

THE BACKUP BIBLE

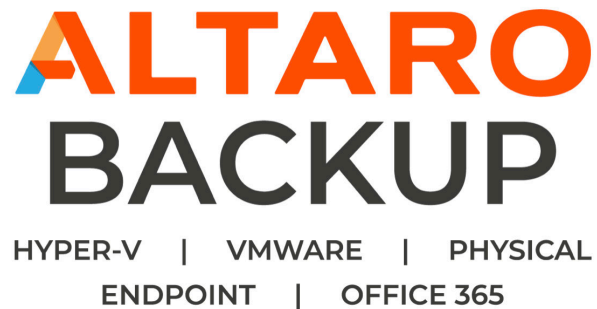
for MSPs



By Eric Siron
Foreword by SMB Nation CEO Harry Brelsford



ALTARO



HORNETSECURITY

ABOUT ALTARO

Altaro is an award-winning developer of a range of backup solutions including [Hyper-V & VMware Backup and Replication](#), [Microsoft 365/Office 365 Backup](#), and Physical Server Backup, Endpoint Backup and coming soon Azure Backup.

Altaro also offers an attractive [Partner Program for VARs](#), resellers and IT consultants as well as dedicated programs for [Managed Service Providers \(MSPs\)](#). With 50,000+ customers in 121+ countries, 10,000 partners and 2,000+ MSPs, Altaro provides affordable enterprise-class functionality coupled with outstanding 24/7 support.

[Download a free 30-day trial](#) of any of the Altaro Backup product range completely without obligation.

Altaro proudly forms [part of the Hornetsecurity group](#).

THE BACKUP BIBLE FOR MSPS

by Eric Siron

PUBLISHED BY ALTARO SOFTWARE

Copyright © 2021 by Altaro Software

<https://www.altaro.com>

Part 4: Providing Backup Services to MSP Customers written by
Andy Syrewicze and Paul Schnackenburg.

Project Editor: *Neal Storan*

All rights reserved. No part of this book may be reproduced or transmitted in any form or by any means without the prior written permission of the publisher or author.

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.

FOREWORD

by Harry Brelsford,

SMB Nation Founder and CEO



To serve and protect customers is what motivates and drives successful Managed Services Providers (MSPs). In a time where infrastructure hardware and software have attained operational maturity, that allows sophisticated MSPs to focus on high-value tasks such as backups and disaster recovery strategic planning resulting in the execution of best practices. And that's what the The Backup Bible by Eric Siron is all about.

In reality, Siron's book serves two masters: MSPs and IT Pros. But as an MSP community leader, I view it as creating the strategic framework and action plan for service providers such as MSPs. Personally, I like the journey that Siron took me on when I reviewed the book. It's a top-down format starting with the "strategy" then working all the way down to what I call "verbs" that are really action oriented tasks.

A few shout-outs to impactful sections of the book. In [Chapter 1](#), Siron leads with a [Disaster Recovery Checklist](#) that is akin to Standard Operating Procedures (SOPs) that wise MSPs adore because the solutions are scalable and replicable in their customer portfolio. Siron weaves in a behavioralist dialog across the tome but I especially liked the section on [Negative Attitudes Toward Disaster Recovery Planning](#) in [Chapter 2](#). [Chapter 4](#) is where I got my "geek on" with advanced technical conversations about fault-tolerant systems,

high-availability, clustering and asynchronous replication. In short, Siron keeps the relationship in balance between the book and the reader with both business and technical prose.

As the founder of SMB Nation, a large community of MSPs, I reached out to one very successful MSP, Anthony Oren (Nero Consulting in New York City, New York). He created Nero over a decade ago focusing on backup and disaster recovery and built his business from there to expand into cybersecurity and other fields. He viewed the book as a business plan for MSPs just starting out or hoping to convert over from the break-fix business model. I concur and proclaim to the readers to “be like Nero!” when reading and implement the advice in The Backup Bible for sure.

Something Oren took away from the book is that the user is the weakest link in backup and disaster recovery scenarios. As mentioned in this book, a business can and should leverage all of the latest and greatest backup technologies available. However, a great backup system is a process and not any single tool mostly because there is a lot more to backup and disaster recovery than just backing up or restoring files that were inadvertently deleted. Like most industries, MSPs should believe the mantra that the “customer is always right”, but unlike most industries in IT this does not translate well when customers’ computer habits are

wrong! IT folks should assume end users have good intentions but IT professionals also understand that every email that flows into a user’s inbox is just a few errant clicks away from them falling for a phishing scam that propagates a malicious software that can potentially encrypt an entire network. Ultimately, it’s the MSPs responsibility to layer their defenses with proactive remote management, timely patching and most of all having multiple backup methodologies in place to keep Bob from Accounting from opening that malware-laced spreadsheet attachment he never asked for but will open anyway.

Bottom line: backup and DR are core services for MSPs and must be done right. Read this book in chunks to absorb the full context of The Backup Bible and enjoy serving and protecting your customers!

Harry Brelsford is the founder of SMB Nation and a career technology professional, author, trainer and advisor/analyst.

CONTENTS

INTRODUCTION	10
Who Should Read This Book.....	11
The Structure of This Book.....	11
About this Book	12
Feedback	12
For MSPs	12
PART 1: CREATING A BACKUP & DISASTER RECOVERY STRATEGY	13
Chapter 1: Getting Started with Disaster Recovery Planning....	14
Chapter 2: Identifying Data Risks and Priorities	17
2.1. Negative Attitudes Toward Disaster Recovery Planning	18
2.2. Assessing the Risks that Necessitate a Disaster Recovery Strategy.....	20
2.3. A List of Common Risks	20
2.4. Determining Key Stakeholders	20
2.5. Data Prioritization	23
2.6. Widen the Search for Essential Data.....	23
2.7. Wrapping Up Risks and Priorities	24

Chapter 3: Recovery Objectives and Loss Tolerances	25
3.1. Establishing Recovery Time Objectives.....	26
3.2. Establishing Recovery Point Objectives	27
3.3. Defining Retention Policies.....	28
3.4. Adjusting RTOs, RPOs, and Retention Policies to Match Practical Restraints	30
3.5. Reviewing Recovery Objectives	31
Chapter 4: Translating your Business Plan into a Technical Strategy	32
4.1. Discovering the Technological Capabilities of Data Protection Systems	33
4.2. First Line of Defense: Fault-Tolerant Systems	34
4.3. Common Fault Tolerant Systems	35
4.4. Second Line of Defense: High Availability	41
4.5. High Availability with Clustering (Microsoft Examples)	41
4.6. High Availability with Asynchronous Replication	43
4.7. The Universal Fail-Safe - Backup	44
4.8. Closing the Planning Phase	46

PART 2: BACKUP BEST PRACTICES IN ACTION	48
Chapter 5: Exploring Backup and Disaster Recovery Technologies ..	49
5.1. Choosing the Right Backup and Recovery Software	50
5.2. Backup Application Features	51
5.3. Trial and Free Software Offerings – What to Look for	52
5.4. Placement of Backup Software.....	53
5.5. Consistency and Application-Awareness	54
5.6. Hypervisor-Aware Backup Software.....	55
5.7. Standard Physical Systems Backup Software.....	56
5.8. Single-Vendor vs. Hybrid Application Solutions.....	56
5.9. Putting It in Action.....	57
Chapter 6: Backup Storage Targets.....	60
6.1. Magnetic Tape in Backup Solutions	61
6.2. Optical Media in Backup Solutions.....	63
6.3. Direct-Attached Storage and Mass Media Devices in Backup Solutions	64
6.4. Networked Storage in Backup Solutions.....	66
6.5. Cloud Storage in Backup Solutions	68
6.6. Putting It in Action.....	71
Chapter 7: Securing and Protecting Backup Data.....	73
7.1. Risk Analysis for Backup	75
7.2. Security by Redundancy	76

7.3. Using Account Control to Protect Your Backups.....	78
7.4. Encrypting Your Backup Data	79
7.5. Isolating Your Backup Systems.....	80
7.6. Putting it in Action.....	84
Chapter 8: Deploying Backup.....	90
Chapter 9: Documenting your Backup System.....	92
9.1. Documentation Procedures and Tools.....	93
Chapter 10: Defining Backup Schedules	95
10.1. Understanding How the Value of Data Affects Backup Scheduling	96
10.2. Understanding How the Frequency of Change Affects Backup Scheduling	96
10.3. Understanding How Backup Application Features Affect Scheduling	96
10.4. Putting it in Action.....	98
Chapter 11: Monitoring and Testing Your Backups.....	102
11.1. Monitoring Your Backup System.....	103
11.2. Testing Backup Media and Data.....	103
11.3. Putting it in Action.....	104
Chapter 12: Maintaining Your Systems	105
12.1. Putting it in Action.....	106

PART 3: DISASTER RECOVERY & BUSINESS CONTINUITY BLUEPRINT	108
Chapter 13: Data Recovery Across the Organization	109
13.1. Disaster Recovery Planning Beyond the Datacenter.....	110
13.2. Restoring Non-Data Services	116
13.3. Wrapping up Non-Technical Planning	119
Chapter 14: Business Continuity and Disaster Recovery Architecture	121
14.1. Using Secondary Sites for Business Continuity and Disaster Recovery	121
14.2. Ongoing Maintenance for Secondary Sites	126
14.3. Analyzing Disaster Recovery Hardware Requirements..	127
14.4. Maximizing Disaster Recovery Architecture	130
Chapter 15: Using Replication to Enable Business Continuity	131
15.1. A Short Introduction to Replication.....	132
15.2. Choosing Replication Solutions	135
15.3. Where Replication fits into your Backup Strategy	138
15.4. Considering Replication Licensing Implications.....	139
15.5. Configuring Replication	139
15.6. Maintaining a Replica	141
15.7. Correcting Problems with Replication	143
15.8. Leveraging Replication in Disaster Recovery	147

Chapter 16: The Impact and Opportunities of Cloud Solutions in Disaster Recovery	148
16.1. Unprotected Cloud Resources	149
16.2. Cloud Backup	149
16.3. Advanced Disaster Recovery Using Cloud Services	150
16.4. Considering Cloud-based Solutions.....	152
Chapter 17: Business Process for Disaster Recovery	153
17.1. Executive Declaration.....	154
17.2. Preparing and Planning for Impacted Personnel	155
17.3. Design Guidelines for Business Continuity Processes.....	159
Chapter 18: Testing Disaster Recovery Systems	160
18.1. Testing Backup Data with Restore Operations.....	162
18.2. Testing Backup Data with Automated Operations....	164
18.3. Geographically Distributed Clusters.....	166
18.4. Do Not Neglect Testing.....	171
Chapter 19: The Ongoing Process	173
19.1 Keep Others Invested	174
19.2 The Road Ahead.....	177

PART 4: PROVIDING BACKUP SERVICES TO MSP CUSTOMERS..... 178

Chapter 20: Why Should You Offer Backup Services

to Your Customers?180

Chapter 21: Considerations for Building a BaaS Practice182

21.1 What Data Should Be Protected?..... 184

21.2 Taking it Further than Just Backup/Recovery..... 184

21.3 Storage Considerations 185

21.4 Selecting a Backup Product..... 186

21.5 Pricing models..... 187

Chapter 22: Defining BaaS Service Level Agreements188

22.1 An Example: A Tiered SLA Model for BaaS..... 190

22.2 A Note Regarding Disaster Recovery Testing..... 191

Chapter 23: Regulatory Considerations193

Chapter 24: Defining BaaS Specifics Within Your MSA196

Chapter 25: Preparing for BaaS Operations.....198

25.1 Training 199

25.2 Defining Processes 199

25.3 The Onboarding Process..... 200

25.4 Defining Backup Schedules 201

25.5 Processes for handling restoration jobs..... 201

25.6 Auditing Operations 202

25.7 The Off-boarding Process..... 202

Chapter 26: BaaS MSP Technical Operations204

26.1 Deploying Backup Operations at Scale 205

26.2 Monitoring and Reporting on Backup at Scale 208

26.3 Backup and Maintenance Operations at Scale 213

26.4 Providing Backup Access to Customers 214

26.5 Returning Backup Data to Customers 216

26.6 Technical Operations Review..... 218

Chapter 27: Never the End219

GLOSSARY 221

APPENDIX - LIST TEMPLATES AND CHECKLISTS 225

Checklist for Meetings in the Planning Phase

Risk Identification

Key Stakeholders

Data Protection Questionnaire

Information Technology Department Checklist

Desired Protection Technologies

Sample Table of Backup Application Test Results

Backup System Deployment Checklist

Sample Backup Documentation for a Small Organization

Sample Backup Documentation for a Larger Organization

Site Description

Customer Backup Requirements Template

ABOUT THE AUTHOR 248

ALTARO DOJO 251

INTRODUCTION



Humans tend to think optimistically. We plan for the best outcomes because we strive to make them happen. As a result, many organizations implicitly design their computing and data storage systems around the idea that they will operate as expected. They employ front-line fault-tolerance technologies such as RAID and multiple network adapters that will carry the systems through common, simple failures. However, few design plans include comprehensive coverage of catastrophic failures.

Without a carefully crafted approach to backup, and a strategic plan to work through and recover from disasters, an organization runs substantial risks. They could experience data destruction or losses that cost them excessive amounts of time and money. Business principals and managers might even find themselves facing personal liability consequences for failing to take proper preparatory steps. At the worst, an emergency could permanently end the enterprise.

This book seeks to guide you through all stages of preparing for, responding to, and recovering from a substantial data loss event. It contains instructions and examples for building and maintaining a thorough data protection system. It also includes discussion material on expanding recovery processes throughout the enterprise.

WHO SHOULD READ THIS BOOK

This book was written for anyone with an interest in protecting organizational data, from system administrators to business owners. It explains the terms and technologies that it covers in simple, approachable language. As much as possible, it focuses on the business needs first. However, a reader with little experience in server and storage technologies may struggle with applying the content. To put it into action, use this material in conjunction with trained technical staff.

THE STRUCTURE OF THIS BOOK

The chapters of this book proceed sequentially through three major sections:

- **Part 1: Creating a Backup and Disaster Recovery Strategy** – discussing the planning stage, how to assess your backup requirements and how to locate the necessary information
- **Part 2: Backup Best Practices in Action** – The deployment of the strategy laid out in Part 1 and putting those plans into reality.
- **Part 3: Disaster Recovery and Business Continuity Blueprint** – how to deal with the inevitable data loss events.

Where all the hard work pays off.

- **Part 4: Providing Backup Services to MSP Customers** – how Managed Service Providers can translate a Backup and DR strategy into a reliable and profitable offering”



The Appendix contains number of templates and checklists that relate to specific sections of this book enabling readers to easily put the theory into action.

Most industry-specific words are defined in text however, if you encounter any unfamiliar terms, check **the Glossary**.

ABOUT THIS BOOK

This book has been designed to cover the theory and practical exercise of backup and disaster recovery planning. It contains references to Altaro Backup solutions where relevant, but it is not specifically about Altaro or a guide to using Altaro products. As such the information contained here is relevant to whichever backup software you choose to deploy.

FEEDBACK

What do you think of this book? We want to hear your feedback so we can continue to improve this publication. Please contact us on dojo@altaro.com with your comments/feedback/suggestions.

FOR MSPS

This version of the Backup Bible has been updated to include an additional section, specifically for Managed Service Providers. It contains all parts of the normal version of The Backup Bible plus [“Providing Backup Services to MSP Customers”](#). This part was co-authored by [Andy Syrewicze](#), [Paul Schnackenburg](#) and [Eric Siron](#).

PART 1: CREATING A BACKUP & DISASTER RECOVERY STRATEGY



CHAPTER 1: GETTING STARTED WITH DISASTER RECOVERY PLANNING



A solid [disaster recovery \(DR\)](#) plan needs time and attention to properly form. Usually, the total investment closely coincides with the size and scope of the organization. Due to the level of effort, many businesses need help improving their process beyond regular [backups](#). Some also struggle with finding a logical starting point.

To get started, you need to build a checklist. It should include clear goals and the activities that will achieve them. You will need to create a custom list that fits your particular needs. You will likely need to refine the list items as you work through it.

You can use the following example checklist as a starting point.

You may not recognize all of the terms used in this list just yet, but you will find definitions later in the book and in the glossary. We will explore each of these topics in further detail as we work through the book. For now, here are the essential items you will need to include in your backup and disaster recovery checklist.

DISASTER RECOVERY CHECKLIST

- ☑ Make the business case for a disaster recovery plan
- ☑ Identify risks
- ☑ Determine key stakeholders
- ☑ Define a data prioritization strategy
- ☑ Discover data protection scope
- ☑ Define recovery objectives and tolerances (RTOs and RPOs)
- ☑ Determine solutions
- ☑ Define capital and operating budgets
- ☑ Create implementation plan
- ☑ Create business continuity plan
- ☑ Create disaster recovery plan
- ☑ Create test plans
- ☑ Create review plan
- ☑ Follow the implementation plan
- ☑ Schedule and follow the review plan

**“ DISASTER RECOVERY PLANNING
IS AN ONGOING PROCESS,
NOT A ONE-TIME EVENT. ”**



Your list will grow beyond this one, usually with several sub-items specific to your particular requirements. Out of all these items, the last, "Schedule and follow the review plan", may be the most important. Disaster recovery planning is an ongoing process, not a one-time event. You will do the bulk of the work during the initial planning phase, but your organization cannot simply abandon the plan after implementation.

Your first item, making the business case, usually spans a few of the above items. You can easily gain acknowledgment of the importance of backup, but organizational commitment is essential to realizing a thorough plan and then carrying it out.

CHAPTER 2: IDENTIFYING DATA RISKS AND PRIORITIES



The lack of understanding around data protection presents a serious barrier to proper planning. Some organizations fail to adequately plan simply because they do not realize its importance. Others do not feel that the danger justifies the effort. The lack of a plan presents the greatest danger of all. This chapter helps you to paint a fuller picture of the risks that a disaster recovery strategy can mitigate.

2.1. NEGATIVE ATTITUDES TOWARD DISASTER RECOVERY PLANNING

No one has conducted in-depth studies into the behaviors and attitudes around disaster planning. We do not know what percentage of organizations minimize or even skip this critical component. Most importantly, we do not conclusively know why system designers tend to reduce the importance of disaster recovery. We do have common anecdotes from individuals that have worked with companies to plan for or recover from disasters. Some reasons frequently cited:

REASONS WHY DISASTER RECOVERY FAILS



**SUCCESS BREEDS
A SUCCESS MENTALITY**



EXPENSE



TIME



SCOPE



**MISUNDER-
STANDING**



**SHORT-TERM
THINKING**

- **Success breeds a success mentality** - The longer an organization survives without experiencing a catastrophe, the less its members believe in the possibility of it occurring. As is true for humanity in general, few people tend to strongly consider emergencies until it's too late.
- **Expense** - Generally, businesses perform infrastructure computing and storage deployments in bulk. They purchase and install several components all at once. In the case of clusters and replicated storage devices, they may have no other options. Usually, planners design the functional portions first, then add in the protection schemes. As the capital expenditure sum climbs, the willpower to spend tends to decline.
- **Time** - Building and implementing a proper disaster recovery strategy requires time. Much of it requires the involvement of principals and senior staff. They may feel that they have better ways to allot their times than sitting in meetings and filling out questionnaires to prepare for an event that might never occur. They may also feel that their technology teams should focus on other endeavors.
- **Scope** - Sometimes, a backup plan does exist, but falls short of organizational needs. Taking a nightly backup certainly grants better protection than doing nothing at all, but that cannot represent the entire strategy.
- **Misunderstanding** - Even in today's world of ubiquitous technology, few people understand the differences between the datacenter and the desktop. Consumers rarely back up their personal computers or devices. They simply do not comprehend the risks. Without an experienced guide, they tend to underestimate the hazards.
- **Short-term thinking** - In most cases, improper planning results from innocent ignorance and naivete. Unfortunately, not everyone will have the organization's best interests in mind. A trusted consultant might try to win a contract by providing a cheap solution with little or no backup functionality. A less-than-scrupulous business manager might decide to check that important "under-budget" box by skimping on backup. Or, a well-meaning principal might adopt a "let's deal with that later when we have a little more money" stance – but later never comes.

As you work on your disaster recovery plan, keep all of these things in mind. Because backup and disaster recovery have no immediate benefit, you will almost certainly face resistance. You need to remain prepared to answer "why" at any time. The next section can help.

2.2. ASSESSING THE RISKS THAT NECESSITATE A DISASTER RECOVERY STRATEGY

If you study computer security, you will have heard of "threat modeling". Essentially, it means that security experts first identify potential threats. They can use that list to predict the extent of possible damage from an attack. That in turn helps them to design a clear strategy for defense and mitigation. You can use a similar approach to building backup and disaster recovery systems.

In the case of disaster recovery, the risks consist of a superset of the security threat model. Malicious actors pose one kind of threat out of many. You also must worry about hardware failures, natural disasters, and human error.

With each risk, you must consider its possible impact. What are the ramifications if an attacker steals data? What would happen

to the organization if a failed storage system causes complete data loss? What are your prospects if a flood makes your entire building unusable? What if someone deletes a critical e-mail that places your organizations in a legally vulnerable position? Each danger type presents a unique challenge for every organization.

At this point, you may only be able to draft a cursory idea of your risks. A proper assessment includes a detailed analysis. However, in all but the smallest companies, these investigations need more than one person. At this stage, you only need enough to make a solid case for spending time and capital on designing and creating a comprehensive backup and disaster recovery solution.

2.3. A LIST OF COMMON RISKS

To help you start your list of risks, consider some of the major risks that all organizations face:

- Data theft
- Physical theft
- Malicious digital attacks (ransomware, viruses, etc.)
- Rogue insiders
- Social instability

- Power failures
- Arson
- Sabotage
- Natural disaster
- Departure of critical staff

Take some time to research risks particular to your industry. You may not add anything to the list, but you might need to adjust its priorities. For instance, if your organization creates software, then "intellectual property theft" will feature prominently. If you transport commodities, then physical threats will rank higher.

This might be the point at which you create and present the business case for undertaking disaster recovery planning. If you need more material, then perform preliminary work on some or all of the other items in the checklist.



Use the [Risk Identification template](#) in the Appendix to help you build a list of all the risks to your organization you need to plan for.

2.4. DETERMINING KEY STAKEHOLDERS

Depending on your organization's size and your position within it, you may not have the authority or knowledge to conduct a deeper investigation on your own. Whatever your role, start with what you consider important. If you're a systems administrator, you may think of your e-mails or files. If you have a more general operational position, you might think of your book of business or inventory. To define a fuller plan, you need a wider view.

To gain that perspective, you will likely need the approval of your organization's executives. Properly analyzing risk requires time and attention. In the absence of an obvious threat or recent catastrophe, you will likely struggle to move this phase of the plan along.

Even people that understand the risks tend to consider it a low-priority task. Set a goal of getting the appropriate people involved in the conversation and ensuring that they have sufficient motivation and opportunity to participate.

