

Conversational Predictable Recovery

Nick Cavalancia (Microsoft MVP)



Learn about:

- Why solid backup isn't enough and why you need predictable recovery
- The 8 things that make your recovery predictable

2nd
MINI
Edition

Sponsored by

COHESITY

Sponsored by Cohesity

Cohesity radically simplifies data management.

We make it easy to protect, manage, and derive value from data – across the data center, edge and cloud. We offer a full suite of services consolidated on one multicloud data platform: backup and recovery, disaster recovery, file and object services, dev/test, and data compliance, security, and analytics – reducing complexity and eliminating **mass data fragmentation**. Cohesity can be delivered as a service, self-managed, or provided by a Cohesity-powered partner.

COHESITY

For more information, visit
www.cohesity.com

Conversational Predictable Recovery (2nd Mini Edition)

by Nick Cavalancia

© 2021 Conversational Geek



Conversational**Geek**[®]

Conversational Predictable Recovery

(2nd Mini Edition)

Published by Conversational Geek® Inc.

www.ConversationalGeek.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 www.ConversationalGeek.com.

Publisher Acknowledgments

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

Author:

Nick Cavalancia

Project and Copy Editor:

Pete Roythorne

Content Reviewer(s):

Raj Dutt

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!

The Need For Predictable Recovery



“Will my recovery work?”

The modern-day expectation put on businesses is one of operational availability. Every employee, contractor, partner, and customer expects that when they open that app, visit that website, or launch that application it’s going to work. Anything

less than that is pretty much unacceptable by today's standards.

This need for business uptime is driving an "always-on" IT environment where backups are no longer the focus; instead the needs of the business translate into recovery service level agreements (SLAs) and define recoverability for an organization.

This shift in thinking has changed the way IT operates backup and recovery; recovery has become a daily and continual part of operations, where SLAs dictate that IT be 100% confident in the recovery outcome.

It's this need for a *predictable* recovery that's so critical; it's all well and good to have a recovery plan in place, but organizations that realize and take the "always-on" expectation seriously should be looking for ways to ensure that their recovery process doesn't just feel like "well, it *should* work", but is firmly sitting in the "it's *definitely* going to work" bucket.

The Recovery Challenge Today

Despite the advances in technology, services, and the cloud, there are still plenty of reasons why organizations need to have a means of recovering everything from a single file, to an entire environment – and everything in-between. The top three critical events experienced by organizations¹ can leverage recovery as a means of remediation:

- Cyberattacks
- Hardware or software failure impacting a business-critical system or application
- Extreme weather or natural disaster

Cyberattacks – specifically *ransomware* – are properly at the top of the list, as organizations experience an average of 23 days of downtime after a ransomware attack². These increased recovery

1. Forrester, *The Biggest Trends Shaping Enterprise Risk Management In 2020* (2020)

2. Coveware, *Q1 2021 Quarterly Ransomware Report* (2021)

times, according to Forrester, tend to lower the likelihood of fully recovering as the amount of time increases³.

With all this, IT is faced with the challenge of figuring out how to meet SLAs and recovery time and recovery point objectives that are getting increasingly smaller – in some cases down to single digit minutes (which impacts your ability to actually recover) – all while the number of workloads, systems, applications, and amount of data is growing.

All of this creates complexity, resulting in impacts on customer trust, business partners, brand reputation, organizational productivity, and – ultimately – profitability.

3. Forrester, *Ransomware Recoverability Must Be a Critical Component of Your Business Continuity Plans* (2019)



Current methods of recovery simply aren't getting it done. According to the same Forrester report, only one-quarter of organizations successfully recover 75-100% of their data after a ransomware attack!

So, it's time for IT organizations to completely shift their approach to recovery, from just planning something that *can* recover the environment, to a method that *will* recover when you need it and where you need it, without compromise.

Ways Organizations Attempt to Get Recovery to be Predictable

There are lots of traditional – and somewhat “legacy” – means of achieving levels of recovery predictability:

- **Backup verification** – some backup and recovery software solutions offer the ability to bootup image-based backups and test logon or application access. This enables you to know that the backup is viable.

- **Orchestration** – usually part of a recovery solution, this helps to automate the process of recovery, ensuring that system and application dependencies are addressed through mandating an order of recovery.
- **Recovery Testing** – This process involves performing a partial or full recovery simulating a real disaster scenario. This may be performed either manually or using orchestration, depending on how you actually plan on performing your recovery.



