available, the entire service is down. So, it's important to consider the dependencies between applications when deciding which tier they should fall into. Typically, inter-dependent systems, services, and applications are grouped as a suite that would all fall into one tier. The systems, services, and applications that make up the registration front-end and LMS would be grouped together and treated as a single recovery set.

Full and incremental backups preserve apps and inter-dependent systems, and should be performed on a regular basis for files that are irreplaceable, have a high replacement cost, or are considered critical.

Step 2: Define Needed BCDR Technology

You'll note this step doesn't start with "see what kind of recoverability your current technology has." Instead, stick with the definitions of what the organization needs to accomplish during recovery, and then work to identify what it will take to reach the BCDR objectives.

Put simply, *what tech needs to be in place to achieve your defined tier RPO's and RTO's?*

While there are lots of options today from which to choose, be laser focused on the continuity needs found in those recovery objectives, scrutinizing whether in-house solutions and skill sets will do or if new solutions, services, infrastructures, and partners will be necessary. These decisions must be balanced by considering emerging technologies that may offer better solutions. Use of virtualization is a given, as is leveraging the cloud (whether for backup storage, as a recovery target, or both).

This is the part where IT may end up saying goodbye to the other non-IT members of the committee, because the conversation will quickly turn to

technical arguments like whether a tier 1 application should simply be recovered using image-level backups, or if it needs to be replicated to a co-location data center.



If you are thinking about extending into the cloud, do not assume that all cloud providers meet the regulatory requirements for higher education. 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.

Also consider any regulations on data retention. For example, a Texas law requires its universities to store ERP data on *physical tape*. Yep – *tape*. So, as you plan for using the latest and greatest tech to execute your BCDR strategy, regulations may offer additional color that may alter the BCDR path.

Step 3: Build the Actual BCDR Plan

Now that you have a clear definition of your application and data pecking order, the criteria defining whether *continuity* or *recovery* is a choice for each, and an idea of what technology is needed to facilitate your BCDR objectives, it's time to build out a plan that includes the following:

- Mapping applications criticality, depending on the criticality of the business units/functions
- The definitions of the recovery tiers with objectives, keeping in mind the allowable downtime
- A list of where each data set, system, or application sits within these tiers (including dependencies)
- Documentation around any replicated applications or services running in preparation for a disaster

- Definitions of specific cloud-based services to be used, including support contact details, credentials, and which tiers will leverage each service
- Lastly, the recovery steps to be taken for each data set, system, or application



Build your recovery steps so they work in both a scenario where you are recovering everything, as well as one where you are recovering a single application

Step 4: Test, Test, Test

I heard it once said, “my BCDR plan is only as good as the paper it’s on.” Without testing, this is absolutely true. Once you have an idea of how simple or complex an actual recovery will be for any given recovery set, define plans to test the recovery, with increasing frequency as criticality approaches tier 0.

The amount of planning and staff involved in the BCDR testing depends on the size and complexity of the infrastructure required. Some companies can turn the connectivity off to a 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.

But if you have applications that need to be available 24 hours a day, 7 days a week, you’ll need to look to a *recovery 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 replicated workloads in a secondary data center within an isolated network. The isolated network prevents duplicate name and IP address conflicts on the production network. Application teams can connect to the recovered 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

Education environments have the unique need to keep two environments running continuously, making BCDR seem daunting. But with proper (and proactive) planning, implementation, and testing, BCDR becomes a part of daily IT operations, with

staff and technology ready at a moment's notice to ensure the highest levels of uninterrupted services. With the right governance in place, coupled with the technology, you can have confidence that you can execute a BCDR plan that meets your agreed upon recovery objectives regardless of the disaster.

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 an essential requirement for the higher education sector (universities and schools) by those who reside within that environment (teachers and students), it is essential that a BCDR plan be in place to ensure downtime is avoided 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.