To start the conversation, use an informal approach. Start asking things like, "Which people would know the most about our risk profile?" and, "Who has the best knowledge of what we need to protect?" Expect to need input from:



- Executives or principals
- Head and leads of IT
- Key stakeholders – these vary greatly between organizations. It might mean department heads or product owners or individuals in major roles.
- Intellectual property creators and proprietors.

With a starting list of names, you have options: individual interviews, forms, or group meetings. You may eventually use all of these things, but you will likely find that brainstorming meetings will get you the farthest in the beginning. However, the risk discovery task neatly connects with several of the following activities. Therefore, you will likely want to read ahead before scheduling anything.



Use the [Key Stakeholders table](#) to help identify the key stakeholders in your organization you'll need to consult in planning your backup strategy.

2.5. DATA PRIORITIZATION

Meetings and discussions about risk will inevitably cover the vital portions of your organization's systems. As you outline your exposure, you can take the opportunity to rank your assets. Most disaster recovery plans will encompass everything, but even in the best cases, restoration takes time. For now, do concern yourself with the rebuild order. Focus on mission-critical applications – what does the organization need for minimal operation?

At this phase, organize your priority list at its highest level. For example, instead of making line items that make sense to administrators, such as "customer database", use business-oriented labels such as "ERP system". You can work out the technical details later. Things will necessarily look different once you translate this list into an implementation document.

As you build up this list, ensure that everyone involved remembers that top priority belongs to the systems that your organization requires for operational performance. Try to avoid using terms like "critical", as not everyone will agree on the definition, and sometimes, you can function for a while without a crucial system. As an example, consider a company that transports freight. No one can dispute the importance

of keeping the electronic customer record system available, but can the operation continue without that longer than it can continue without the system that maintains contact with delivery and pickup drivers? The question to ask of every system: "What is the business impact of an outage?" For now, you may need to keep those answers short.



Use the [Checklist for Meetings in the Planning Phase](#) in the Appendix to help you plan meetings with key stakeholders and extract the necessary data.

2.6. WIDEN THE SEARCH FOR ESSENTIAL DATA

Meetings alone will not uncover everything that you need to protect. They serve as a starting point for the attendees. They will need to look within their departments. To complete the data protection model, key staff in each department must create a thorough inventory.

The search should not restrict itself to digital assets. Your organization may predate the advent of digital record keeping, or it may fall under the purview of regulations that require physical copies. [Business continuity](#) and disaster recovery will mean protecting those items as well.

2.7. WRAPPING UP RISKS AND PRIORITIES

Business continuity and disaster recovery both mean working through and after major problems, regardless of how they occur. Smaller events need different responses. For instance, you might need to [restore](#) a single database after an accidental deletion. So, you need to know how an accidental (or malicious) deletion might happen.

As you and your colleagues work through the discovery phase, you might find mitigation strategies that allow you to reduce exposure to your unique risks. Where possible, choose prevention over response. You will probably not completely remove many items from your list of concerns but take every advantage that you can.

Be mindful of course-altering events. For instance, if your organization centers around physical product in a warehouse, and a disaster completely destroys the facility and all of its contents, then you probably won't concern yourself as much with a pick-up scheduling application.

As your risk and priority models take shape, you will naturally build up an idea for the tolerances and expectations that you have in your disaster and data recovery planning. You might be able to clearly define

all of those in the same meetings. However, they often require a more detailed examination of the supporting systems. Department managers may need to break to gather input from daily operators. To help you through this portion, the next chapter covers the processes related to recovery objectives and tolerances.

CHAPTER 3: RECOVERY OBJECTIVES AND LOSS TOLERANCES



With sufficient funding and infrastructure, any system could theoretically achieve near-constant uptime through any situation. Reality dictates a more conservative outlook. To establish a workable budget and a practical plan, you will need to determine your organizational tolerances for outages and loss. This chapter explores the related terminology and processes.

In the previous chapter, you were instructed to ask about the business impact of a system outage. Now, you will need to have the key employees of each system explore the question more deeply. If the term "business impact" does not convey the desired level of urgency, ask questions such as "How much does this system cost us per hour when offline?" and "How many hours of work would we need to recover after losing one hour's worth of data?" You need to build up a set of objectives: recovery time and recovery points.

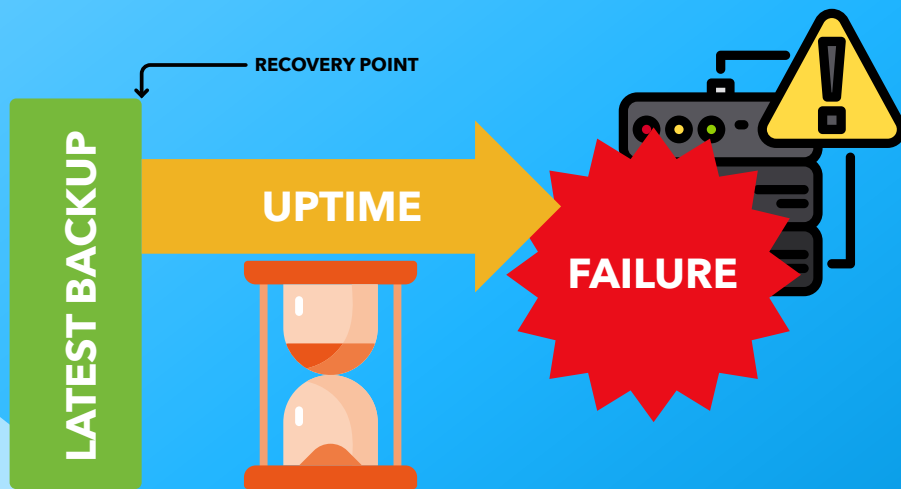
3.1. ESTABLISHING RECOVERY TIME OBJECTIVES

Use the simple question, "How long can we operate without this system?" to get your teams started. The term [Recovery Time Objective \(RTO\)](#) applies to the goals set by this enquiry. RTOs establish the desired maximum amount of time before a system returns to a defined usable state.

RECOVERY TIME OBJECTIVE (RTO)



RECOVERY POINT OBJECTIVE (RPO)



Complex systems can have different RTOs. For instance, you might set an RTO of four hours to restore a core electronic records system after a major failure but set a separate objective of one hour after a minor glitch. Your objectives might also set differing levels for acceptable functionality. Your organization might consider a functioning receipt printer in the customer service area as a meaningful success metric; it can have its own RTO as a part of a larger recovery objective.

RTOs should feature prominently on your long-term disaster recovery planning. Rely on the managers and operators of the individual systems to provide guidance. Use executives to resolve conflicting priorities. You may also need them to grant you the ability to override decisions in order to ensure proper restoration of functionality.

3.2. ESTABLISHING RECOVERY POINT OBJECTIVES

RTOs apply mainly to functionality. Events that trigger recovery actions also tend to cause data loss. Your organization will need to establish tolerance. Of course, no one wants to lose anything, which will make these discussions difficult.

Because most backups occur at specific time intervals, you use them as the basis for [Recovery Point Objectives \(RPO\)](#). An RPO sets the maximum acceptable time span between the latest backup and the data loss event. This determination coincides with the work to determine the business impact of an outage. A system's downtime not only prevents its users from retrieving or utilizing its contents, but also represents matching post-recovery work: they will need to recreate any data that was not in the backup and they will need to complete postponed operations.

You will need to establish multiple RPOs for most systems. Not all events will have the same impact, so you must set expectations accordingly. For instance, you have options for continuously created [replicas](#) and backups. Those work well as buffers against physical hardware failure. They work poorly against malicious attacks, especially encrypting ransomware. You can establish a tiered recovery approach to address the various risks. As an example:

1. First-line hardware failure or malfunction: RPO of zero hours, using continuous replication
2. Corrupted data: RPO of 1 hour after corruption detection, using on-site hourly backups

3. Site destruction: RPO of 24 hours, using off-site daily backups and cloud hosting providers

Consider the possible outcomes of each risk category as you work out RPOs. You don't just the data to restore; you need something to restore it on. If you need to acquire replacement hardware or bring in third parties for assistance, that might add time. If you have a secondary site available, add an RPO item for recovery to that location. Also include an item that addresses cross-facility challenges. Do not forget to account for availability of critical staff.

You can [read more about RTO and RPO on the Altaro DOJO](#).

3.3. DEFINING RETENTION POLICIES

Your teams have a final major decision to make: how long to keep data. These decisions are highly dependent on the nature of the business and the data. For European-based operations you may also have to consider GDPR requirements. If you don't have immediate answers, use two major guidelines:

1. Legal requirements. As an example, you may need to keep records of taxable events for a specific number of years.

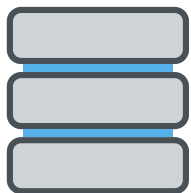
2. How long will the data have value?

Use the answers from these questions to create "retention policies".

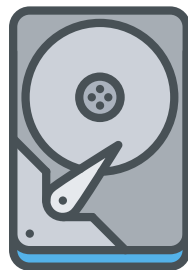
A retention policy dictates how long data must be retrievable. You will likely need more than a single company-wide policy. "Forever" may seem like an obvious answer for some things but ensure that everyone understands that data storage has an associated cost.

Data retention has two tiers: live storage and backup storage.

In disaster recovery planning, IT often only considers the backup tier. However, remember that a backup captures existing data, regardless of its age. So, if a live database has records that go back a decade, then the most recent backup contains ten-year-old information. Therefore, both your current live data and yesterday's backup satisfy a ten-year retention policy.



**RETAIN
IN ONLINE
STORAGE**



**RETAIN
IN OFFLINE
STORAGE**

To accommodate both the live and the backup tiers, retention policies must consider two things:

- Purge policies for active data
- Probability of unnoticed undesired deletion

Some electronic records systems prevent true deletion from databases without a purge action. It might move "deleted" records into a historical table or it may have a flag that removes them from visibility in client applications. Such safeguards reduce the probability of accidents. They can help against malicious deletion as well. Remember that individuals with administrative access can usually override application-level security. For the greatest safety, assume that you will not achieve your retention policies for live data. You can relax that expectation for non-critical data. Factor in the results of impact analysis from the earlier exercises. Ask, "If we lost this data forever, how would it impact the organization?".

3.4. ADJUSTING RTOS, RPOS, AND RETENTION POLICIES TO MATCH PRACTICAL RESTRAINTS

Shorter RTOs and RPOs almost always require greater financial and technical resources. Short backup intervals consume more [media](#) space and network bandwidth. Lengthy retention policies increase storage and administrative costs. Layered approaches to cover the various risk profiles can multiply those needs.



2 hour RPOs = 500 GB per day



4 hour RPOs = 250 GB per day

Backup operations place a load on the production system, which might add more strain than your current equipment allows. Replication and continuous backup technologies need more technical expertise than typical nightly backups. Staff must periodically test the validity of backup data, adding effort and overhead.

Make all of these constraints clear during early planning meetings. As executives and department heads express their wishes for speedy and RTOs and short RPOs, ensure that they understand that costs will rise accordingly. They may need to adjust their expectations to match.

Your plans will also need to factor in time and expense to re-establish infrastructure after a failure. You may need to replace physical systems. Vital foundational infrastructure, such as domain controllers, take automatic precedence over anything that depends on them. Adjust RTOs and RPOs for other systems accordingly.

The backup software that you choose will play a role in your RTO and RPO restrictions. [Altaro VM Backup](#) provides highly customizable backup scheduling options as well as [Continuous Data Protection \(CDP\)](#). You need fine-grained flexibility such as this to balance your backup needs against your available resources.

3.5. REVIEWING RECOVERY OBJECTIVES

The major activities of this chapter include input from all sectors of the business. Through interviews, questionnaires, and meetings, you can assemble an organizational view of what you need to protect. Next, you need to determine how you will implement that protection. You have not completely finished working with the non-technical departments, but you can allow them to gather the necessary data while you move to a different phase of the project. In the next chapter, you will explore the ways that you can use technology to achieve the desired disaster recovery strategy.

CHAPTER 4: TRANSLATING YOUR BUSINESS PLAN INTO A TECHNICAL STRATEGY



With the input of business-oriented personnel, you can determine how IT will deliver an appropriate business continuity design. To that end, you need to discover the capabilities of the technologies available to you. Once you know that, you can predict the costs. You can take that analysis back to the business groups to build a final plan that balances what your organization wants for disaster recovery against its willingness to pay for it.

Mapping out your backup requirements will then help you plan software subscriptions fulfil your needs. Altaro recognizes the need for multiple backup solutions and as such currently provides backup for [Hyper-V & VMware virtual machines](#), [Office 365 \(including Sharepoint and OneDrive for Business\)](#), [physical servers](#), endpoints, and with further solutions including Azure backup on the way.

4.1. DISCOVERING THE TECHNOLOGICAL CAPABILITIES OF DATA PROTECTION SYSTEMS

At this point, you have an abstract list of high-level business items. Few backup solutions target line-of-business applications. So, you need to break that list down into items that backup and replication programs

understand. To attract the widest range of customers, their manufacturers specify services and products that most organizations use.

Common protections include:

- Windows Server and Windows desktop
- UNIX/Linux systems
- Database servers
- Mail servers
- Virtual machines
- Cloud-based resources
- Physical hardware configurations

You'll need to create a map from the prioritized business-level items to their underlying technologies. Bring in technical experts to ensure that you don't miss anything. Gather input on what needs to happen in order to recover the various systems in use at your organization. Many require more effort than a simple restore-from-backup procedure.

Some examples:

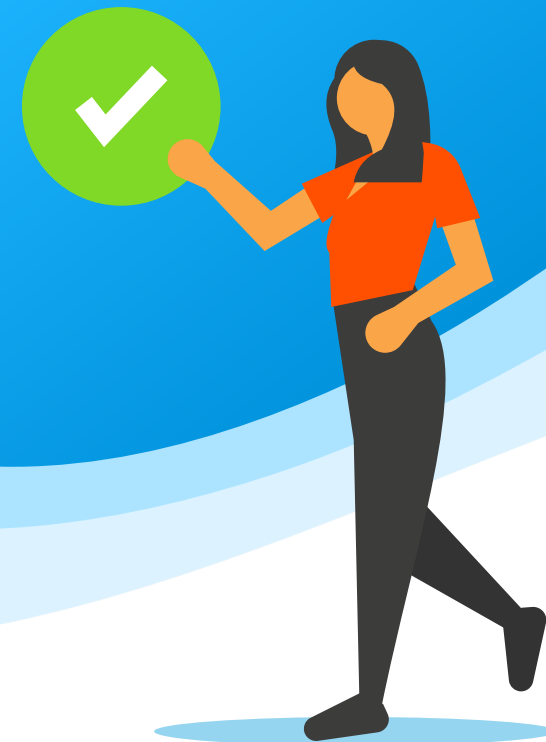
- Active Directory
- Log-based SQL recovery
- Mail servers
- Multi-tier systems
- Cluster nodes

Take input from line-of-business application experts as well as server and infrastructure experts. Seek out the experience of those that have faced a recovery situation with the systems that you rely on most. You might find exceptions or special procedures that would surprise generalists.

4.2. FIRST LINE OF DEFENSE: FAULT-TOLERANT SYSTEMS

Ideally, you would never need to enact a recovery plan. While you can never truly eliminate that possibility, you can reduce its likelihood with fault-tolerant systems. [Fault Tolerance](#) refers to the ability to continue functioning with a failed component. A truly fault-tolerant system should allow an operator to replace the deficient part and return to full operational status without any service interruption.

“ A TRULY FAULT-TOLERANT SYSTEM SHOULD ALLOW AN OPERATOR TO REPLACE THE DEFICIENT PART AND RETURN TO FULL OPERATIONAL STATUS WITHOUT ANY SERVICE INTERRUPTION. ”



Most fault-tolerant systems mostly function at low level, usually on the internal components of computer systems. To provide protection, they usually employ some method of hardware-level data duplication. In the event of a failure, they use the redundant copy to continue providing expected functionality. However, until someone replaces the defective part, the system does not provide redundancy. Further failures will result in an outage and possibly data loss.

4.3. COMMON FAULT TOLERANT SYSTEMS

Storage technologies make up the bulk of fault tolerant systems. Not coincidentally, they also have the highest failure rate. You can protect short-term storage (main system memory) and long-term storage (spinning and solid-state disks).

4.3.1. SYSTEM MEMORY FAULT TOLERANCE

To provide full fault tolerance, memory controllers allow you to pair memory modules. Every write to one module makes an identical copy to the other. If one fails, then the other continues to function by itself. If the computer also supports memory hot-swapping and technicians have a way to access the inside without unplugging anything, then a replacement can be installed without halting the system.

Of course, system memory continues to be one of the more expensive components, and each system has a limited number of slots. To make your memory installation fault-tolerant, you must cut your density in half. To compensate, you will need to buy twice as many server computers to carry the same load. Doubling the number of hosts presents more of a cost than most organizations want to undertake.

Fortunately, memory modules have a low rate of total failure. It is much more likely that one will experience transient problems which can be addressed by cheaper solutions. Server-class computer systems usually support error-correcting code (ECC) memory modules. ECC modules incorporate technologies that allow for detection and correction of memory errors. Some vendors provide proprietary technologies to defend against problems.

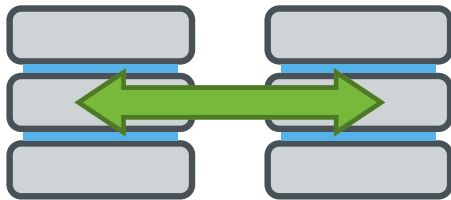
In most cases, you will choose ECC memory over fully fault-tolerant schemes. ECC cannot defend against module failure, but such faults occur rarely enough to make the risk worthwhile. ECC costs more than non-ECC memory, but it still has a substantially lower price tag than doubling your host purchase.

4.3.2. HARD DRIVE FAULT TOLERANCE

Hard drives, especially the traditional spinning variety, have a high failure rate. Since they hold virtually all of an organization's live data, they require the most protection. Due to the pervasiveness of the problem, the industry has produced an enormous number of fault-tolerant solutions for hard drives.

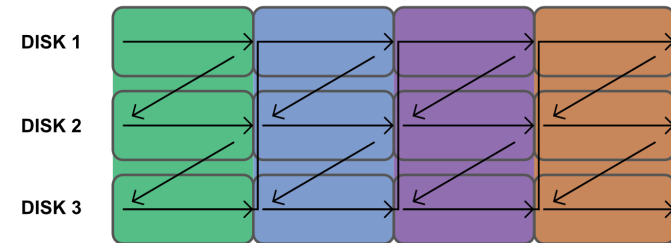
RAID (redundant array of independent disk) systems make up the bulk of hard drive fault tolerance designs. These industry-standard designs use a combination of the following technologies to protect data:

- **Mirroring** - Every bit written to one disk is written to the same location on at least one other disk. If a disk fails, the array uses the mirror(s).

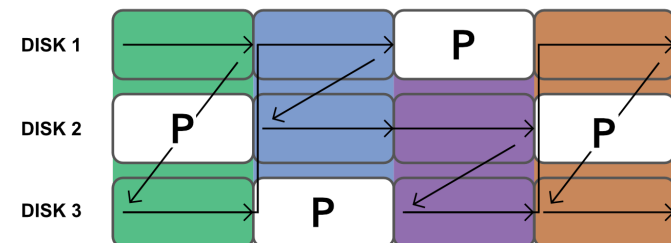


- **Striping** - A block size is set for the array. When data written to a disk fills a block, writes continue on the next disk in the array. After filling a block on the last disk, writes wrap around to a new

block in the first disk. Striping alone does not provide any data redundancy.



- **Parity** - Parity also uses a striping pattern, with a major difference. One or more blocks in each stripe holds parity data instead of live data. The operating system or array controller calculates parity data from the live data as it writes the stripe. If any disk in the array fails, it can use the parity data in place of the live data. A parity array can continue to function with the loss of one disk per parity block per stripe.



If you wish to use RAID, you can choose from a number of "levels". Each level of RAID provides its own balance of redundancy, speed, and capacity. With the exception of RAID-0 (pure striping for performance, no redundancy), all RAID levels require you to sacrifice space for protection. Disks present a relatively low expense when compared to system memory, and you have many expansion options beyond the base capacity of a system chassis. So, while RAID presents a higher cost per stored bit than single disk systems, it is usually not prohibitive.

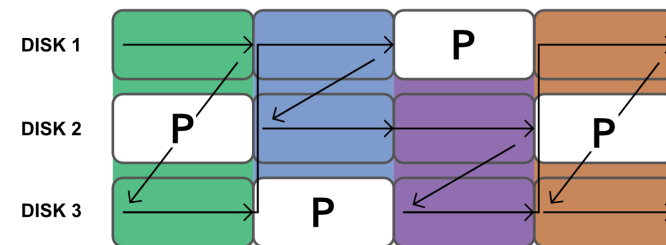
Many RAID levels have fallen out of favor due to insufficient protection in comparison to others, and some types simply consume too much space for cost efficiency. You will typically encounter these types:

- **RAID-1** - A simple mirror of two disks. Provides adequate protection, slightly lower than normal write speeds, higher than normal read speeds, and a 50% loss of capacity.



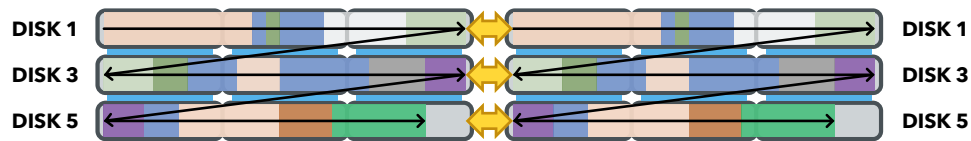
- **RAID-5** - A stripe with a single parity block. Requires at least three disks. Each stripe alternates which disk holds the parity data so that in a failure scenario, parity calculations only need to occur

for 1/n stripes. Can withstand the loss of a maximum of one disk. Provides adequate protection, above normal write speeds, above normal read speeds, and a loss of 1/nth capacity. Not recommended for arrays that use very large disks due to the higher probability of additional disk failure during rebuilds and the higher odds of a failure occurring between patrol reads (scheduled reads that look for bit failures).



- **RAID-6** - Like RAID-5, but with two parity blocks per stripe. Requires at least four disks. Safer than RAID-5, but with similar concerns on large disks. Slower than RAID-5 and a capacity loss of 2/n.
- **RAID-10** - Disks are first paired into mirrors, then a non-parity stripe is written on one side of the mirror set, which is then duplicated to the corresponding mirror disk. Can function with the loss of one disk in each mirror but cannot lose two disks

in the same mirror. Provides better performance and a higher safety rate than parity schemes, but at a loss of 50% of total drive capacity.



Due to the preponderance of drive failures and reduced performance of standardized redundancy schemes, many vendors have introduced proprietary solutions that seek to address particular shortcomings in RAID. Whereas RAID works at the bit and block levels, most vendor-specific systems add on some type of metadata-level techniques to provide protection or performance enhancements.