While recovery simulation testing is the best traditional means of ensuring recovery, according to Forrester's *2021 State of Business Continuity Preparedness* report, nearly half (47%) don't do any full recovery simulation testing, limiting their recovery predictability.

The challenge with these methods, as it pertains to predictability, is that they are all based on old standards for backup and recovery. To truly achieve

predictable recovery, organizations need to adopt new methods and standards.

So, let's spend the remainder of this book focusing on what organizations need to have as part of their recovery strategy to ensure it's truly predictable.

Building a Predictable Recovery Strategy

Organizations with data centers are leading the charge to get to *predictable recovery*, as they have made significant investment into IT infrastructure. But, it's also important that businesses of all sizes and infrastructure implementations (so, data center or not) are equally thinking about achieving the same levels of predictable recovery to see their operations be as close to “always on” as possible.

According to Gartner⁴, solutions designed to backup and recover should have the following capabilities:

4. Gartner, *Magic Quadrant for Data Center Backup and Recovery Solutions* (2020)

1. **Point-in-Time Backups** – An ability to capture a copy of your data isn't exactly revolutionary. But Gartner puts in two additional requirements with this capability. First, there needs to be an ability to backup “across heterogeneous enterprise workloads”, along with – and this is *really* important – a “100% backup success rate.”
2. **Use of Secondary (Cloud) Storage** – Backup data should be automatically written out to a secondary copy residing in the cloud.
3. **Inclusion of Cloud Workloads** – Cloud platforms, infrastructure, and SaaS data should ideally be able to be included in your backup strategy with the same solution addressing your datacenter.
4. **Search** – Solutions should be able to quickly find needed datasets, regardless of where they are located.

It's likely you were checking the boxes in your head with these three requirements, measuring them against your existing solution. But, to get to

predictable, there are some very specific aspects of your recovery strategy and your backup and recovery solution – some found in the Gartner requirement list, some not – that need to both be in place *and* utilized.

And, in many ways, it's all about the details – meaning, you need to make sure all of the following are part of your recovery strategy if you want to truly achieve *predictable*.

Recovery That Meets Business Objectives

Recovery needs to be in perfect alignment with how the business defines operations and, therefore, what needs to be up and running to be considered *operational*. Recovery SLAs should be defined based on agreed upon business objectives, clearly establishing what needs to be running, within how fast a timeframe, and with how much acceptable loss of data (if any).

Immutable Backups

Ransomware attackers today know you have backups, so they take steps to find on-premises

backup files, compromise NAS and SAN devices, and take any and all measures necessary to locate and delete or encrypt backups to ensure you pay the ransom. Your recovery is obviously not predictable if you can't predict that your backups will be available and intact. Using cloud-based storage for backups is the easy answer to achieving immutable backups.

A 100% Backup Success Rate

I want to re-emphasize this, as it's critical to recovery being predictable. Now, the reality is you may never get to a perfect backup success rate. But it's important to have and use solution capabilities that empower you to automatically test backups and their viability to recover as close as technically possible to 100% of the time.

An Ability to Search Data... Everywhere

This goes beyond a simple list of completed backup jobs and, instead, is about your ability to search throughout your entire environment for data that may be needed for recovery. This can include accessing backups that reside on-premises, on

secondary storage, in an alternate data center, or in one of potentially many cloud storage instances across multiple cloud vendors and data sources. Ideally, searches should be actionable, allowing you to both find and immediately recover needed data.

An Understanding of the Health of Your Data

It's imperative that you know whether the data you've backed up is good enough to recover from. This is more than just the traditional backup verification I mentioned earlier; this is about having a comprehensive understanding of what state the data was in *before* it was backed up, and what state it will be in *after* it's recovered. Take the following examples:

- **Before** – If your organization suffered a ransomware attack, how do you know that the backups are clean for recovery? Some ransomware authors use attack loops (where ransomware sits idle for, literally, months so that you make many, many backups) in an attempt to ensure a successful infection and no recovery

recourse. Predictable recovery includes an ability (in this case) to validate backups are not infected. There's also a small possibility, depending on backup medium used, that a backup can become corrupted, making it useless.

- **After** – Once a system or application is recovered, are there any previous vulnerabilities that need, in essence, *re-patching* to bring it, once again, up to date? It's imperative that you understand what's changed *since* the backup to ensure the viability of a recovery.