You have an overwhelming number of choices when it comes to fault-tolerant disk storage, so keep a few anchor points in mind:

- Storage vendors naturally want you to buy their highest cost equipment. Use planning tools to predict your capacity and performance needs before you start the purchasing process. Businesses frequently overestimate their space and performance requirements.”

- You can almost always expand your storage after initial implementation. You do not need to limit yourself to the capacity of a single chassis as you do with system memory.
- Solid-state disks have a substantially lower failure rate than spinning disks. You can leverage hybrid systems that incorporate both as a way to achieve an acceptable balance of performance, redundancy, and cost.

The most important point: downtime costs money. Storage redundancy directly reduces the odds of an unplanned outage.

4.3.3. ADVANCED STORAGE FAULT TOLERANCE

The advent of affordable, truly high-speed networking (ten gigabit and above) has brought exciting new options in storage protection. Today's networking speeds exceed even high-end storage equipment. Once the sole purview of high-end (and very high-cost) storage area network (SAN) devices, you can now acquire chassis-level, and even datacenter-level, storage redundancy at commodity prices.

These technologies depend on real-time, or synchronous, replication of data. In the simplest design, two storage units mirror each other.

" DOWNTIME COSTS MONEY. "



Systems that depend on them can either connect to a virtual endpoint that can fail over as needed, or they connect to one unit at a time in an active/passive configuration. In more complex designs, control systems distribute data across multiple storage units and broker access dynamically. We discuss real-time replication more completely in the chapter titled "Using Replication to Enable Business Continuity".

The most advanced examples of these technologies appear in relatively new [hyper-converged](#) solutions. These use software to combine the compute layer with the storage layer on standard server-class computing hardware. In most cases, they involve a hypervisor to control the software layer and proprietary software to control storage.

While costs for distributed storage and hyper-converged systems have declined dramatically, they remain on the higher end of the expense spectrum. Unlike traditional discrete systems, you will need significant infrastructure and technical expertise to properly support them.

You can consider the duplicated data in this fashion as a [hot copy](#). It's updated instantaneously and you can fail over to it quickly. Some synchronous replication systems even allow for transparent failover or active/active use.

4.3.4. APPLICATION AND OPERATING SYSTEM FAULT TOLERANCE

At the highest layer, you have the ability to mirror an operating system instance to another physical system. To make that work, you must run the instance under a hypervisor capable of mirroring active processes. It's a complex configuration with many restrictions. Few hypervisors offer it, it won't work universally, it won't survive every problem, and the performance hit might make it unworkable for the applications that you want to protect most.

At a more practical level, some applications allow a measure of fault tolerance through tiering. For instance, you can often run a web front-end for a database. You can use load-balancers that instantly move client connections from one web server to another in the event of failure. Some database servers also allow for multiple simultaneous instances that can instantly redirect connections to a functioning node. These technologies have greater functionality and feasibility than operating system fault tolerance.

4.3.5. CAVEATS OF FAULT-TOLERANCE

As you explore options for fault-tolerance, you'll quickly notice that it comes at substantial cost. Almost all of the technologies will require you to purchase at least two of everything. Most of them will necessitate additional infrastructure. All of them depend on expertise to install, configure, and maintain. Those costs always need to be scoped against the cost of equivalent downtime.

The primary purpose of fault tolerance is to rely on duplicates to continue functioning during a failure. That has a negative side effect: your fault-tolerant solution might duplicate something that you don't want. For example, if ransomware attacks your storage system, having RAID or a geographically redundant SAN will not help you in any way. Even in the absence of a malicious actor, redundant systems will happily copy accidental data corruption or delete all instances of a vital e-mail on command.

While fault tolerance will serve your organization positively, it cannot stand alone. You will always need to employ a backup solution for asynchronous data duplication. However, you have options between fault-tolerance and backup. Those technologies reside in the [high availability](#) category.

4.4. SECOND LINE OF DEFENSE: HIGH AVAILABILITY

You can't use fault tolerance for everything. Some systems have no way to implement it. Some have a prohibitively high price tag. Instead, you can deploy high availability solutions. High availability has a more nebulous definition than fault tolerant. It applies less to actual technologies and more to outcomes. Where fault-tolerance means working through a failure without interruption, high availability measures uptime against an expected metric.

As an example, your organization sets a target of 99.99% annual availability for a system that they want working at all times. To achieve that, you would need to ensure that the system does not experience more than a few minutes of total downtime over a one-year period. 365 days times 99.99% equals 364.9635 days of uptime, which allows a little less than 48 minutes of downtime per year. That's an aggressive goal, but not necessarily unachievable. Some high availability metrics only count unscheduled downtime, others apply to any downtime. These methods have radically different needs, so ensure that you know your policy.

From the technology angle, any tool that specifically helps to improve uptime falls under the high availability umbrella. All fault-tolerant technologies qualify. However, you also have some that allow a bit of downtime in exchange for reduced cost, wider application, and simpler operation. Among these, clustering is generally the most common.

4.5. HIGH AVAILABILITY WITH CLUSTERING (MICROSOFT EXAMPLES)

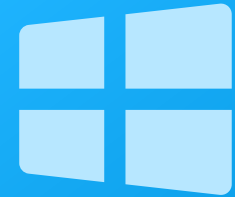
[Clustering](#) involves using multiple computer or appliance nodes, usually in an active/passive configuration, to host a single-instance resource. Some examples that depend on Microsoft's failover clustering technology:

- **Microsoft SQL** - A clustered Microsoft SQL database runs on one of many nodes. In a planned failover, the database becomes unavailable for a few seconds while its active node stops and one of the passive nodes start. In the event of active node failure, the database is offline for a few seconds while a passive node starts it. Active transactions might drop in an unplanned failover.

- **Hyper-V** - A clustered virtual machine can quickly move online (Live Migration) or offline (Quick Migration) to another node in a planned failover. If its active node fails, the virtual machine crashes but another node can quickly restart it.
- **File server** - The standard clustered Microsoft file server hosts through an active node, with planned and unplanned failovers occurring quickly. Microsoft also provides a scale-out file server, which operates in a more fault-tolerant mode.
- **Storage Spaces Direct** - Commonly called "S2D", [Storage Spaces Direct](#) is Microsoft's distributed file system offering. It can also operate on Hyper-V nodes, providing a complete hyper-converged infrastructure.

You will find clustering technologies in other operating systems, hypervisors, and physical appliances. Remember that these differ from fault-tolerance in that they allow some downtime. However, they should greatly reduce downtime risks when compared to standalone systems.

HIGH AVAILABILITY WITH CLUSTERING MICROSOFT EXAMPLES



FILE SERVER



4.5.1. CAVEATS OF CLUSTERING

Clustering provides a duplicate of the compute layer. It ensures that a clustered workload has somewhere to operate. It does not make any copies of data. Without additional technology, a critical storage failure can cause the entire cluster to fail.

Because of the necessity of hardware duplication, clustering costs at least twice as much as operating without a cluster. You might also need to purchase additional software features in order to enable a clustered configuration. Clustering requires staff that know how to install, configure, and maintain it.

You must also take care that the backup solution that you choose can properly protect your clustered resources. Solutions such as [Altaro VM Backup](#) protect Hyper-V and VMware clusters. You can sometimes successfully employ a backup solution that doesn't interoperate with your high availability solution, but it will require significantly more administrative effort.

4.6. HIGH AVAILABILITY WITH ASYNCHRONOUS REPLICATION

You can employ technologies that periodically copy data from one storage unit to another. Asynchronous replication can use a snapshotting technique to maintain complete file system consistency. Some replication applications use a simple file-copy mechanism, which works well enough for basic file shares but not for applications.

Some applications have their own asynchronous replication built in. Microsoft's Active Directory will automatically send updates between domain controllers. Most SQL servers have a set of replication options. Microsoft Hyper-V can create, maintain, and control virtual machine replicas.

You can consider data created by asynchronous replication as a [warm copy](#). It does require some sort of process to bring online after a failure, but you can place it in service quickly.

We cover replication in much greater detail in the "Using Replication to Enable Business Continuity" chapter.

4.6.1. CAVEATS OF ASYNCHRONOUS REPLICATION

Unlike clustering, asynchronous replication requires some human interaction to switch over to a copy after a failure. Clustering technologies use some sort of control technique to prevent [split-brain](#) situations in which two copies run actively and simultaneously. Most replication systems have no built-in way to do that. So, if you choose to implement replication, ensure that you plan accordingly.

Replication shares the main drawbacks of clustering: it requires duplicated hardware, special software, and expertise. It also does not protect against data corruption, including ransomware.

4.7. THE UNIVERSAL FAIL-SAFE - BACKUP

Out of all available disaster recovery and business continuity technologies, only backup is both sufficient on its own and necessary in all cases. You can safely operate an organization without any fault tolerant or high availability technologies, but you cannot responsibly skip backup.

Before you start shopping, ensure that you understand common backup terms. [The glossary](#) contains definitions for all of the following terms that are important to understanding this section:

- Full backup
- Differential backup
- Incremental backup
- Media
- Delta
- Crash-consistent
- Application-consistent
- Restore
- Rotation

As you explore backup solution choices, you need to use the plan created by your business teams as a guideline. You want to try to satisfy all requirements for data protection and retention. Consider these critical components of backup technologies:

- Backups must create a complete, standalone duplicate of data
- Backups must maintain multiple unique, non-interdependent copies of data
- Backups should complete within your allotted time frame
- Backups should provide [application-consistent](#) options
- Backups should work with the type of backup media that you want to use
- Backups should work with your cloud providers, both to protect your cloud resources and to back up to your cloud storage account(s), as desired

The above list only constitutes a bare minimum. Realistically, all backup vendors know that they need to hit these targets, so only a few will miss. Usually, those are the built-in free options

or small hobbyist-style projects. You will find the greatest variances among the last two items.

Products will distinguish themselves greatly in operation and in optional features. You should avail yourself of trial software to experience these for yourself. Some things to look for:

- **Ease of operation (especially restores)** - In a disaster, you cannot guarantee the availability of your most technically proficient staff, so your backup tool should not require them.
- **Speed of operations** - Backup and restore operations need to complete in a reasonable amount of time. However, they cannot sacrifice vital functionality to achieve that. Most backup vendors utilize some sort of deduplication technology to reduce time and capacity needs, but you absolutely must have a sufficient number of non-interdependent copies of your data.
- **Retention lengths** - Most backup applications allow an infinite number of backups – except in their free editions. If your organization won't allow you to spend money on backup software, that might prevent you from achieving their requirements.

- **Support for the products that you use** - As mentioned earlier in this book, very few backup applications know anything about line-of-business software. However, they should handle the operating systems and hypervisors that you use. Some will have advanced capabilities that target common programs, such as mail and database servers. If you choose a solution that does not natively handle your software, ensure that you know how to use it to perform a proper backup and restore.
- **Offsite support** - Because you will use backup to protect against the loss of your primary business location, your backup tool needs to have some method that allows you to take backup data offsite. Traditionally, that meant some sort of portable media. Today, that also means transmitting to an alternative location or a cloud provider.
- **Support for alternative hardware** - After a disaster, you probably won't have the luxury to restore data to the same physical hardware that it protected. Make sure that your backup application can target replacement equipment.
- **Technical support options** - Hopefully, you'll never need to call support for your backup product. However, you don't know who

might need to perform a restore. That task might fall to a person that will need help. You also need to consider future product updates and the possibility of bugs that need attention. Ensure that you understand your backup provider's support stance and process. If possible, try to talk to them before purchase.

Consider data created by backup as a "[cold](#)" copy. You must take some action to transition the data from its backup location before you can use it in production. It usually has a much higher time distance from the failure point than replication.

4.8. CLOSING THE PLANNING PHASE

You have now seen all of the basic concepts and have enough knowledge to tackle the planning phase of your disaster recovery strategy.



The [Information Technology Department Checklist](#) in the Appendix covers the major IT operations you'll need to ensure you consider in your backup strategy.



Use the [Data Protection Questionnaire](#) in the Appendix to form the basis of your investigation into your specific backup requirements based on your wider business operations and IT infrastructure.

The next chapter shifts the focus from planning to acquisition and implementation. You may wish to complete your planning phase before moving on. If you have no familiarity with the technologies used in backup and disaster recovery, then you might benefit from reading ahead.

PART 2:

BACKUP BEST PRACTICES

IN ACTION



CHAPTER 5: EXPLORING BACKUP AND DISASTER RECOVERY TECHNOLOGIES



The preceding chapters worked through the initial design phase of a business continuity strategy. Now, we transition to implementation. The connection point is usually when you have received the bulk of your hardware and software purchase and can put it to use.

If you have not even submitted orders yet, that's ideal. If you already have everything, that's fine as well. You must design the architecture, which you might find easier to perform before you decide what to buy.

In simple terms, you must move from deciding what to protect on to deciding how to protect it. For some things, your organization might choose to use printed hard copies. Those survive power outages and need no technical expertise to use. You will need to find a way to adequately keep these items safe. Consider their risk from events such as fire, flood, and theft. If the contents of the documents are vital but not a risk to security, then perhaps creating and distributing multiple copies is the best answer. Technology may not help much for these types of problems.

To guard your digital information, you need three major things:

- Backup software
- Backup storage
- Security strategy

If you start by selecting your backup application, that can guide you toward the most appropriate hardware platform and security approach. You could also start with a physical storage system that you like, but this may restrict your options for software solutions.

In the past, companies rarely put much thought or effort into backup security. Soon, they learned – the hard way – that bad actors found enough value in data backups to steal them. That prompted the backup industry to introduce security features into their products.

Later, ransomware authors began targeting backup applications to prevent them from saving victims' data.

We will tackle all these concerns across the next few chapters.

This chapter focuses on the topic of software used in backup and disaster recovery.

5.1. CHOOSING THE RIGHT BACKUP AND RECOVERY SOFTWARE

Your software selection will have monumental long-term impact on your disaster recovery and business continuity operations. Once you successfully implement your choice of application(s), inertia will set in almost immediately. Most vendors offer renewal pricing substantially

below their first-year cost, which makes loyalty attractive. Switching to another provider might prove prohibitively expensive. Even if you get attractive pricing from a competitor, you still need to invest considerable time and effort to make the switch. For these reasons, you should not rush to a determination.

At its core, every single backup application has exactly one purpose: make duplicate copies of bits. Any reasonably talented scripter can build a passable bit duplication system in a short amount of time. Due to the ease of satisfying that core function, the backup software market has a staggering level of competition. With so many available choices, you get some good and some bad news. The good news: you have no shortage of feature-rich, mature options to choose from. The bad news: you have no shortage of feature-rich, mature options to choose from.

You likely will not try out more than a few vendors before you either run out of time or become overwhelmed. In the upcoming sections, you will find many pointers to help you quickly pare down your options to a reasonable subset before installing your first trial package.



Use the [Sample Table of Backup Application Test Results](#) in the Appendix to help you map out the pros and cons of backup software before choosing the right solution.

5.2. BACKUP APPLICATION FEATURES

To distinguish themselves in a marketplace crowded with dozens of other companies trying to sell a product that performs the same fundamental role, backup program manufacturers spend a great deal of time on the supporting features. Like anyone else, they tend to brag about whatever they feel that they do especially well. So, you can often get an overall feeling about a product just by looking at its marketing literature. If they frequently use words like “simple” and “easy”, then you should expect to find a product that will not need a lot of effort to use. If you see several references to “fast” and “quick” and the like, then the application likely focuses on optimizations that reduce the amount of time to perform backup or restore operations. Businesses that work from a value angle tend to use words like “affordable” and “economical”. Words like “trusted” and “leader” tend to indicate a mature product with a dedicated following.

So, if you go to the homepage of a backup vendor and see phrases containing words that speak to you, then you are almost certainly in that company's target market. At the very least, they think that they have something to offer that fits your needs. You will have to do more work to determine if their product lives up to the promise. However, if you see nothing that addresses your primary concerns, take that as a warning sign. For instance, if you mostly want a stable product with responsive support that you can afford, you might want to avoid a company that prides itself on bleeding-edge capabilities, places its support links after everything else, and makes it difficult to even find pricing. It's important to match the scope of the solution with your unique deployment characteristics and business requirements rather than simply opt for the cheapest or most feature rich.



Use the [Desired Protection Technologies list](#) in the Appendix to ensure your backup software covers the most important technologies you'll need to protect your data.

5.3. TRIAL AND FREE SOFTWARE OFFERINGS - WHAT TO LOOK FOR

Every major backup application manufacturer offers a trial, and most offer a limited but free version of their products. You should take advantage of these opportunities. With so many quality products on the market, avoid anything that you cannot try prior to purchase.

As you test the software, use your plan from part one as your guide. If the program cannot satisfy anything on that list, then you must gauge the importance of that deficit. Find out if the program provides an alternative method to achieve the goal. If it does not, then you must choose between augmenting this program with another or skipping the product altogether.

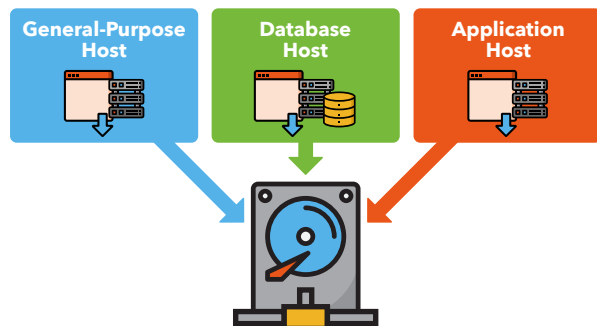
As for free software, it works perfectly well for trial purposes. However, exercise extreme caution if you intend to use it long-term. Commercial software companies need income to survive, so they invariably build their free tiers in some way that showcases the power of their software but still makes the paid tiers desirable. You can even find a few completely free programs provided by contributors out

of the goodness of their hearts. These are rarely enterprise-ready and almost never maintained for very long. In all cases, you cannot expect to receive significant support for free products. Think long and hard before deciding to entrust your organization's disaster recovery and business continuity to such tools.

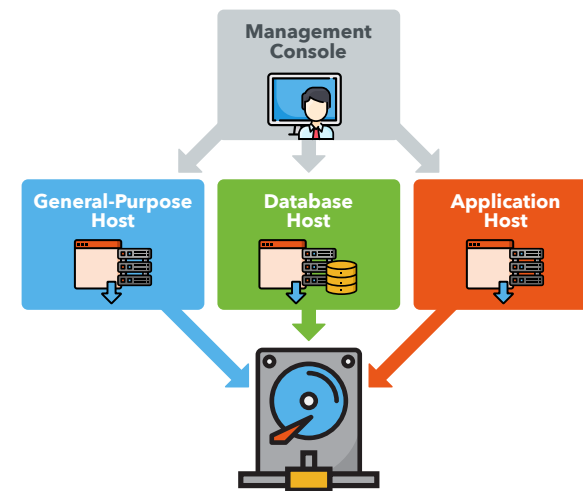
5.4. PLACEMENT OF BACKUP SOFTWARE

As you look through your software options, you will find considerable differences in deployment and management behaviors. Take note of their installation requirements and procedures. Common options:

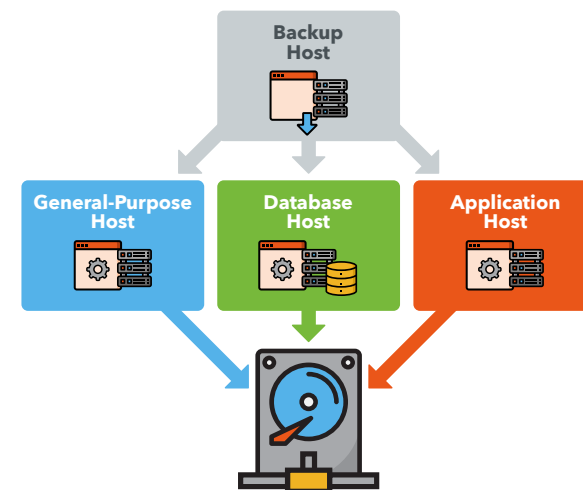
- Per-host installation, data direct to storage, no centralization



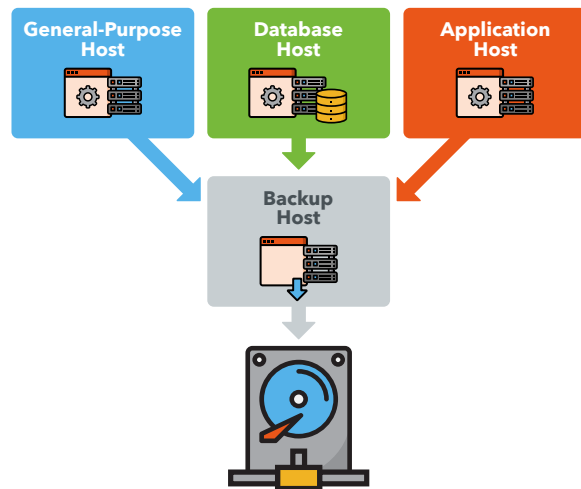
- Per-host installation, data direct to storage, managed from a central console



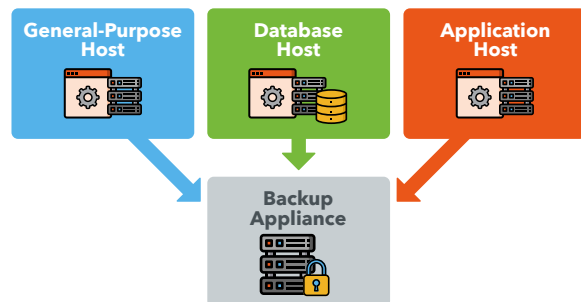
- Central installation, agents on hosts, data direct to storage



- Central installation, agents on hosts, data funneled through the central system



- Appliance-based installation, agents on hosts, data stored on or funneled through appliance

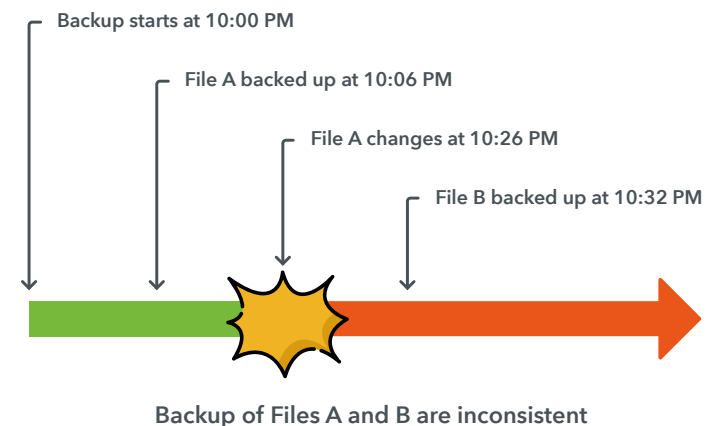


You will find other architectures. Before you purchase anything, ensure that you understand how to deploy it. If you will need to rack a physical appliance or make capacity for a virtual appliance, you do not want

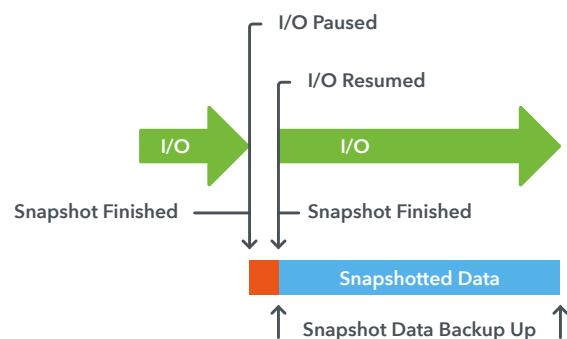
that to catch you by surprise. If your preferred program requires a dedicated server instance, that may have licensing implications beyond the backup application's cost.

5.5. CONSISTENCY AND APPLICATION-AWARENESS

In the past, we could not capture a consistent backup. Operations would simply read files on disk in order as quickly as possible. But, if a file changed after the backup copied it but before the job completed, then the backup's contents were inconsistent. If another program had a file open, then the backup would usually skip it.



Microsoft addressed these problems with Volume Shadow Copy Services (VSS). A backup application can notify VSS when it starts a job. In response, VSS will pause disk I/O and create a “snapshot” of the system. The snapshot isolates the state of all files as they were at that moment from any changes that occur while the backup job runs. The backup signals VSS when it has finished backing up, and VSS merges the changed data into the checkpoint and restores the system to normal operation. With this technique, on-disk files are completely consistent. However, it cannot capture memory contents. If you restore that backup, it will be exactly as though the host had crashed at the time of backup. For this reason, we call this type of backup “[crash-consistent](#)”. It only partially addresses the problem of open files.



All Files Consistent; Memory Contents and Paused Transactions Lost
(Crash-Consistent)

VSS-aware applications can ensure complete consistency of the files that they control. Their authors can write a component that registers with VSS (called a “VSS Writer”). When VSS starts a snapshot operation, it will notify all registered VSS writers. In turn, they can write all pending operations to disk and prevent others from starting until the checkpoint completes. Because it has no active I/O at the time the backup is taken, the backup will capture everything about the program. We call this an “application-consistent” backup.

As you shop for backup programs, keep in mind that not everyone uses the terms “crash-consistent” and “application-consistent” in the same way. Also, Linux distributions do not have a native analog to VSS. Research the way that each candidate application deals with open files and running programs.

5.6. HYPERVISOR-AWARE BACKUP SOFTWARE

If you employ any hypervisors in your environment, you should strongly consider a backup solution that can work with them directly. You can back up client operating systems using agents installed just like physical systems if you prefer. However, hypervisor-aware backup applications can appropriately time guest backups to not overlap

and employ optimization strategies that greatly reduce time, bandwidth, and storage needs.

When it comes to your hypervisors, investigate applications with the same level of flexibility as [Altaro VM Backup](#). You can install it directly on a Hyper-V host and operate it from there, use a management console from your PC, or make use of [Altaro's Cloud Management Console](#) to manage all of your backup systems from a web browser. Such options allow you to control your backup in a way that suits you.

5.7. STANDARD PHYSICAL SYSTEMS BACKUP SOFTWARE

Few organizations have moved fully to virtualized deployments. So, you likely have physical systems to protect in addition to your virtual machines. Some vendors, such as Altaro, provide a separate solution to cover physical systems. Others use customized agents or modules within a single application. However, some companies have chosen to focus on one type of system and cannot protect the other.

5.8. SINGLE-VENDOR VS. HYBRID APPLICATION SOLUTIONS

In small environments, administrators rarely even consider using solutions that involve multiple vendors. Each separate product has its own expertise requirements and licensing costs. You cannot manage backup software from multiple vendors using a single control pane. You may not be able to find an efficient way to store backup data from different manufacturers. Using a single vendor allows you to cover the most systems with the least amount of effort.

On the other hand, organizations with more than a handful of servers almost invariably have some hybridization – in operating systems, third-party software, and hardware. Using different backup programs might not pose a major challenge in those situations. Using multiple programs allows you to find the best solution for all your problems instead of accepting one that does “enough”.

I once had a customer that was almost fully virtualized. They placed high priority on a granular backup of Microsoft Exchange with the ability to rapidly restore individual messages. Several vendors offer that level of coverage for Exchange in addition to virtual machine backup.

Unfortunately, no single software package could handle both to the customer's satisfaction.

To solve this problem, we selected one application to handle Exchange and another to cover the virtual machines. The customer achieved all of their goals and saved substantially on licensing.

5.9. PUTTING IT IN ACTION

Using the above guidance and the plan that you created in part one, you have enough information to start investigating programs that will satisfy your requirements. Let's now see how that would look like.

5.9.1. PHASE ONE: CANDIDATE SOFTWARE SELECTION

Begin by collecting a list of available software. You will need to find a way to quickly narrow down the list. To that end, you can apply some quick criteria while you search, or you can build the list first and work through it later. Maintain this list and the reasons that you decided to include or exclude a product.

" I once had a customer that was almost fully virtualized. They placed high priority on a granular backup of Microsoft Exchange with the ability to rapidly restore individual messages. Several vendors offer that level of coverage for Exchange in addition to virtual machine backup. Unfortunately, no single software package could handle both to the customer's satisfaction.

To solve this problem, we selected one application to handle Exchange and another to cover the virtual machines. The customer achieved all of their goals and saved substantially on licensing. "

Case Study



Create a table to use as a tracking system. As an example:

Product	Version	Within	Physical	Hyper-V Virtual Machines	Backup to Cloud	Active Support
Product A	7.0	Yes	Yes	No	No	No
Product B: Advanced Edition	2.1	No				
Product B: Started Edition	2.1	Yes	Yes	No	Yes	Yes
Product C	4.9	Yes	Yes	Yes	Yes	Yes

It might seem like a bit much to create this level of documentation, but it has benefits:

- **Historical purposes:** Someone might want to know why a program was tested or skipped
- **Reporting:** You may need to provide an accounting of your selection process
- **Comparisons:** Such a table forms a feature matrix

Because this activity only constitutes the first phase of selection, use criteria that you can quickly verify. To hasten the process, check for any deal-breaking problems first. You can skip any other checks for that

product. While the table above shows simple yes/no options, you can use a more nuanced grading system where it makes sense. Keep in mind that you want to shorten this list, not make a final decision.

5.9.2. PHASE TWO: IN-DEPTH SOFTWARE TESTING

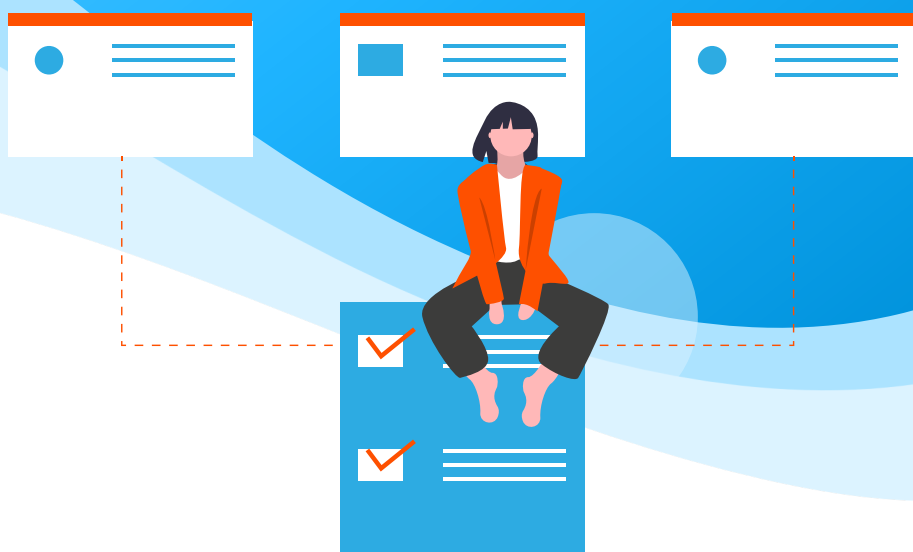
You will likely spend the most time in phase two. Phase one should have left you with a manageable list of programs to explore more completely. Now you need to spend the time to work through them to find the solution that works best for your organization. Keep in mind that you can use multiple products if that works better than a single solution.

For this phase, you will need to acquire and install software trials.

Some recommendations:

- Install trialware on templated virtual machines that you can quickly rebuild
- Use test systems that run the same programs as your production systems
- Test backing up multiple systems
- Test encryption/decryption
- Test complete and partial restores

“ DUE TO THREATS OF MALWARE AND MALICIOUS ACTORS, ALL BUSINESS CONTINUITY PLANS SHOULD INCLUDE SOME SORT OF IN-HOUSE SOLUTION THAT YOU CAN TAKE OFFLINE AND OFFSITE. ”



Extend the table that you created in phase one. If you used spreadsheet software to create it, consider creating tabs for each program that you test. You could also use a form that you build in a word processor. Make sure to thoroughly test each program. Never assume that any given program will behave like any other.

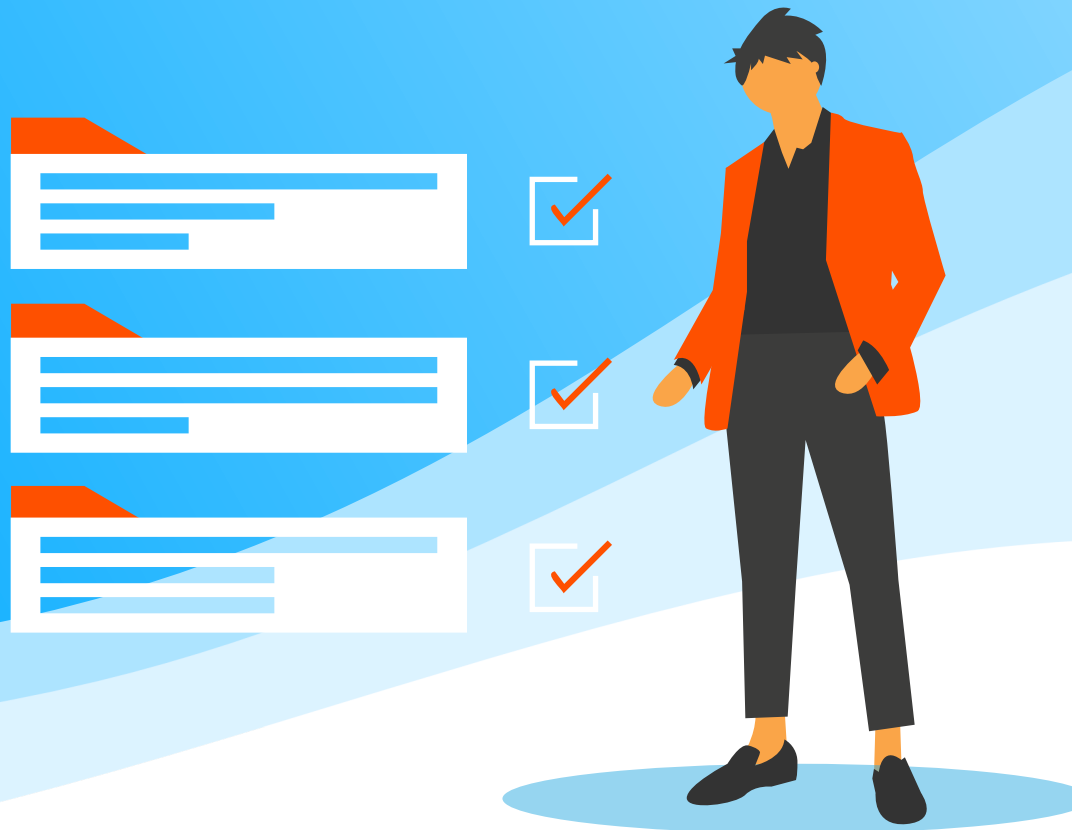
5.9.3. PHASE THREE: FINAL SELECTION

Hopefully, you will end phase two with a clear choice. Either way, you will likely need to notify the key stakeholders from phase one of your selection status. If you need additional input or executive sign-off to complete the process, work through those processes.

Unless you chose a completely cloud-based disaster recovery approach, you will still need to acquire hardware. Remember that, due to threats of malware and malicious actors, all business continuity plans should include some sort of in-house solution that you can take offline and offsite. If you have not yet made your hardware choices, then you might want to work through the next chapter before closing phase three of software selection.

CHAPTER 6:

BACKUP STORAGE TARGETS



The days of tape-only solutions have come to an end. Other media has caught up to it in cost, capacity, convenience, and reliability. You now have a variety of storage options. Backup applications that can only operate with tape have little value in modern business continuity plans.

Unless you will buy everything from a vendor or service provider that designs your solution, make certain to match your software with your hardware. Use the software's trial installation or carefully read through the manufacturer's documentation to determine which media types it works with and how it uses them.

Backup software targets the following media types:

- Magnetic tape
- Optical disc
- Direct-attached hard drives and mass media devices
- Media-agnostic network targets
- Cloud storage accounts

6.1. MAGNETIC TAPE IN BACKUP SOLUTIONS

I.T. departments have relied on tape for backup since the dawn of the concept of backup. The technology has matured surprisingly well and mostly kept up with the pace of technology. However, the physical characteristics of magnetic tape place a severe speed limit on backup and restore operations.

PROS OF MAGNETIC TAPE:

- Most backup software can target it
- Tape media has relatively low cost per gigabyte
- Reliable for long-term storage
- Lightweight media, easy to transport offsite
- Readily reusable

CONS OF MAGNETIC TAPE:

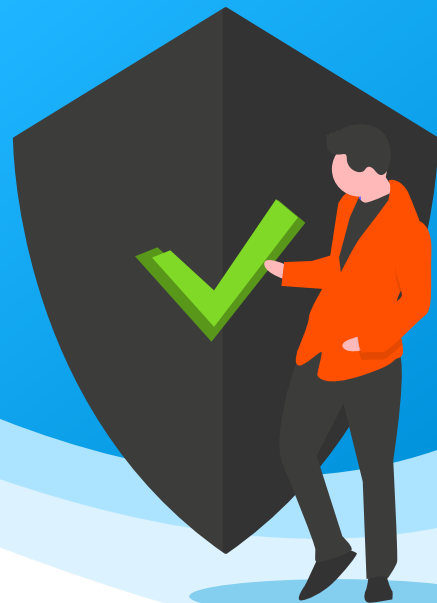
- Extremely slow
- Tape drives have a relatively high cost
- Media susceptible to magnetic fields, heat, and sunlight

“ I have seen many techniques for tape management through the years. One of the worst involved a front desk worker who diligently took the previous night’s tape offsite each night - leaving it on their car dashboard. It would bake in the sunlight for a few hours each night. So, even though the company and its staff meant well, and dutifully followed the recommendation to keep backups offsite, they wound up with warped tapes that had multiple dead spots.

At the opposite end, one customer used a padded, magnetically shielded carrying case to transport tapes to an alternative site. There, they placed the tapes into a fireproof safe in a concrete room. I was called upon once to try to restore data from a tape that was ten years old. It took almost a week to find a functioning tape drive that could accommodate it. That was the only complication.

The tape was still readable. “

Case Study



For most organizations, the slow speed of tape presents its greatest drawback. You can find backup applications that support on-demand features such as operating directly from backup media. That will not happen from a tape.

Tape has a good track record of reliability. Tapes stored on their edges in cool locations away from magnetic fields can easily survive ten years or more. Sometimes, the biggest problem with restoring from old tape is finding a suitable, functioning tape drive.

I have seen many techniques for tape management through the years. One of the worst involved a front desk worker who diligently took the previous night’s tape offsite each night – leaving it on their car dashboard. It would bake in the sunlight for a few hours each night. So, even though the company and its staff meant well, and dutifully followed the recommendation to keep backups offsite, they wound up with warped tapes that had multiple dead spots.

At the opposite end, one customer used a padded, magnetically shielded carrying case to transport tapes to an alternative site. There, they placed the tapes into a fireproof safe in a concrete room. I was called upon once to try to restore data from a tape that was ten years old. It took almost a week to find a functioning tape drive that

could accommodate it. That was the only complication. The tape was still readable.

6.2. OPTICAL MEDIA IN BACKUP SOLUTIONS

For a brief time, optical technology advances made it attractive. Optical equipment carries a low cost and interfaces well with operating systems. It even supports drag-and-drop interactivity with Windows Explorer. They were most popular in the home market. Some optical systems found their way into datacenters. However, magnetic media quickly regained the advantage as capacities outgrew optical media exponentially.

PROS OF OPTICAL MEDIA:

- Very durable media
- Shelf life of up to ten years
- Inexpensive, readily interchangeable equipment
- Drag-and-drop target in most operating systems
- Lightweight media, easy to transport offsite

CONS OF OPTICAL MEDIA:

- Very limited storage capacity
- Extremely slow
- Few enterprise backup applications will target optical drives
- Poor reusability
- Wide variance in data integrity after a few years

When recordable optical media first appeared on the markets, people found its reliability attractive. CDs and DVDs do not care about magnetic fields at all and have a higher tolerance for heat and sunlight. Also, because the media itself has no mechanism at all, they survive rough handling better than tape.

However, they have few other advantages over other media types. Even though the ability to hold 700 megabytes on a plastic disc was impressive when recordable CDs first appeared, optical media capacities did not keep pace with magnetic storage. By the time recordable DVDs showed up with nearly five gigabytes of capacity, hard drives and tapes were already moving well beyond that limit.

Furthermore, people discovered – often the hard way – that even though optical discs have hardy structural material, their data-retaining material has much shorter life. Even though a disc may look just fine, its contents may have become unreadable long ago. Recordable optical media has a wide range of data life, from a few years to several decades. Predicting media life span has proven difficult.

Because of its speed, low capacity, and need for frequent testing, you should avoid optical media in your disaster recovery solution.

6.3. DIRECT-ATTACHED STORAGE AND MASS MEDIA DEVICES IN BACKUP SOLUTIONS

You do not need to limit your backup solutions to systems that distinguish between devices and media. You can also use external hard drives and multi-bay drive chassis. Some attach temporarily, usually via USB. Others, especially the larger units, use more permanent connections such as Fiber Channel. These types of systems have become more popular as the cost of magnetic disks has declined. They have a somewhat limited scope of applications in a disaster recovery solution, but some organizations can put them to great use.

“ BECAUSE OF ITS SPEED, LOW CAPACITY, AND NEED FOR FREQUENT TESTING, YOU SHOULD AVOID OPTICAL MEDIA IN YOUR DISASTER RECOVERY SOLUTION. ”



PROS OF DIRECTLY ATTACHED EXTERNAL DEVICES:

- Fast
- Reliable for long-term storage
- Inexpensive when using mechanical drives
- Easily expandable
- High compatibility
- Use as a standard file system target

CONS OF DIRECTLY ATTACHED EXTERNAL DEVICES:

- Difficult to transport
- Additional concerns when disconnecting
- Mechanical drives have many failure points
- Expensive when using solid-state drives
- Not a valid target in every backup application

Portability represents the greatest concern when using directly attached external devices for backup. Unlike tapes and discs, the media does not simply eject once the backup concludes. With USB devices, you should notify the operating system of pending removal so that it has a chance to wrap up any writes, which could include metadata operations and automatic maintenance. Directly connected Fiber Channel devices usually do not have any sort of quick-detach mechanism. In an emergency, people should concern themselves more with evacuation than spending time going through a lengthy detach process. In normal situations, people tend to find excuses to avoid tedious processes. Expect these systems to remain stationary and onsite.

Once upon a time, such restrictions would have precluded these solutions from a proper business continuity solution. However, as you will see in upcoming sections, other advances have made them quite viable. With that said, you should not use a directly attached device alone. Any such equipment must belong to a larger solution.

You may run into some trouble using external devices with some backup applications. Fortunately, you should never run into any modern programs that absolutely cannot backup to a disk target. However, some may only allow you to use disk for short-term storage.

Others may not operate correctly with removable disks. If you purchase your devices before your software, make certain to fully test interoperability.

Even though mechanical hard drives have advanced significantly in terms of reliability, they still have a lot of moving parts. Furthermore, designers of the typical 3.5-inch drive did not build them for portability. They can travel, but not as well as tapes or discs. Even if you don't transport them, they still have more potential failure points than tapes. Do not overestimate this risk, but do not ignore it, either.

6.4. NETWORKED STORAGE IN BACKUP SOLUTIONS

Network-based solutions share several characteristics with directly attached storage. Where you find differences between the two, you also find trade-offs. You could use the same pro/con list for networked solutions as you saw above for direct-attached systems. We place the emphasis on different points, though.

In the “pros” column, networked storage gets even higher marks for its expandability. Typically, storage units built for the network sport multiple bays. They usually have more slots than common USB-based systems. You can start with a few drives and add more as needed.

Some even allow you to connect multiple chassis, physically or logically. In short, you can extend your backup storage indefinitely with such solutions.

You will pay a higher overall cost per gigabyte for network-attached storage due to the network components. The need for manufacturers to build in the necessary infrastructure to present their devices as network citizens tends to have a side effect: even more features. Almost all provide some level of security filtering. Less expensive devices, typically marketed simply as “Network-Attached Storage” (NAS), may not provide much more than that. Higher-end equipment, commonly called “Storage Area Network” (SAN), boasts many more features. You can often make them SAN storage show up in connected computers much like directly attached disks. In all, the more you pay, the more you get. Unfortunately, though, cost increases more rapidly than features.

What you gain in capacity and features, you lose in portability. Many NAS and SAN systems are rack-mounted, so you cannot transport them to a safe offsite location without significant effort. But, because these devices have a network presence, you can place them in remote locations. Using them requires some sort of site-to-site network

connection, which introduces security concerns, possible reduction in speed, and more points of failure.

Even though placing networked storage offsite involves additional risks, it also presents opportunity. Most NAS and SAN devices have replication technology. You can back up to a local device and configure it to automatically replicate to remote site(s). If your device cannot perform replication, or if you have different devices and they cannot replicate to each other, your backup software may have its own replication methods.

6.4.1. USING COMMODITY COMPUTING EQUIPMENT AS BACKUP STORAGE

Up to this point, we have talked about network-attached devices only in terms of dedicated appliances. SANs have earned a reputation for carrying price tags that exceed their feature sets. In the best case, that reduces your budget's purchasing power. More commonly, an organization cannot afford to put a SAN to its fullest potential – if they can afford one at all.

As a result, you now have choices in software-based solutions that run on standard server-class computing systems. Some backup

applications can target anything that presents a standard network file protocol, such as NFS or SMB. Software vendors and open-source developers provide applications that provide network storage features on top of general-purpose operating systems. These solutions fill the price and feature space between NAS and SAN devices. They do require more administrative effort to deploy and maintain than dedicated appliances, however.

When I built my first backup solution with the intent of targeting a dedicated appliance, I quickly learned that hardware vendors emphasize the performance features of their systems. Since I only needed large capacity, I priced a low-end rack-mount server with many drive bays filled with large SATA drives. I saved quite a bit over the appliance options.