According to the Coveware *2020 Q4 Ransomware Marketplace* report, those organizations that paid the ransom for the decryption key still lost, on average, 3% of their data. This is a great example of why you need to be certain that your backup data is viable!

It's also worth mentioning that this process of understanding the current health state of your backup data needs to be automated; today those

administering backup and recovery rely on manual – non-scientific – methods. Automation is likely the only means to ensure a comprehensive review of all data (before and after, as is appropriate).

Backups That Ensure Data Consistency

Data consistency revolves around two issues. The first is ensuring that you have multiple copies of your data that are consistent with one another to ensure proper recovery. The second (and more important) issue is achieving this with no single point of failure.

Replication of some kind is built into nearly every backup solution that propagates backup data to the cloud. So, I'm assuming nearly all of you are meeting the "use of secondary storage" requirement from Gartner as well as the still relevant *3-2-1 Backup Rule*.

But most backup data replication methods use a similar model – the primary node ingests the backup data, acknowledges the data is received, and *then* propagates it to the secondary node. If the

primary node fails post-acknowledgement/pre-propagation, your data is lost.

So, for example, if your current solution backs up to on-premises storage first and then performs a separate replication process to your secondary cloud storage, you run the risk that the primary on-premises storage can fail before the cloud storage is updated. Predictable recovery demands that there be a means by which data is purposely distributed in real-time to all backup storage mediums.

What's needed is a propagation method (likely as a combination of hardware and software) designed around pretty much guaranteeing that no data is ever lost.

An Ability to Recover at Scale

Some organizations may find themselves needing to recover 25, 50, or even 100 VMs as part of a single recovery effort, so predictability can't stop with just the recovery of a single system or application. To be truly predictable, you need to know that your recovery strategy can recover *at scale*, potentially including any and all files, applications, and VMs.

With slightly more than half of organizations performing full recovery simulation tests, this same fraction of organizations are the only ones with any level of assurance that they can recover at scale. Predictable recovery takes all the previously mentioned parts of your recovery strategy and puts them all to the test simultaneously.

Recovery at scale means you can not only perform a recovery *en masse*, but that you also know if you have all the needed backups, that the backups have near-100% integrity, they can be found (and used) regardless of storage location, they are virus-free and can be easily patched, and that your chosen recovery source is consistent and up-to-date.

An Awareness That You May Need Help

While the concept of backing up and recovering seems simple enough, many organizations lack the expertise and experience to properly craft a backup strategy designed around predictable recovery. Leveraging a Backup as a Service (BaaS) provider may benefit the organization by offloading the responsibility of designing, establishing, monitoring

maintaining, and verifying backups to experts in predictable recovery.

Making the Shift to Predictable

Changes in expectations relating to organizational availability have created the need for new recovery strategies that go beyond simply having defined recovery objectives measured in minutes.

Organizations need to employ predictable recovery strategies that proactively ensure backups are accessible, intact, consistent, and valuable in an effort to meet the recovery requirements that align with an organization's business objectives.

By moving towards predictable recovery, organizations increase the certainty of their ability to recover through continual testing and validation of the recovery data, and by using the right people, processes, and technology in concert to make sure recovery is possible – no matter how big the disaster.

COHESITY




Predictable and Rapid Recovery

Reduce Downtime and Stop Data Loss

cohesity.com

© 2021 Cohesity, Inc. All rights reserved. Cohesity, the Cohesity logo, SnapTree, SpanFS, DataPlatform, DataProtect, Helios, and other Cohesity marks are trademarks or registered trademarks of Cohesity, Inc. in the US and/or internationally. Other company and product names may be trademarks of the respective companies with which they are associated. This material (a) is intended to provide you information about Cohesity and our business and products; (b) was believed to be true and accurate at the time it was written, but is subject to change without notice; and (c) is provided on an "AS IS" basis. Cohesity disclaims all express or implied conditions, representations, warranties of any kind.

In our modern “always-on” IT environments, recovery has become a daily and continual part of operations, with IT staff looking for ways to be 100% confident in the recovery outcome. In this book, I’ll look at what makes recovery predictable and how you can achieve this every time.



About Nick Cavalancia

Nick Cavalancia is a Microsoft MVP, a Technical Evangelist by trade, and is a 25+ year IT veteran who regularly speaks and writes for some of today’s most recognizable companies.



ConversationalGeek®

For more books on topics geeks love visit

conversationalgeek.com