“ When I built my first backup solution with the intent of targeting a dedicated appliance, I quickly learned that hardware vendors emphasize the performance features of their systems. Since I only needed large capacity, I priced a low-end rack-mount server with many drive bays filled with large SATA drives. I saved quite a bit over the appliance options. ”



6.5. CLOUD STORAGE IN BACKUP SOLUTIONS

Several technological advances in the past few years have made Internet-based storage viable. Most organizations now have access to reliable, high-speed Internet connections at low cost. You can leverage that to solve one of the most difficult problems in backup: keeping backup data in a location safe from local disasters. Of course, these rewards do not come without risk and expense.

PROS OF CLOUD BACKUP:

- Future-proof
- Offsite from the beginning
- Wide geographical diversity
- Highly reliable
- Effectively infinite expandability
- Access from anywhere
- Security

CONS OF CLOUD BACKUP:

- Dependencies outside your control
- Expensive to switch vendors
- Possibility of unrecoverable interruptions
- Speed

To keep their promises to customers, cloud vendors replicate their storage across geographical regions as part of the service (cheaper plans may not offer this protection). So, even though do you need to worry about failures in the chain of network connections between you and your provider and about outages within the cloud provider, you know that you will eventually regain access to your data.

That gives cloud backup an essentially unrivaled level of reliability.

The major cloud providers all go to great lengths to assure their customers of security. They boast of their compliance with accepted, standardized security practices. Each has large teams of security experts with no other role than keeping customer data safe.

That means that you do not need to concern yourself much with breaches at the cloud provider's level. However, you will need

to maintain the security of your account and access points. As with any other Internet-based resource, the provider must make your data available to you somehow. Malicious attackers might target your entryway instead of the provider itself. So, you still accept some responsibility for the safety of your cloud-based data.

When using cloud storage for backup, two things have the highest probability of causing failure. Your Internet provider presents the first. If you cannot maintain a reliable connection to your provider, then your backup operations may fail too often. Even if you have a solid connection, it might not have sufficient bandwidth to support your backup needs. For the latter problem, you can choose backup software such as [Altaro VM Backup](#) that provides compression and deduplication features specifically to reduce the network load.

Your second major concern is interim providers. While you can trust your cloud provider to exercise continuous security diligence, many third-party providers follow less stringent practices. If your backup system transmits encrypted data directly to a cloud account that you control, then you have little to worry about. Verify that your software uses encryption and keep up on updates, and you will have little to worry about beyond the walls of your institution. However, some

providers ship your data to an account under their control that they resell to customers. If they fall short on security measures, then they place your data at great risk. Vet such providers very carefully.

Note that “cost” did not appear on either the pro or con list. Cost will always be a concern, but how it compares to onsite storage will differ between organizations. Using cloud storage allows you to eliminate so-called “capital expenditures”: payments, usually substantial, made up-front for tangible goods. If you have an Internet connection, you will not need to purchase any further equipment. You also wipe out some “operational expenses”: recurring costs to maintain goods and services. You will need to pay your software licensing fees, and your cloud provider will regularly bill you for storage and possibly network usage. However, you will not need to purchase storage hardware, nor will your employees need to devote their time to maintaining it. You transfer all the hassle and expense of hardware ownership to your provider in exchange for a lower overall fee.

Although tempting, you should not transfer your entire backup load to a cloud provider. So, you should still expect some capital expense and local maintenance activities.

“ DUE TO THE RISKS AND SPEED LIMITS OF RELYING ON AN INTERNET CONNECTION, IT STILL MAKES THE MOST SENSE TO KEEP AT LEAST SOME OF YOUR SOLUTION ONSITE. ”



6.6. PUTTING IT IN ACTION

The previous section helped you to work through your software options. If you have made a final selection, then that has at least some control over your hardware purchase. If not, then you can explore your hardware options and work back to picking software. This is how it should play out.

6.6.1. STEPS TO PERFORMING HARDWARE SELECTION

Truthfully, your budget plays the largest restrictor in hardware options. So, start there. Work through the features that you want to arrive at your project scope. Your general process looks like this:

1. Determine budget
2. Establish other controlling parameters
 - c. Non-cloud replication only works effectively if you have multiple, geographically distant sites
 - d. Inter-site and cloud replication need sufficient bandwidth to carry backup data without impeding business operations
 - e. Rack space
3. Decide on preferred media type(s). The above explanations covered the pros and cons of each. Now you need to decide what matters to your organization:
 - d. Cost per terabyte
 - e. Device/media speed
 - f. Media durability
 - g. Media transportability
4. Prioritize desired features:
 - e. Deduplication
 - f. Internal redundancy (e.g. RAID)
 - g. External redundancy (e.g. hardware-based replication)
 - h. Security (e.g. hardware-based encryption, access control)

If you find that the cost of a specific hardware-based feature exceeds your budget, then your software might offer it. That can help you to achieve the coverage that you need at a palatable expense.

Once you have concluded your hardware selection, you could proceed to acquiring your software and equipment. However, it makes sense to work through the next portion on security before making any final decision. You might decide on a particular course for securing data that influences your purchase.

CHAPTER 7: SECURING AND PROTECTING BACKUP DATA



Multiple high-profile breaches have made everyone painfully aware of the need for data security. The theft of unencrypted backup tapes from a few major organizations widened the scope to include backup. Unfortunately, information technology departments have not done much to improve protection of cold data. Since attackers typically target active data and online systems, technology professionals and data security firms focus efforts there. For many years, businesses have avoided compromise of backup systems more by luck than by effort. In the age of ransomware, that luck will run out in dramatic fashion.

While we often think of “data security” in terms of warding off malicious intrusion, it has broader scope. Data damaged by accident or catastrophe is just as lost as data encrypted by ransomware. You must protect your information from all dangers, not just evil intent.

**“ WHILE WE OFTEN THINK OF ‘DATA SECURITY’
IN TERMS OF WARDING OFF MALICIOUS
INTRUSION, IT HAS BROADER SCOPE.**

**DATA DAMAGED BY ACCIDENT
OR CATASTROPHE IS JUST AS LOST
AS DATA ENCRYPTED BY RANSOMWARE.**

**YOU MUST PROTECT YOUR INFORMATION
FROM ALL DANGERS, NOT JUST EVIL INTENT. ”**



7.1. RISK ANALYSIS FOR BACKUP

The beginning of this book urged you to perform risk analysis for your production systems. If you did that, then you already have an idea of the importance of the various items that you back up. Most of that priority structure transfers directly to the backup copies. However, you should consider all your backup data as a target. Large organizations often segregate data in backups due to time or capacity constraints, but many coalesce all of it into one place. If you decide not to encrypt the backups of data that has no value to a thief, such as documents that you make available to the public for free, then an attacker may find a way to use that as a keyhole to get to your encrypted data.

As you think of risks to your backup data, remember one of the primary reasons that backup belongs to your disaster recovery solution: it can help your data to survive physical loss or damage to your live systems. Geographical dispersion provides a direct answer to those concerns. A proper protection system places significant distance between at least some of your backup data and its home site.

Different geographical locations face unique threats. Coastal facilities must suffer through hurricanes. Heavily forested areas deal with more fires. Inland plains deal with tornadoes. Dense urban areas sometimes

go through periods of destructive civil unrest or worse. Think through the ways that your business continuity system protects you from the realistic dangers that you face.

7.1.1. RANSOMWARE RISKS TO BACKUP

Ransomware creates a unique challenge. Where traditional attacks try to steal or destroy your data, ransomware wants to prevent you from accessing it. Standard disaster recovery technique easily thwarted early ransomware. Administrators would simply wipe out the live system entirely and rebuild from the latest backup.

As ransomware proved itself a uniquely lucrative vector for malicious actors, it received greater development efforts. Where the initial iterations of this type of malware would try to spread following the techniques of viruses and worms, newer programs can specifically target backup software. If uncaught, they will encrypt all data that they can reach. Such risk should influence your backup deployment.

7.2. SECURITY BY REDUNDANCY

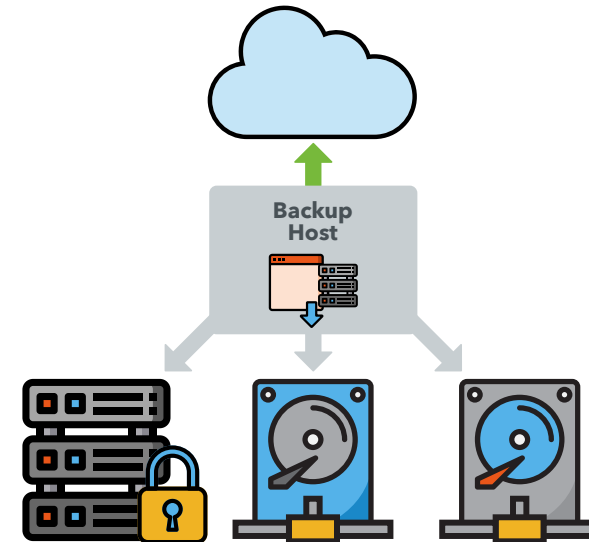
We use backup primarily because it makes a distinct copy of our live data. To solidify that protection, we need to have further redundancy within our backups. Each unique copy greatly reduces the odds of a permanent loss. In “[Exploring Disaster Recovery Technologies](#)”, we covered the multiple lines of defense against data loss. Your disaster recovery system must have its own separate tiers.

7.2.1. PROTECTING YOUR BACKUPS WITH MULTIPLE TIERS

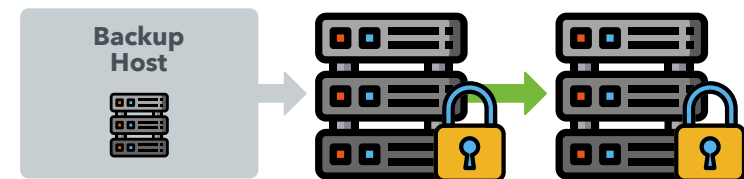
Storage cost per terabyte continually declines as technology advances. You can take advantage of that to create backups of your backups. Whereas your [rotation](#) schemes and [full backup](#) scheduling schemes will prevent corruption of deduplicated data from causing overwhelming loss, they do little to protect data that only exists on a single backup.

You have several ways to address this problem:

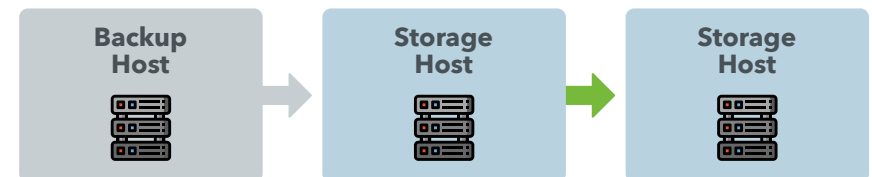
- Multiple copies in separate locations made by your backup software



- Replication of backup data using built-in NAS/SAN features



- Replication of backup data using external software



“ NO MATTER WHAT TECHNOLOGY AND TECHNIQUE YOU USE TO STORE YOUR BACKUPS, NEVER RELY ON A SINGLE COPY OF ANYTHING. ”



You can use multiple approaches as suits your needs and the technology available to you. For instance, you might have your backup software place its data on a NAS and then use a storage replication technology to copy it to another system. An older solution, called disk-to-disk-to-tape, would use backup software to keep recent data on tapes and then transfer it to disk as it aged.

Where possible, try to use the capabilities of your backup software. If someone needs to take over your deployment after your departure, you want them to leave them with the fewest complications possible. While you retain control, you do not want a convoluted system that makes your maintenance activities difficult.

7.2.2. THE ROLE OF RETENTION AND ROTATION POLICIES IN BACKUP SECURITY

In and of themselves, retention policies do not impact redundancy. However, they do set how far you can stretch your media. If you have very long retention policies, then you will require more media capacity to achieve the same frequency of full backups. No matter what technology and technique you use to store your backups, never rely on a single copy of anything. Prefer to shorten your retention policy rather than sacrifice having sufficient full backups.

To make the most use of your backup media and storage space, you will establish a rotation practice to reuse it. If you have a tape-based system, then you might opt for a scheme that reuses some tapes but keeps others for long periods of time. If you use a disk-based system, then you might rotate through removable drives or periodically exclude some backups from deduplication. Utilize these techniques in a way that balances the economics of media consumption with the value of multiple full copies.

Rotation can really shine when you leverage it as protection against malware. If malware impacts your backup solution, then it will encrypt anything that the program touches. Only your offline media will remain safe. You will need to exercise vigilance over your backup solution so that you can catch infections before one makes its way through your rotation.

7.3. USING ACCOUNT CONTROL TO PROTECT YOUR BACKUPS

Backup has a special role in your information technology environment, but it has the same foundational needs as all your other systems. So, you can apply common security practices to it.

Start by creating a unique account to run backups and lock it down.

Restrict its abilities to taking backup. If your backup application allows it, consider using different accounts in different contexts.

Exercise restraint; do not make an unmanageable mess.

Follow the same practices that you should for all vital service accounts:

- Maintain very tight control over the account – treat it like a domain administrator account.
- Place the account in an organizational unit that grants control to the fewest possible people
- Assign viable password rotation and complexity policies to the password
- Change the password immediately if anyone with access leaves the organization
- Use a properly secured password vault

7.4. ENCRYPTING YOUR BACKUP DATA

You can easily reduce the risk of your data falling into the wrong hands by employing encryption. If someone steals a tape or cracks into your cloud account, it will not gain them much if they find encrypted data for which they have no key.

All modern backup software should natively include some form of encryption. Avoid any that does not. When you try the software, ensure that you understand how it implements encryption. If you intend to rely on an application's deduplication and other storage saving features, run comparisons to determine how encryption impacts them.



While encryption does greatly strengthen the security of your backups, do not rely on it alone. If someone steals an encrypted copy of your data, then they have a copy of your data. If your attacker has the expertise, time, and willingness, they will eventually break even the best ciphers with the longest keys.

“ Years ago, almost no one encrypted their backup tapes. Some organizations would keep them in a safe deposit box at a bank or in some other kind of secured storage. Only a few took any precautions to secure them during transport or to validate their inventory. Thieves realized that with these lax controls, it was easier to steal data from tapes than to directly attack the company. After a few high-profile breaches, encryption became a much more desirable feature in backup applications. Some countries have passed laws that require certain types of institutions to encrypt sensitive data.

Fortunately, I have never been involved with any organization that had tapes stolen. However, many clients lost track of their tapes. They could usually locate their most recent backups fairly quickly, but anything older than that sometimes... disappeared. Remember that some data, such as a social security number, never truly expires. The data on your archived backup media deserves the same level of protection as the data on your active systems. ”

We expect to have many years before anyone breaks through current cryptographic schemes, but we cannot know what vulnerabilities remain hidden or how imminent technological advances will impact code breaking. Employ all available security measures.

Remember to take special care of the keys used to encrypt your backups. They represent the weakest links with this strategy. Use similar techniques to protect them to those that you implement for important account passwords.

Years ago, almost no one encrypted their backup tapes. Some organizations would keep them in a safe deposit box at a bank or in some other kind of secured storage. Only a few took any precautions to secure them during transport or to validate their inventory. Thieves realized that with these lax controls, it was easier to steal data from tapes than to directly attack the company. After a few high-profile breaches, encryption became a much more desirable feature in backup applications. Some countries have passed laws that require certain types of institutions to encrypt sensitive data.

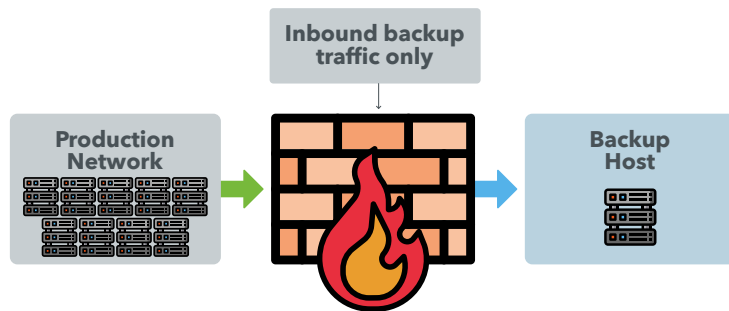
Fortunately, I have never been involved with any organization that had tapes stolen. However, many clients lost track of their tapes. They could usually locate their most recent backups fairly quickly, but anything older than that sometimes... disappeared. Remember that some data, such as a social security number, never truly expires. The data on your archived backup media deserves the same level of protection as the data on your active systems.

7.5. ISOLATING YOUR BACKUP SYSTEMS

Take steps to reduce the surface area of your backup. In some way, backup touches everything in your environment, but the reverse does not need to be true. Isolation techniques range from simple to highly complex; you will need to balance the risk of not employing a method against the effort of implementing it.

7.5.1. SHIELDING BACKUP SYSTEMS WITH FIREWALLS

Your backup application should have the ability to reach out to other systems, but almost nothing needs to reach into its system. You can put up barriers to external access easily using firewalls.



Every modern operating system includes a native firewall. Several third parties provide add-on software firewalls. Your anti-malware software might include a firewall module.

Hardware firewalls bolster software firewalls immensely. Most smaller organizations typically employ them only at the perimeter, but they can add substantial security to your internal networks as well.

Even inexpensive devices provide isolation and protection. You can also configure routers and switches with VLANs or network address translation to provide additional isolation layers.

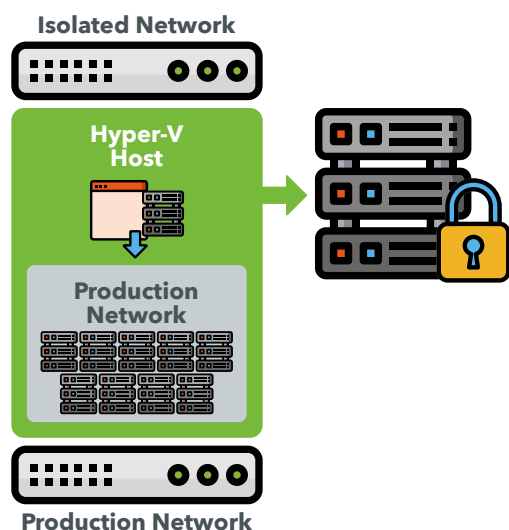
7.5.2. AIR-GAPPING FOR ISOLATION

Among all the methods of isolating, air-gapping represents the strongest. However, it also requires the most effort to implement. Before choosing this route, take the time to understand its ramifications. It should not be undertaken on a whim or without input from executive decision makers.

The simplest description of air-gapping is that there is no remote connectivity into a given system. A most extreme example is offline root certification authorities. Administrators create them, publish their public keys and revocation lists, then take them offline and disconnect them. Some even go to extra steps, such as removing their hard drives and placing them in locked storage. To access such a system, a human must perform manual steps that involve physical actions and security measures.

You cannot realistically apply such a drastic procedure to your backup system. However, it serves as a beginning. Start there and add the minimum elements to make backup operational. Backup servers need to be powered on and have some way to retrieve data from and push restores to their targets, but nothing else. To make maintenance easier, the system should have some way of sending notifications to administrators.

With all of that configured, you can function without any way to access the backup server remotely. So, you can set it up to only allow access from a physical console.



The more that you use virtualization in your environment, the easier you make it to use air-gapping. You can configure the hypervisor and backup in one network and everything else in another. If they have no overlap or interconnect, then you have created a proper air gap. You may even choose to go so far as to create an Active Directory domain just to hold these systems. That way, you can benefit from centralized management without connecting your production network to your management network.

The greatest risk with an air-gapped system is its enormous inconvenience. Preventing remote connections includes blocking valid administrative duties, too. It makes patching very difficult. It has no ability to transfer data to a remote location either, which means that you lose replication capability. You have only two choices: cope with these restrictions or do something that breaks the air-gapping. A poorly air-gapped system is more vulnerable than one that was designed from the start to participate on the network. If you cannot commit to a completely disconnected system into perpetuity, then connect your backup system and build defenses around it.

7.5.3. CARING FOR OFFLINE DATA

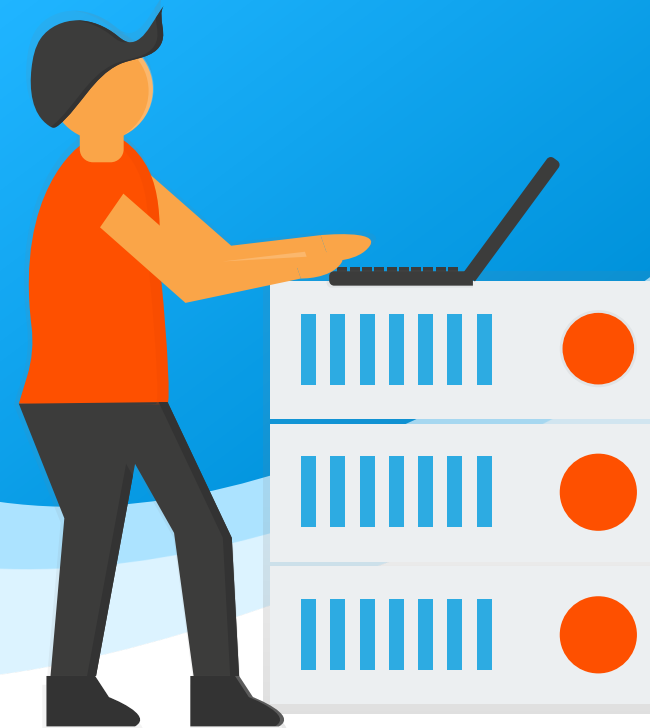
The cold data that lives on data tapes and detached hard drives often does not get the protection that it deserves. Usually, IT departments start out with a protocol to care for them, but over time, they lose diligence. We covered encryption in an earlier section, which can serve as a last-ditch safeguard. However, you must make every effort to prevent unauthorized access.

At its core, your approach involves establishing a “chain of custody” for all your backup media. If only one person has responsibility for the media, then that person must follow a defined practice for safely transporting and storing it. Treat this media as though the life of your company depends on it – because it does. Some organizations can even justify outsourcing these tasks to a security company.

Technological advances, reduced costs, and increased convenience have made fully online backup systems viable. Today, you can easily replicate backup data to geographically remote locations without an exorbitant investment. However, the ever-growing threat of ransomware means that you must periodically create offline copies. In the past, that could only mean data that was completely inaccessible by any automated means. That is still the safest option. However, you can take advantage of modern technology to create alternative approaches. You could manually upload backup data to a location that requires two-factor authentication, for instance. Whatever measures you put in place, ensure that they isolate the remote site in such a way that no compromise of your online backup system or password vault will put offline data at risk.

“ IF ONLY ONE PERSON HAS RESPONSIBILITY FOR THE MEDIA, THEN THAT PERSON MUST FOLLOW A DEFINED PRACTICE FOR SAFELY TRANSPORTING AND STORING IT.

TREAT THIS MEDIA AS THOUGH THE LIFE OF YOUR COMPANY DEPENDS ON IT — BECAUSE IT DOES. ”



7.6. PUTTING IT IN ACTION

Think of security as a continual process, not a one-time event. We will cover the hardware portions in the next chapter on deployment.

This portion will cover these security actions:

- Perform a risk analysis
- Set policies for software-level/media redundancy
- Establish backup encryption policy
- Determine practices and policies for creating and protecting offline data

You saw the concepts behind these activities throughout this chapter.

Now you must put them into practice.

7.6.1. RISK ANALYSIS ACTIVITIES

Much of risk analysis involves asking questions. You should gather input from multiple sources. Usually, one person does not have the visibility to know all likely risks. You can use formal meetings, informal queries, or any other approach that works for your organization. Categorize and list what everyone comes up with.

Share them with key stakeholders, as they might bring up other ideas.

Some starting thoughts:

- Internal vs. external categories
- Malicious vs. accidental damage
- Targeted risks (employee data, client data, account numbers, etc.)
- Equipment failures
- Weather
- Electrical outages

You must keep this list up to date with an explanation of how your solution mitigates each.

7.6.2. CREATING BACKUP REDUNDANCY POLICIES

You will need to have made your software and hardware selections before you can craft your redundancy approach. The features and media types used in your systems will have great influence on your decisions. However, your primary goal must be to create sufficient standalone copies to survive loss or damage.

To qualify as “standalone”, a backup copy must not require any other backup data to exist in order for you to restore it. Furthermore, to provide security, the copy must only exist in an offline, disconnected state.

Due to the inconvenience of offline backups, you can either build a schedule that mixes online with the occasional offline complete backup, or you can set a special schedule for offline backups.

You also need to plan for how long these offline copies will exist. The cost, longevity, and ease of storing the media tends to have the greatest control over that. If possible, simply keep them until you can no longer restore from them. Reality often dictates otherwise.

If you can schedule full backups, then you might come up with a schedule such as:

- **Monthly:** full, offline
- **Weekly:** full, online

Some modern backup software that relies heavily on deduplication technology does not allow for scheduled full backups. Instead, these depend on other policies to set the oldest possible backup and calculate deduplication from that point forward. Therefore, they consider full backups to be special, so you will need to perform them manually.

The inconvenience of such backups, especially for an already busy IT department, will likely prevent their weekly occurrence. Create a palatable policy that balances the security of multiple full copies against the ease of creating and storing them.

7.6.3. ESTABLISHING AN ENCRYPTION POLICY

You will need to build your backup encryption policy around the way that your backup hardware or software utilizes encryption.

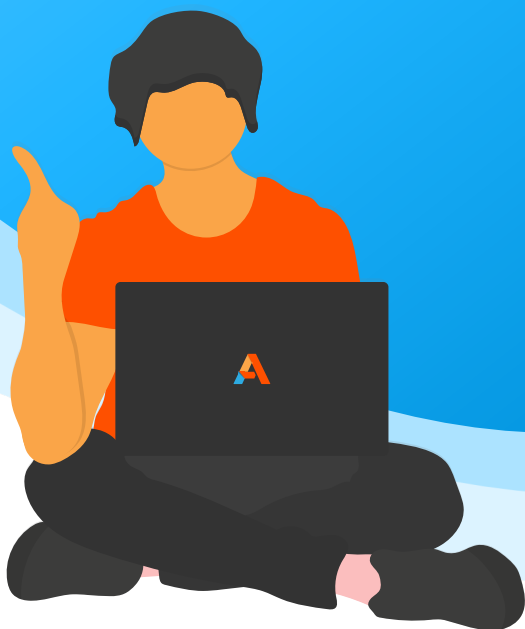
Most software requires a single secret key for encryption.

You have three major points for this type:

- Where will you store the key?
- Who will have access to the key?
- How will you ensure that the key will survive catastrophe?

It's easy to overlook the safety of access points such as keys when assessing vulnerabilities. Always include the loss of your backup encryption key in your risk analysis.

**“ ALWAYS INCLUDE THE LOSS
OF YOUR BACKUP ENCRYPTION
KEY IN YOUR RISK ANALYSIS. ”**



The location of your key directly dictates access. Since you need it to remain available in the event of a total loss, then your best option is likely a cloud-based password vault. There are multiple software companies that provide such services. Microsoft’s Azure has a “Key Vault” product and Amazon Web Services offers “AWS Secrets Manager”. Find the solution that works for your organization.

Any backup created with a particular encryption key will always need that key. So, if you change the key, you still need to keep a historical record for as long as a key has protected data.

Your hardware may offer some encryption capabilities. These features are manufacturer dependent. You will need to learn how it works before you can create a policy. If you prefer, you can simply use the software’s protection and forgo hardware-level encryption.

7.6.4. SHIELDING BACKUP WITH PHYSICAL AND NETWORK PROTECTIONS

Leverage your infrastructure and network systems to build a layer of protection around your backup systems easily and efficiently. You will need to defend at layers one, two, and three.

1. Implement layer one (physical) protections
 - a. Place backup hosts and devices in secure locations
 - b. Create a chain-of-custody process for backup media
 - c. If possible and cost effective, do not directly share switching hardware between backup systems and other systems
2. Implement layer two (Ethernet) protections
 - a. Establish a VLAN just for your backup systems, or
 - b. Use dedicated physical switches for your backup systems and connect them to the rest of your production network through a router
3. Implement layer three (TCP/IP) protections
 - a. If you isolate with a VLAN or dedicated router, create an IP subnet just for backup
 - b. Set up a firewall at the edge of the backup network that blocks all externally initiated ingress traffic

- c. Configure the software firewall on backup hosts with a similar configuration to the previous firewall

All, or most of the previous isolation techniques should fit within even modest budgets. For greater protection, you have additional options.

- Install intrusion prevention and detection solutions
- Configure network monitoring

Data moving into your backup network will fit easily recognizable patterns. With even a rudimentary monitoring system, you should have no trouble spotting suspicious traffic.

7.6.5. FULLY ISOLATING BACKUP SYSTEMS

Perform a complete risk analysis before you even consider an air-gapped approach. If you do not face significantly high exposure threats from malicious actors. Complete isolation may appear simple, but it presents substantial long-term challenges for administrators. Review the discussion above and consult with executives, key stakeholders, and others in your IT department with deployment or maintenance responsibilities.

To create an air-gapped backup system:

1. Designate an IP subnet for your air-gapped network
2. Decide on a workgroup or management domain configuration
 - a. If you will use a management domain, create and configure it before proceeding
 - b. Connect your hypervisor and backup hosts to a physical network that has no uplink
3. Ensure that the physical network links that you use for virtual machines does not provide any layer two or layer three connectivity to the hypervisor's management operating system
4. Create a policy and an accountability process for acquiring and applying software updates

The only practical risks to a properly air-gapped system are internal actors and breakout attacks against your hypervisors or container hosts. It still makes some sense to use anti-malware software as well as intrusion prevention and detection systems.

“ DUE TO THE REQUIREMENT OF FULL ISOLATION, AN AIR-GAPPED BACKUP SYSTEM ONLY WORKS FOR HYPERVISOR-BASED BACKUPS. ”

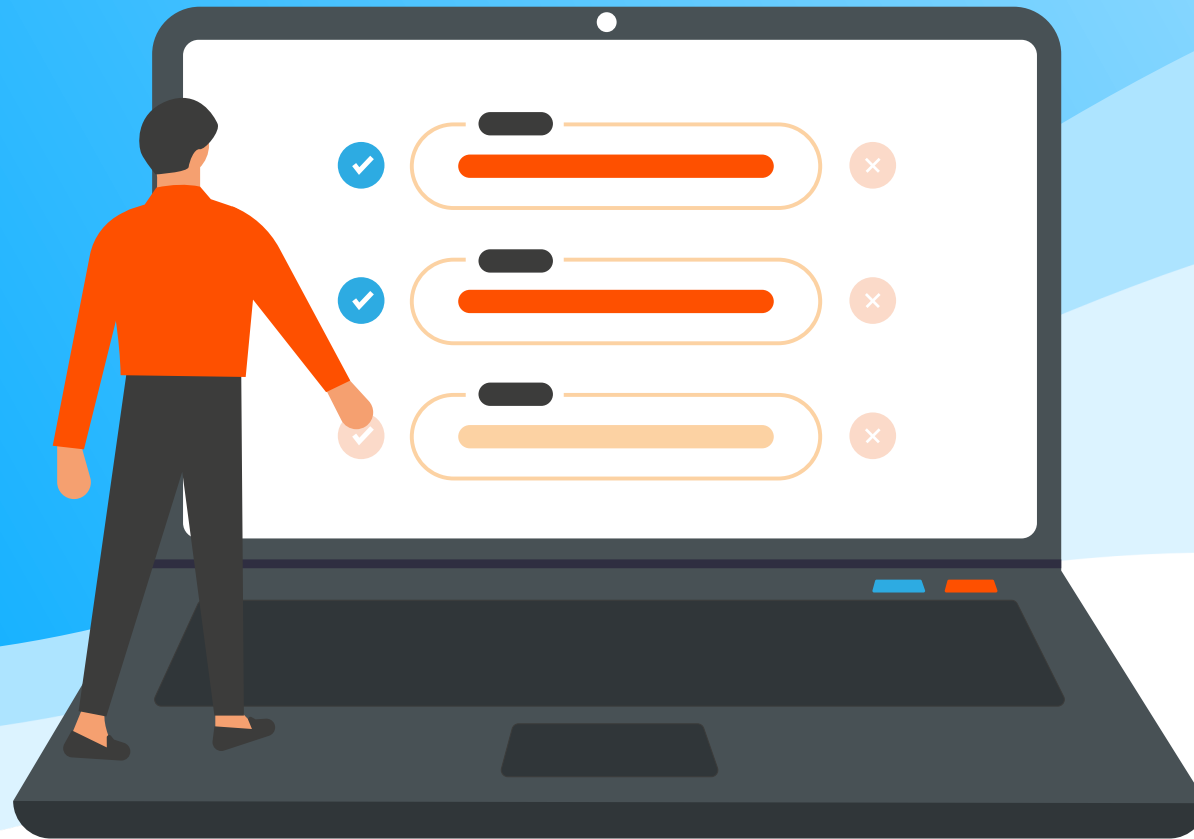


**“ REMEMBER THE ONE RULE
FOR AN AIR-GAPPED SYSTEM:
IT CANNOT PARTICIPATE
ON ANY NETWORK CONNECTION
EXCEPT THOSE DEDICATED
TO THE BACKUP SYSTEM. ”**



Remember the one rule for an air-gapped system: It cannot participate on any network connection except those dedicated to the backup system. It cannot connect to the Internet in any way. If you cannot permanently guarantee absolute isolation, then you should instead follow the steps in the previous section to allow your backups to participate on the network with adequate protections.

CHAPTER 8: DEPLOYING BACKUP



If you have completed all previous tasks, then you have enough to begin deploying your backup solution. The process is straightforward.

1. Acquire software and hardware
2. Place backup hardware
3. Install backup software into your test environment
4. If your software uses agents, push to test systems
5. Fully test your hardware and software. Verify all functionality, even the portions that you might not expect to use.
6. Document the test environment setup and installation process. Take special note of anything that did not go as planned and how you remedied the problem.
7. Install backup software into your production environment
8. If your software uses agents, push to a representative sampling of systems
9. Test expected functionality on the sample systems
10. Document the deployment into production, including fixes and workarounds

11. Continue deploying agents, if necessary, until you have covered your entire environment
12. Capture your initial full backup to store offline and transport it to the offsite location
13. Train staff on usage and have them practice
14. Document backup and restoration processes



Note there is [a print-friendly version of this checklist](#) in the Appendix.

Keep your documentation handy. As you work through the post-deployment phases, you will establish other processes and uncover other challenges. Keep careful track of any knowledge that could possibly aid yourself or others in the future. Make the documentation simple to follow, as you have no guarantee that the task of restoration will fall to someone familiar with any of your equipment or software.

First, you need to establish your recurring backup routine. You will see how to do that in an upcoming section. If you have added a replication mechanism to your backup solution, you will implement that once you have completed your initial deployment. Part three will cover replication in more detail.

CHAPTER 9: DOCUMENTING YOUR BACKUP SYSTEM



The notes that you gather during your test deployment can be informal. Use any method that will serve to prepare you for the production deployment. You need to put much more effort into the documentation that you create for your permanent installation.

9.1. DOCUMENTATION PROCEDURES AND TOOLS

If you don't have a formal documentation process for your IT activities, start one. It's certainly a best practice to document all systems, but you absolutely must take full stock of the one that forms the backbone of your disaster recovery process.

Disaster recovery planning process works differently from other systems, so its documentation has unique goals:

- Must be accessible in the absence of typical digital storage, such as an on-premises file server
- Non-technical people must be able to comprehend it
- A sufficient number of people, including personnel outside of IT, must have knowledge of the documentation and access to it

You can use any tools that you like, provided that that can produce documentation that satisfies those goals. Figure out your criteria for meeting them. Ask non-technical stakeholders for their opinions. Someday, they may need to refer to your documentation without your guidance.

If you have access to the desktop Microsoft Office version of OneNote, you will find it more than a capable tool for most needs. It shares many features with Word, so you can create headers and lists. You can paste almost any type of content, but more like PowerPoint, you can position content anywhere on the canvas. Even better, it has a simple, built-in system for categorizing and organizing information using tabs and pages. You can quickly create links to another page, section, or even notebooks. The bad news: Microsoft has ended development on the product as of OneNote 2016. They have “replaced” it with a free version. The new one adds several glitzy features, but does not allow saving files locally, will not open files created in the desktop version, and does not have a printing feature. So, for as long as you have access to the desktop version, you have one great tool.

A few other ideas:

- **Microsoft Word.** It has several features in common with OneNote. It's more difficult to organize and does not have the same free-form layouts, but almost everyone knows how to use it and you can work around its shortcomings.
- **Online Markdown sites.** You can use something like GitHub. They are automatically safe from anything that happens to your site, you can configure access control with two-factor authentication, Markdown is easy to learn, and such sites have exceptional organization and change history tracking. The downside is that non-technical, and even some technical, users do not easily understand how to use them.
- **Purpose-built applications.** Several software makers sell software designed specifically for documenting IT projects. These might also present some difficulty for non-technical users, so make sure to vet them carefully.

Just like selecting when choosing backup software, make sure your documentation software is in line with your organizational requirements.

You will also want to acquire software for taking screenshots. Building good documentation requires plentiful visual examples. Windows 10 includes its own powerful snipping tool. You can download the open-source program [Greenshot](#), which has several convenient features, such as “Capture last region”, which is extremely helpful when taking screenshots of wizard pages. You can also choose from a few high-quality paid screen capturing tools.

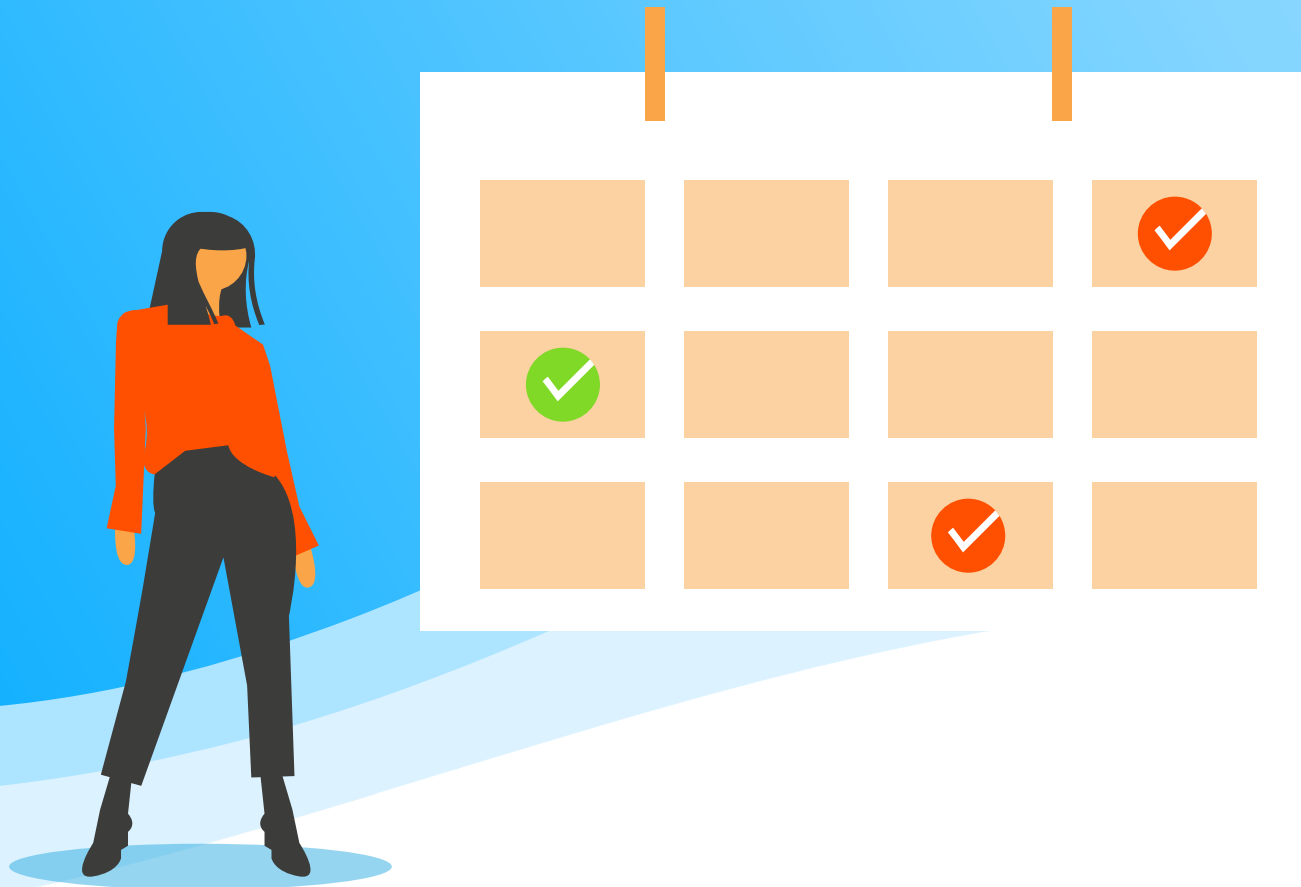


The Appendix contains two templates to help you create your own backup documentation:

[Sample Backup Documentation for a Small Organization](#) and
[Sample Backup Documentation for a Larger Organization](#).

CHAPTER 10:

DEFINING BACKUP SCHEDULES



Now you understand your organization's data protection needs and you have the means to implement. To bring it to life, you need to design the schedules for your backups. Unless you have very little data or a high budget for backup, you will use more than one schedule. You will use three metrics: value, frequency of change, and application features.

10.1. UNDERSTANDING HOW THE VALUE OF DATA AFFECTS BACKUP SCHEDULING

The frequency of your full backup schedule directly determines how many copies of data that you will have over time. The more copies you have of any given bits, the greater the odds that at least one copy will survive catastrophe. So, if you have data that you cannot lose under any circumstances, then your schedule should reflect that.

10.2. UNDERSTANDING HOW THE FREQUENCY OF CHANGE AFFECTS BACKUP SCHEDULING

Data that changes frequently may need an equally frequent backup. As you recall from part one, recovery point objectives (RPO) set the maximum amount of time between backups, which establishes the boundaries of how much recent data you can lose. You must also consider how often that data changes independently of RTO.

If you have data that does not change often, then you might consider a longer RPO. If you only modify an item every few months, then it might not make sense to back it up every week. However, that might have unintended consequences. As an example, you set a monthly-only schedule for your domain controller because you rarely have staff turnover and only replace a few computers per year. Then, you hire a new employee and supply them with a PC the day after a backup. If anything happens to Active Directory during that month, then you will lose all that new information. Your schedule needs to consider such possibilities.

10.3. UNDERSTANDING HOW BACKUP APPLICATION FEATURES AFFECT SCHEDULING

You will find that modern commercial backup applications have more in common than different. They all have some way to schedule jobs. Each one uses some way to optimize backups. The exact features in the solutions that you use will influence how you schedule.

The following list provides a starting point for you to determine how to leverage the features in your selected program:

- **Virtual machine awareness.** If your backup software understands how to back up virtual machines, then you can allow it to handle efficient ordering. If not, then you will need to schedule to back up the guest operating systems such that the jobs do not overwhelm your resources.
- **Space-saving features.** If your backup tool can preserve storage space, that has obvious benefits. Everything involves trade-offs – ensure that you know what you give up for that extra space. Some common considerations:
 - Traditional [differential](#) and [incremental backups](#) complete more quickly than the full backups they depend on. They mean nothing without their source full backup. Design your schedule to accommodate full backups as time and space allow.
 - Newer [delta](#) and deduplication techniques save even more space than differential and incremental jobs but require calculation and tracking in addition to the requisite full backups. They should not use significant CPU time, but you need to test it. Also check to see if and how your application tracks changes. Some will use space on your active disks.
 - If you have extra space in your storage media, then do not depend overly on these technologies. Create more full backups if you can.
- **Time-saving features.** Many of the features in the previous bullet point save time as well as space. As with space, do not try to save time that you do not require.
- **Replication.** Replication functions require bandwidth, which can cause severe bottlenecks when crossing Internet links. If a replication job does not complete before the next job begins, then you might end up with unusable backups.
- **Media types.** Due to the wide variance in performance of the various backup media types, the option(s) that you choose will determine how you schedule backups and what space-saving features they use. For instance, if you need to back up several terabytes to tape and a full backup requires twelve hours to perform, then you will only run a full backup when you have twelve hours available.
- **Snapshot features.** If your backup application integrates with VSS or uses some other technique to take crash-consistent or application-consistent backups, then you have greater scheduling

options. Backup uses system resources and you do not want one job to conflict with another, but snapshotting allows you to run backups while systems are in use.

You should have become well-acquainted with your backup program during the deployment phase. Take the time to fully learn how your backup program operates. Keep in mind the need for periodic full backups.

10.4. PUTTING IT IN ACTION

Remember that, if possible, you should take a complete backup of all your data at least once per day. Since that would quickly exceed any rational quantity of time and media, you must make compromises.

Guidelines for backup scheduling:

- Full backups need time and resources, even with non-interrupting snapshot technologies. Try to schedule them during low activity periods.
- Full backups do not depend on other backups. Therefore, they have greatest value after major changes. As an example, some organizations have intricate month-end procedures.

**“ IF POSSIBLE, YOU SHOULD TAKE
A COMPLETE BACKUP OF ALL YOUR
DATA AT LEAST ONCE PER DAY. ”**



Taking a backup immediately afterward could save a lot of time in the event of a restore.

- Incremental, differential, delta, and deduplicated backups require relatively little time and space compared to full backups, but they depend on other backups. Use them as fillers between full backups.
- If your backup scheme primarily uses online storage, make certain to schedule backups to offline media. If that is a manual process, implement an accountability plan.
- Just as administrators tend to perform backups at night, they also like to schedule system and software updates at night. Ensure that schedules do not collide.

10.4.1. GRANDFATHER-FATHER-SON SAMPLE PLAN

“Grandfather-father-son” (GFS) archiving is common and is available through most backup software solutions such as [Altaro VM Backup](#). It works best with rotating media such as tapes. One typical example schedule:

- **“Grandfather”**: full backup taken once monthly. Grandfather media is rotated annually (overwrite the January 2020 tape with January 2021 backup, February 2020 with February 2021 data, etc.). One “grandfather” type per year, typically the one that follows your organization’s fiscal year end, is never overwritten, following data retention policy.
- **“Father”**: full backup taken weekly. “Father” media is rotated monthly (i.e., you have a “Week 1” tape, a “Week 2” tape, etc.).
- **“Son”**: incremental or differential backups are taken daily, and their media overwritten weekly (i.e., you have a “Monday” tape, a “Tuesday” tape, etc.).

The above example is not the only type of GFS scheme. The relationship of the different rotation components is how it qualifies. You have one set of very long-term full media, one shorter-lived set of full media, and rapidly rotated media. Some implementations do not keep the annual media. Others do not rotate the monthly full, instead keeping them for the full backup retention period. Some do not rotate the daily media every week. Your organization’s needs and budget dictate your practices.

With a GFS scheme, you are never more than a few pieces of media away from a complete restore. Remember that a “differential” style backup needs the latest “son” media and the “father” immediately preceding whereas an “incremental” style backup needs the latest “father” media and all of its “sons”.

The downside of a GFS scheme is that you quickly lose the granular level of daily backups. Once you rotate the daily, then anything overwritten will, at best, survive on the most recent monthly or perhaps an annual backup. The greatest risk is to data that is created and destroyed between full backup cycles.

10.4.2. ONLINE MEDIA SAMPLE PLAN

If your backup solution uses primarily online media, then the venerated GFS approach might not work well. Most always-online systems do not have the same concept of “rotation”. Instead, they age out old data once it reaches a configured retention policy expiration period.

For these, your configuration will depend on how your backup program stores data. If it uses a deduplication scheme and only keeps a single full backup, then you have little to do except configure backup frequency and retention policy. Due to the risks posed by having only one complete

copy of your data, you must enforce periodic full backups to offline media. You should also consider some form of replication, whether to the cloud or an alternative working site.

10.4.3. CONTINUOUS BACKUP SAMPLE PLAN

Many applications have some form of “continuous” backup. They capture data in extremely small-time increments. As an example, [Altaro VM Backup](#) has a [Continuous Data Protection \(CDP\) feature](#) that allows you to set a schedule as short as five minutes.

Scheduling these types of backups involves three considerations:

- How does the backup application store the “continuous” backup data?
- How quickly does the protected data change?
- How much does the protected data change within the target time frame?

If your backup program takes full, independent copies at each interval, then you could run out of media space very quickly. If it uses a deduplication-type storage mechanism, then it should use considerably less.

Either way, your rate of data churn will determine how much space you need.

For systems with a very high rate of change, your backup system might not have sufficient time to make one backup before the next starts. That can lead to serious problems, not least of which is that it cannot provide the continuous backup that you want.

You can easily predict how some systems will behave; others need more effort. You may need to spend some time adjusting a setting, watching how it performs, and adjusting again.

10.4.4. MIXED BACKUP PLAN EXAMPLE

You do not need to come up with a one-size-fits-all schedule.

You can set different schedules. Use your RTOs, RPOs, retention policies, and capacity limits as guidance.

One possibility:

- Primary line-of-business database server: continuous, six-month retention
 - Primary file server: standard GFS with five-year retention
 - E-mail server: uses a different backup program that specializes in Exchange, daily full, hourly differential, with five-year retention
 - All: replicated to remote site every day at midnight
 - All: monthly full offline, following retention policies
- Remember to document everything!
- Domain controllers: standard GFS with one-year retention
 - Primary line-of-business application server (app only): monthly full, scheduled after operating system and software updates, with three-month retention

CHAPTER 11: MONITORING AND TESTING YOUR BACKUPS



As you might expect, setting up backup is just the beginning. You will need to keep it running into perpetuity. Similarly, you cannot simply assume that everything will work. You need to keep constant vigilance over the backup system, its media, and everything that it protects.

11.1. MONITORING YOUR BACKUP SYSTEM

Start with the easiest tools. Your backup program almost certainly has some sort of notification system. Configure it to send messages to multiple administrators. If it creates logs, use operating system or third-party monitoring software to track those as well.

Where available, prefer programs that will repeatedly send notifications until someone manually stops it or it detects problem resolution.

Set up a schedule to manually check on backup status. Partially, you want to verify that its notification system has not failed.

Mostly, you want to search through job history for things that didn't trigger the monitoring system. Check for minor warnings and correct what you can. Watch for problems that recur frequently but work after a retry. These might serve as early indications of a more serious problem.

11.2. TESTING BACKUP MEDIA AND DATA

You cannot depend on even the most careful monitoring practices to keep your backups safe. Data at rest can become corrupted.

Thieves, including insiders with malicious intent, can steal media.

You must implement and follow procedures that verify your backup data.

Keep an inventory of all media. Set a schedule to check on each piece.

When you retire media due to age or failure, destroy it. Strong magnets work for tapes and spinning drives. Alternatively, drill a hole through mechanical disks to render them unreadable. Break optical media and SSDs any way that you like.

Organizations that do not track personal or financial information may not need to keep such meticulous track of media. However, anyone with backup data must periodically check that it has not lost integrity.

The only way you can ever be certain that your data is good is to restore it. Establish a regular schedule to try restoring from older media.

If successful, make spot checks through the retrieved information to make sure that it contains what you expect.

Use this section as a basic discussion on testing best practices. We will revisit the topic of testing in a dedicated chapter toward the end of the book.

11.3. PUTTING IT IN ACTION

The activities in this section will take time to set up and perform.

Do not allow fatigue to prevent you from following these items or tempt you into putting them off.

- Configure your backup system to send alerts on failed jobs at least
- Establish an accountability for manually verifying that the backup program is functioning on a regular basis
- Configure a monitoring system to notify you if your backup software ceases running
- Establish a regular schedule and accountability system to test that you can restore data from backup. Test a representative sampling of online and offline media.

Monitoring backup, especially testing restores, is admittedly tedious work. However, it is vital. Many organizations have suffered irreparable damage because they found out too late that no one knew how to properly restore data. Too many do not realize until they've lost everything that their backup media did not successfully preserve anything. Some have had backups systems sit in a failed state for months without discovering it. A few minutes of occasional checking can prevent such catastrophes.

CHAPTER 12:

MAINTAINING YOUR SYSTEMS



The intuitive scope of a business continuity plan includes only its related software and equipment. When you consider that the primary goal of the plan is data protection, then it makes sense to think beyond backup programs and hardware. Furthermore, all the components of your backup belong to your larger technological environment, so you must maintain it accordingly.

Fortunately, you can automate common maintenance. Microsoft Windows will update itself over the Internet. The package managers on Linux distributions have the same ability. Windows also allows you to set up an update server on-premises to relay patches from Microsoft. Similarly, you can maintain internal repositories to keep your Linux systems and programs current. In addition to the convenience that such in-house systems provide, you can also leverage them as a security measure. You can automatically update systems without allowing them to connect directly to the Internet.

In addition to software, keep your hardware in good working order. Of course, you cannot simply repair modern computer boards and chips. Instead, most manufacturers will offer a replacement warranty of some kind. If you purchase fully assembled systems from a major systems vendor, such as Dell or Hewlett-Packard Enterprise, they offer warranties

that cover entire systems as a whole. They also have options for rapid delivery or in-person service by a qualified technician. If at all possible, do not allow out-of-warranty equipment to remain in service.

12.1. PUTTING IT IN ACTION

Most operating systems and software have automated or semi-automated updating procedures. Hardware typically requires manual intervention. It is on the system administrators to keep current.

- Where available, configure automated updating. Ensure that it does not coincide with backup, or that your backup system can successfully navigate operating system outages.
- Establish a pattern for checking for firmware and driver updates. These should not occur frequently, so you can schedule updates as one-off events.
- Monitor the Internet for known attacks against the systems that you own. Larger manufacturers have entries on common vulnerabilities and exposures (CVE) lists. Sometimes they maintain their own, but you can also look them up at: <https://cve.mitre.org/>. Vendors usually release fixes in standard patches,

but some will issue “hotfixes”. Those might require manual installation and other steps.

- If your hardware has a way to notify you of failure, configure it. If your monitoring system can check hardware, configure that as well. Establish a regular routine for visually verifying the health of all hardware components.

Maintenance activities consume a large portion of the typical administrator’s workload, so these procedures serve as a best practice for all systems, not just those related to backup. However, since your disaster recovery plan hinges on the health of your backup system, you cannot allow it to fall into disrepair.

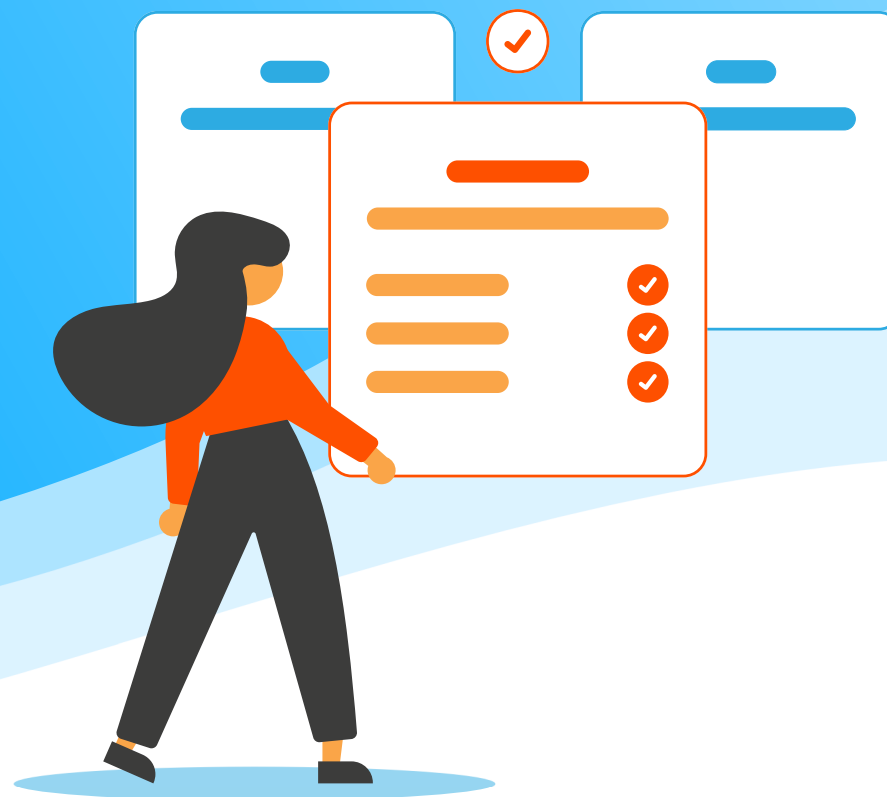
PART 3: DISASTER RECOVERY & BUSINESS CONTINUITY BLUEPRINT



CHAPTER 13:

DATA RECOVERY ACROSS

THE ORGANIZATION



When we talk about disaster recovery and business continuity planning, we devote most of our time and space to backup operations. For technologists and systems administrators, that makes sense. However, this is not the full story. For almost everyone else in a business, other things matter more. We touched some of these points briefly while discussing the planning phase. Here, we will expand on them. Because IT often drives the majority of continuity and recovery planning, it may fall to you to motivate the other business units to participate.

13.1. DISASTER RECOVERY PLANNING BEYOND THE DATACENTER

During risk analysis, you likely found many hazards that would damage much more than data and systems. Fire threatens practically everyone and everything. Floods have almost as much ubiquity; even if you build your business atop a mountain, miles away from a river, you still have plumbing. Some sites need to worry about civil unrest and terrorism. Simple mistakes can have far-reaching implications. Your disaster recovery strategy must consider all plausible (and perhaps some implausible) risks. Just as you plan to recover systems and data, you must also think of your people, buildings, equipment, and other property.

“ YOUR DISASTER RECOVERY STRATEGY MUST CONSIDER ALL PLAUSIBLE (AND PERHAPS SOME IMPLAUSIBLE) RISKS. ”



“ IT DUTIES, IT INTERSECTS EVERY SINGLE PART OF THE BUSINESS AND THEREFORE SHOULD SERVE AS THE BASIS AND HUB FOR THE OTHER DEPARTMENTS TO PREPARE THEMSELVES AND TO MAINTAIN OPERATIONAL CONTINUITY. ”



Use your findings as a basis to bring the non-technical groups into the process. You might need to convince executives in order to gain the leverage that you need. Remember that the best data protection and recovery schemes mean nothing if the organization has no plan to continue operations. A question template to capture interest: "How do we satisfy customers/meet our contractual obligations/continue making money in the event of...?"

Even though this book mainly concerns IT duties, IT intersects every single part of the business and therefore should serve as the basis and hub for the other departments to prepare themselves and to maintain operational continuity.

13.1.1. ESTABLISH PLAN SCOPE

During the initial discovery phase, you involved the leaders of other departments and teams. In addition to their technical requirements, they also have knowledge of the personnel, locations, and items needed to carry out their group's function. All of that information is vital to understanding the business' critical IT operations. Hopefully, those sorts of things were included in the early stages of planning. The sample checklists provided in the appendixes have a handful of questions related to people and things. To fully encompass your organization, your disaster recovery plan must expand further.

13.1.1.1. COVERING LOCATIONS

We will revisit the subject of business sites a few times during this discussion. Right now, we need to think in terms of which locations to include in your plan and what to record about them. It's unlikely that you would exclude any working site, but your company might own some empty or unused properties. When you update your plan in the future, you may need to remove locations that the organization no longer owns or controls.

When you place sites into the scope of your disaster recovery plan, you should centralize their role within that context. Labels like "main campus" don't mean much to someone trying to address a catastrophe. For every place listed in your plan, include:

- **Where:** Identify locations by address (e.g., 7904 West Front Road). If your organization has an informal shorthand and anyone likely to read your disaster recovery document will understand it, you can use that as well (e.g., "Pinewood B" to represent the B building on your Pinewood Road campus).
- **Normal operations purpose:** Succinctly and clearly note the typical purpose of the site (e.g., "primary accounting site"). Small companies

may not list anything other than "main operations" or "main office" or the like.

- **Disaster recovery operations purpose:** List the expected use of the site during or after a disaster if it survives. Be creative and get others involved, such as head of the team or department responsible for building maintenance. Use entries such as "overflow site for accounts receivable", "alternative shipping and receiving facility", and "tornado shelter".

This exercise will expose many things that staff should consider.

Where will employees go if you lose a site? What about customers?

Does every location have a disaster response plan? Where can we shelter people? If a loss impacts the loading dock, where else can we ship and receive freight? Look out for other vital information to include.

13.1.1.2. PROTECTING EQUIPMENT

The term "equipment" connotes different things, depending on your industry and business function. Your disaster recovery plan needs to account for all kinds. If it's tangible and isn't land, a building, or inventory that you sell as part of your normal business operations, then it might

also have a place here. The equipment inventory portion of a disaster recovery document addresses these concerns:

- What was lost or damaged in a disaster, and what is still serviceable?
- What qualifies for an insurance claim?
- What leases, loans, and rentals were impacted?
- What bulk small items would need replacements? (e.g., pens, paper stock)

Make your search broad. You need to cover vehicles, printers, desks, and anything else that your business would need in order to return to full functionality after a destructive event.

You will likely separate this portion from your major IT inventory. Personal computers, devices, and printers could logically appear in either place. You might also wish to create separate lists for different departments, sites, etc. Use any organizational method that makes sense and would have value in the hectic aftermath of a calamity.

13.1.1.3. HANDLING BUSINESS INVENTORY

If your business works with inventory flow (retail, manufacturing, etc.), then you will need a section for it, separated from the equipment entry. Due to its fluid nature, precision gains you nothing but work. Instead, document its supporting structure. Examples:

- Link to sites (warehouse, retail outlet, etc.)
- Methods to account for damaged, destroyed, or stolen inventory
- Disposal processes for damaged and destroyed inventory
- Mitigations in place (fire suppressant systems, backup generators for freezers, etc.)
- Contacts to repair, replace, and reset mitigations
- Information on insurance coverage and processes

Remember to stay on target: you're not trying to duplicate your inventory system. You only want to incorporate inventory management into your disaster recovery plan.

13.1.1.4. DELEGATING INFORMATION COLLECTION

One person will not bear the responsibility of this entire list. Depending on the size and makeup of your organization, you may need to involve people from several departments. You will have the best luck with subject-matter experts, especially if they have already done some of this work. For instance, most manufacturing companies that existed before computers have had equipment and inventory control systems in place for a long time.

Plan to encounter some resistance. Department heads might see this as an encroachment into their territory. Some people just won't want any more work to squeeze into their day. Have some points ready to show the benefits of cooperation.

- Disaster recovery planning is a unified effort for the entire company, not a way to wrest control
- In the event of a disaster, every department will be expected to take responsibility for their part. A plan will help to establish the rules in advance.
- The central plan only needs sufficient information to coordinate a recovery effort. If a department already has a detailed procedure,

then you might only need to document its input and output junction points to other departments. However, make it clear that you cannot simply use something like, "to recover the billing department, call Bill." A disaster recovery plan cannot depend on any individual.

- If all else fails, you should have gotten executive sign-off when beginning this project. Leverage that as a last resort.

Just as you have learned that business continuity planning is an ongoing process, not an event, you will need to make the leaders of the other business units aware. Set up a quarterly or biannual schedule for everyone to bring their updates and to review the document.

13.1.1.5. PREPARING AND SAFEGUARDING PERSONNEL

Human life and safety concerns make this portion of your planning stand out. For the bulk of your physical items, you will have to wait to respond until after the disaster. People need immediate attention and care. As this book primarily targets technology workers, these responsibilities may not fall to you. Either way, the organization's disaster recovery plan must include people-related preparedness and response points. Elements to consider:

- Building exit points
- Evacuation routes
- Emergency power and light sources
- First aid kits
- Fire suppression devices and systems
- Extreme weather shelter
- Staff drills
- Notification/call trees
- Check-in procedures

Some of these items might be excessive for your situation. For instance, if you have a one-room shop with two employees, then your evacuation route is "door", your call "tree" is a flat line, and drills are not the best use of your time. If something doesn't make sense to include on your plan, skip it or use some sort of placeholder to fill in as your company grows. No one should disregard this step entirely. All organizations of all sizes can find something of value here.

Do not consider this list comprehensive. Compare it to your identified risks. Gather input from others. If your business already has safety or emergency management teams, join forces. They have likely worked through all of this and you only need to find logical connection points to your data and systems recovery plans. Even if you don't have a designated team, someone might have done some of this work to comply with regulatory demands.

13.1.1.6. PRIORITIZING AND DOCUMENTING DEPARTMENTAL TIE-INS

Take special care with organizational seams. In the example that mixed IT and sales, the sales staff will depend on other teams (depending on the disaster conditions). They will be limited in their response until those other components fall into place. Without a properly predefined process, IT will likely receive a call every few minutes from a different person asking, "Are we up yet?" To prevent that, specify how response teams will utilize the notification system. Designate points of contact between the departments. When you take those spurious calls, notify the caller who in their department was selected to disseminate information. These small preparations reduce frustration and interruptions.

13.2. RESTORING NON-DATA SERVICES

Just as we need to document and prepare for disaster to strike our business units, we also need to have a path back to functionality. We've touched on this in previous points of the text where it made sense. Recovery requires a substantial amount of time and effort to perform. It should have a matching level of representation in your plan.

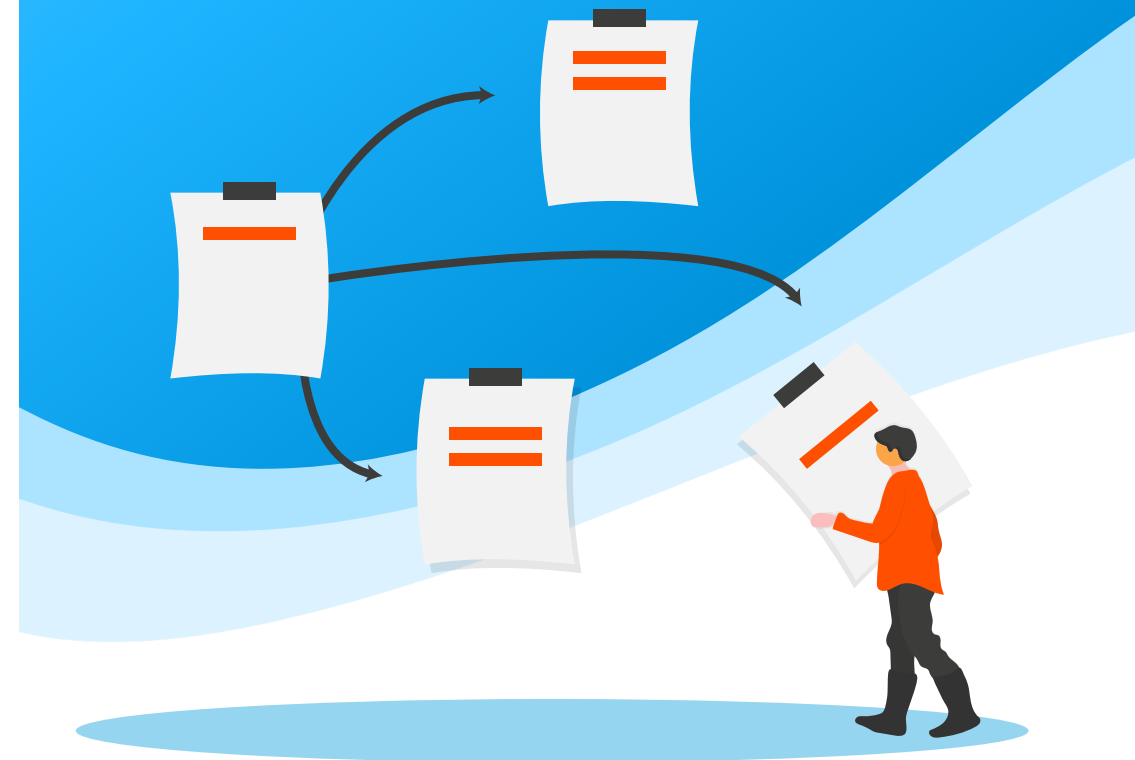
If you worked straight down through this section, then you have already begun to think about the necessary components. Buildings need to be made usable, equipment needs to return to service, inventory needs to flow. All those activities represent individual pieces. As they start up and come online, employees will move to the next tier: coordinate the pieces to enable your business' functions and services.

With modern digitized infrastructure, IT often provides a backbone for everyone else. Unfortunately, the world won't stop just because the computers don't work. You need an alternative course of action.

" WITH MODERN DIGITIZED INFRASTRUCTURE, IT OFTEN PROVIDES A BACKBONE FOR EVERYONE ELSE.

UNFORTUNATELY, THE WORLD WON'T STOP JUST BECAUSE THE COMPUTERS DON'T WORK.

YOU NEED AN ALTERNATIVE COURSE OF ACTION. "



13.2.1. DOWNTIME PROCEDURES

Each department or team needs to develop downtime procedures. Define a minimum time period that staff has no access to the system before switching procedures. If IT can bring things back in a few minutes, the switch over probably consumes too much time.

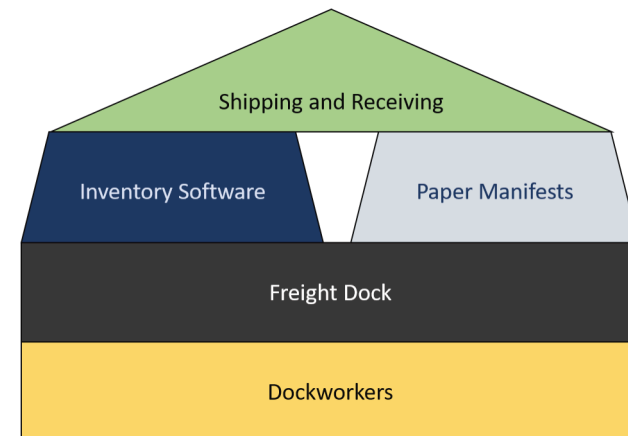
The upcoming "Business Process" section will expand these concepts. Right now, we mostly want to present its tie-ins to the non-technology portions of the business. A few ideas to get you started on developing downtime procedures:

- Prepared paper forms, receipts, and ledgers
- Traditional telephone lines as backup to VoIP
- Cell phones or handheld radios for internal communications

To reiterate, these procedures specifically apply to continuing business processes through a system failure. Do not mix them with response and recovery activities. As an example, "Periodically add gasoline to the backup generator" is fine to include. Something like, "Notify carriers of alternative pick-up and drop-off sites" should go elsewhere. Staff that does not need to devote their time to resumption efforts will follow downtime procedures.

13.2.2. DEPENDENCY HIERARCHIES

Create visual aids that illustrate the differences between necessary and desired dependencies. For simple designs, text trees might suffice. However, you already have tools that can easily produce suitable graphics. As an example, consider the following diagram of a shipping/receiving department's operation:



The structure of this graphic implies that the "Shipping and Receiving" function must have people to work the freight and a place to load and unload. The placement of the service atop two separated blocks shows that it can operate with either but must have at least one. Inventory software appears first (in a left-right culture), indicating that it has preference over the paper-based solution. If some other service relies

on the shipping and receiving function, then start another diagram for it with a single "Shipping/Receiving" block at its bottom.

An image like this works well for simple hierarchies and is easy to make. This particular graphic was built in Microsoft PowerPoint and exported as an image. PowerPoint also includes several standard flowcharting shapes for more complicated trees. If you need even more detail, you can bump up to a dedicated diagramming product such as Microsoft Visio.

13.2.3. ORGANIZING DISPARATE DOCUMENTATION

With everything in this section added in with your IT documentation, your disaster recovery plan may become unwieldy. To minimize the problem:

- Decide on a logical document organization strategy and apply it uniformly
- Use your tools' features to create a navigable table of contents
- Split the document

If possible, build a template for departments to follow. If you find that needs are too diverse for a single form, you can create multiple forms from the same source, or you can have a generic starting point with

custom content. Maintain the core principle that the people who initially create and innately understand the document may not be the same people that put the document into action. Recovery efforts might be coordinated by project managers who don't know what some of the words mean. Clarity, consistency, and uniformity matter.

Microsoft Word and most common PDF generation tools can establish intra-document links. Some can automatically create a table of contents from headers and sub-headers. Although a major disaster may make it possible that no one will have a digital copy, the pervasiveness of small devices and highly redundant cloud storage heavily reduce those odds. When connecting different parts of a document, ensure that the linkage exists as both a clickable item and with words (e.g. "voice engineering (page 14) must have restored telephone services before inbound sales can resume").

Consider splitting the document. This action breaks and sometimes removes links, so use it only as a last resort. You will need to maintain an absolute minimum of one complete digital and one complete physical document. Each revision will require you to repeat splitting operations. Even if you only revise one section, a shift in page numbering could throw off all unsynchronized copies and splits. As a mitigation,

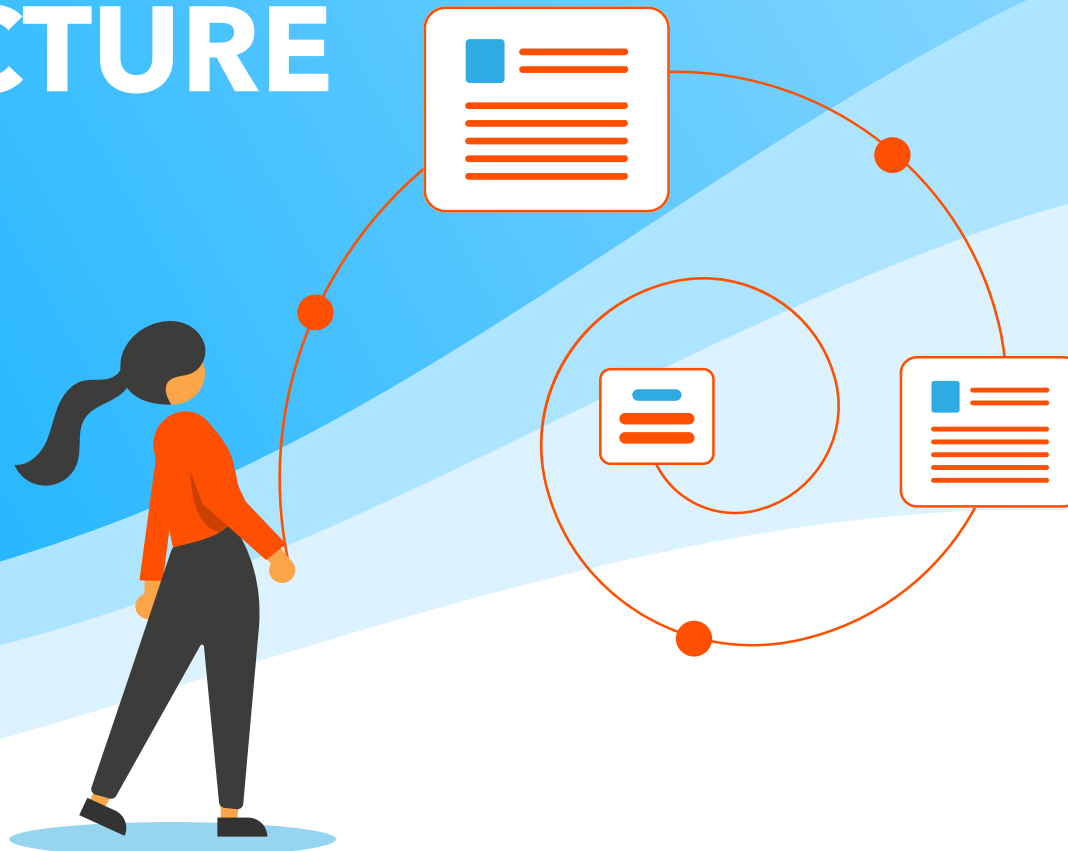
you can use only monolithic digital copies but allow people to selectively print the pages that pertain to them. You can also use a document management system or multi-document software, such as SharePoint, OneNote, or a wiki. Tools of that kind often have challenges with creating hard copies, though.

13.3. WRAPPING UP NON-TECHNICAL PLANNING

Now you can gather the other people needed for the tasks in this section. If they were not included in previous steps, make sure that they understand the goal: surviving, working through, and recovering from major and minor disasters that affect their departments, their customers, and their interactions with other parts of the business.

As they start work gathering information and preparing their documentation, you can move on to the technical details of recovery.

CHAPTER 14: BUSINESS CONTINUITY AND DISASTER RECOVERY ARCHITECTURE



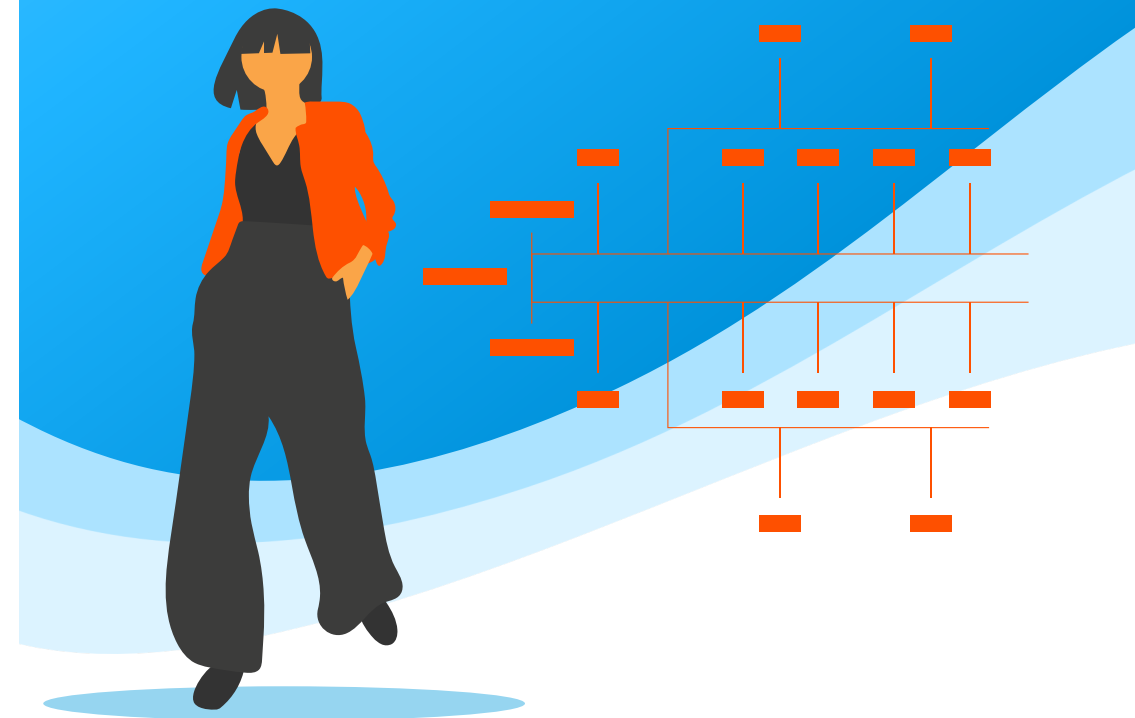
We have worked through designing and configuring a backup strategy. We have dispatched others on a quest to define their needs and roles in a disaster situation. Now we need to focus on the critical aspect for the IT admin - to architect a solution that will carry technology solutions through. The most appropriate label for this portion is "business continuity". We want to enable our systems to maintain enough functionality to support business processes through adverse situations. This expands well beyond simple backup.

14.1. USING SECONDARY SITES FOR BUSINESS CONTINUITY AND DISASTER RECOVERY

Geographically distant alternative operating sites are the most direct way to achieve business continuity in a disaster situation. Larger businesses may have additional locations that they can designate as secondary to others. In order to qualify as a secondary or alternative site, the location must have sufficient computing capability and space for dislocated personnel to stand in for the primary site. Logistical constraints may force operations to run at diminished capacity but set your goal at reasonable continuation of business functions. As we will see in the "Business Process for Disaster Recovery" section, recent technology advances have alleviated the pressure for secondary sites to match primaries.

" WE WANT TO ENABLE OUR SYSTEMS TO MAINTAIN ENOUGH FUNCTIONALITY TO SUPPORT BUSINESS PROCESSES THROUGH ADVERSE SITUATIONS.

THIS EXPANDS WELL BEYOND SIMPLE BACKUP. "



14.1.1. EVALUATING SECONDARY SITE VIABILITY

Merely owning another building does not automatically mean that you can use it as a disaster recovery location. Most importantly, the sites must have enough geographical distance between them that a single disaster won't disable both. A secondary site must also not require a prohibitive level of effort to make into a functional workspace. An empty warehouse will not replace a datacenter and call center in short order, for example.

If you don't already have business justification for a secondary site, then you might need the same employees to operate out of both locations. In that case, they must be close enough that traveling to the alternative site doesn't constitute a hardship. If the most probable types of concerns in your region are concentrated, such as building fires and tornadoes, then a few miles should suffice. If hurricanes, tsunamis, wildfires, or earthquakes threaten you, then you might face a greater challenge.

Always consider the primary business function of a site. Sadly, secondary sites simply cannot save some. If you distribute widgets from your main warehouse and a fire eliminates all the inventory, would a backup site accomplish anything meaningful? If you can file insurance claims against

the loss and redirect suppliers and carriers to another warehouse, then you can answer, "yes". If you cannot find a way for a backup site to continue the business functions of its primary, then it adds overhead without value.

Finally, verify that a site has access to all utility services. You need reliable electricity, plumbing, and Internet. Needs will vary based on the level of readiness that you expect of a site. The next sections cover the different levels for secondary sites.

14.1.2. HANDLING SPLIT RESPONSIBILITIES

With the limitless variety of configurations, one book cannot cover all possibilities. This section is written as if all sites perform all roles (operations, finance, computing, etc.). Reality ranges somewhere between that and locations dedicated to specific activities.

In your documentation, rather than pairing one site to another, you can match a function of a site to the location that can act as its secondary. For example, a site that is just a datacenter might fail over to a building that houses server hardware and sales staff.

The computing equipment at the alternative location could use the dedicated datacenter as its secondary, but the sales functions would need to target somewhere else.

14.1.3. PLANNING HOT SECONDARY SITES

A [hot site](#) can take over for a primary site very quickly. It has sufficient hardware onsite and operational and receives regular data updates directly from the primary site. Enabling such a feat requires detailed planning, high quality equipment, frequent maintenance, and constant monitoring. You will need regular, perhaps permanent, onsite staff. That staff must know how to keep the inter-site replication operating and how to fail over from the primary and back.

As you might expect, this level of functionality carries a significant cost. It works best in companies that have enough resources and volume to justify multiple locations even before considering business continuity. To operate as a hot secondary site, the location must have:

- Sufficient connectivity to the primary site during normal operations to support replication, measure speed and stability
- Server hardware powerful enough to operate in the absence of the primary site
- Physical space for personnel
- Suitable connectivity for failover conditions; think of computer and voice networking

An upcoming chapter dives deeper into replication. For now, understand that a hot site needs nearly constant data updates from its main site. That means that you need a fast and sturdy data connection between them. At the high end, you can order direct fiber runs between locations. Research the available options for point-to-point services. If you cannot find or afford such services, you will need to use general Internet connectivity instead. If possible, utilize two different providers per site. For maximum value, separate providers should use different infrastructure. It greatly reduces your redundancy if both follow the same circuits to your building or if they route through the same intermediate facilities. Rural installations have great susceptibility to outages from ditch-digging accidents. Identify such concerns and plan mitigations and workarounds.

Place a premium on data security implementation. If you can afford a point-to-point technology, then you have a lower risk profile for data interception. For the greatest protection, encrypt traffic as it traverses sites. Have devices under your control perform the encryption and decryption. Even lower-end equipment frequently supports site-to-site VPN technology. Forcing all traffic that crosses the line through an encrypted tunnel prevents the need to police all communications separately.

As a bonus, you can alleviate the CPU load on computing equipment by allowing your replication software to skip its own encryption functions.

Be mindful of the computing and data storage needs of a hot site.

It will require at least as much as the primary, and perhaps more.

It may become a "data dump" for archival purposes. As a secondary site ages without handling a catastrophe, it might find some of its resources "temporarily" repurposed. You will probably not have any real power to stop that from happening, and these "temporary" activities tend to become permanent. Make certain to maintain a minimum level of functionality and capacity at each secondary site.

Employee spaces need to be prepared to accept personnel at any time.

Prepare it like any other work site. It needs:

- Power
- Water
- Lighting
- Seating
- Desktop computing

- Voice support
- Air handling
- Physical security needs (e.g. locks, access control systems, security cameras)

You might face some struggles acquiring maintenance support to make this viable. While the data recovery portions of a plan obviously fall to IT, these types of business continuity responsibilities fall outside its purview. Your business managers will be reluctant to devote resources like this to any building that does not have a continuous personnel presence. Even if you get sign-off in the first year, that does not preclude someone from looking back in a few years and deciding that it was wasteful and that the resources should go somewhere else. In those cases, you might lose your alternative site entirely. If you are uncertain that you can maintain a hot site into perpetuity, strongly consider implementing a warm or [cold site](#) instead.

14.1.4. PLANNING WARM SECONDARY SITES

[Warm sites](#) mainly differ from hot sites in the lack of continuous data updates. We only treat that as a convention, not an unbreakable definition. In practical usage, a warm site may simply mean the closest that an organization gets to having a hot site.

A warm site has two major distinctions from a hot site:

- A warm site needs more than a few minutes' effort to resume operations from the primary.
- The inter-site network connection between a primary site and a warm site does not need to pass any special quality tests

Because a warm site does not receive continuous updates, you must have a plan in place to transfer data to the site when needed. You can achieve that by having employees transport backup tapes or drives to the site and restoring them on the hardware there. You can relay data through a cloud provider. Since your plan cannot depend on the presence of any specific individual, use the most generic descriptions and instructions possible. Anyone that the task might fall to must understand their responsibilities before needing to undertake them.

The site does need to meet all the other tests that apply to a hot site. But, if it can't function as an alternative location, then it fails the test entirely. However, you have more flexibility as the architecture and definition of a warm site include an expectation that it will take some time to spin up. In order to properly distinguish itself from a cold site, it must have adequate onsite computing abilities to resume business functions from the primary site.

14.1.5. PLANNING COLD SECONDARY SITES

Cold sites have the widest definition of the three alternative sites. Anything that could replace primary site functionality can qualify. Like warm sites, they lack an active replication scheme. They differ from warm sites in that they do not contain enough computing hardware. Such a site requires significantly less cost and effort to maintain, especially at hardware refresh intervals.

This savings comes with a risk trade-off. If you lose your primary site due to a localized building fire, then you can probably get replacement hardware quickly. If the calamity is widespread and affects a large number of businesses, you might face significant supply and delivery challenges. At the same time, if both your primary and secondary sites

“ IF ALL GOES WELL, YOU WILL NEVER NEED TO USE A SECONDARY SITE.

UNFORTUNATELY, SUCH GOOD FORTUNE CAN ALSO CAUSE A LOSS OF INTEREST AND LONG-TERM UNWILLINGNESS TO SINK FURTHER FUNDS INTO IT.

DON'T ALLOW A FALSE SENSE OF SECURITY TO OVERLOOK DILIGENT PREPARATIONS. ”



exist within the danger zone, you might work from the odds that one of the buildings remains usable. In that situation, it might make sense to only gamble with the contents of one facility.

A cold site must pass most of the non-computing tests of a hot site without the always-on restrictions. The time waiting for computers to arrive and data restoration to complete also gives you time for office furniture delivery. The power and environmental systems must function before people start, so find out if your utility companies can make that happen quickly enough that you do not need to maintain them when not in use.

Cold sites require a meaningful amount of time to begin work. They reduce your ability to continually conduct business. The next chapter will explore the technologies that can greatly mitigate these shortcomings.

14.2. ONGOING MAINTENANCE FOR SECONDARY SITES

If all goes well, you will never need to use a secondary site. Unfortunately, such good fortune can also cause a loss of interest and long-term unwillingness to sink further funds into it. Don't allow a false sense of security to overlook diligent preparations. You must include all secondary locations in the regular updates of your plan. Ask:

- Does the site still have sufficient hardware to take over for the primary?
- Do we know that the power, water, lighting, and environmental systems function?
- Do current employees know how to get to the site?
- Do current employees understand their role in transitioning to the alternative site?
- Do we have monitoring in place that guarantees the quality of replication?
- Are we replicating everything? Have we added any systems since we last answered this question?

Site maintenance goes well beyond the functions of IT. Keep the relevant departments invested. When you perform reviews of the disaster recovery plan, invite them to provide updates.



Use the [Site Description template](#) in the Appendix to help document the specifics of your primary and secondary sites to create a full understanding of your disaster recovery plan.

14.3. ANALYZING DISASTER RECOVERY HARDWARE REQUIREMENTS

In a perfect world (perfect except for disasters, anyway), you would establish secondary sites as complete mirrors of their primaries. Budgets and managerial tolerances rarely make that possible. So, you'll need to document hardware needs to enable disaster recovery. If you can afford secondary sites, then determine acceptable hardware levels. Whether you have only one site or many, you must have access to the necessary hardware to make disaster recovery possible.

14.3.1. END-USER INFRASTRUCTURE AND SYSTEMS

Some things have no room for reduction. Every knowledge worker will need a computing device. Every one of those devices will need some way to attach to the network. On the non-technical side, each person will need a chair and a work surface. The business managers responsible for related operations will need to participate in planning.

They can provide headcounts and need assessments relevant to replacement and secondary site concerns.

End-user networking will require a physical survey of any secondary sites. You can determine port counts easily, but even a cold site should not wait for cabling. You might also uncover conditions that dictate a different deployment strategy, such as per-floor local hardware instead of home runs from each endpoint.

Inter-site and Internet connectivity need planning as well. If you want the secondary to act as a hot site, you will need enough bandwidth, reliability, and security to safely transmit data from the primary. If the site has another use when not in business continuity mode, then its current Internet connection may not have sufficient bandwidth to accommodate overflow employees from the primary. Consult employees and have them think through what they need to conduct a normal days' business. Plan for printing, faxing, and other needs.

14.3.2. SERVER INFRASTRUCTURE AND SYSTEMS

For single-site recovery planning, usually you only need the specifications of your hardware. If you buy using any sort of account with a vendor,

they probably maintain a purchase history. However, they probably don't know the purpose of any of it. For best results, include a hardware catalog in your disaster recovery planning. Specify the hardware's purpose, then its specifications. For general purpose equipment that exists to extend coverage, such as end-user aggregator switches and printers, you can use locations instead.

If you will use hot or warm secondaries, they will need to have server systems onsite. Take care when configuring standby server hardware. There will be temptation to purchase lower-powered equipment than what you have in the primary site. Since you may need to run at reduced functionality, that seems logical. However, you may fail to receive sufficient funding for the secondary site when it comes time to refresh the hardware at the primary site. If that's a concern, then consider using somewhat higher-end equipment than you strictly need.

You might need to add switching, routing, firewall, and load-balancing equipment for any servers that will only operate in a failover condition. Having enough between sites to enable replication does not mean that it can also suddenly take on the load of dozens or hundreds of users performing their daily roles.

14.3.3. INTER-SITE HARDWARE

Beginning the use of multiple sites for disaster recovery requires additional equipment. You have many architectural decisions to make, and some might challenge your networking teams' current knowledge levels. Among the things to consider:

- Will you use point-to-point networking to enable replication?
- Will you maintain a constant direct Internet connection at the secondary, or will you have it physically connected but only have your provider turn it on when needed?
- If you have a constant point-to-point connection, will you also have constant Internet connections at the secondary?
Will you require the remote sites to tunnel through the primary for their Internet access and only enable direct-connect in the event of an emergency?
- Do you have the necessary hardware to perform the desired functions at each location?
- Does your staff have the networking knowledge to configure this as desired?

- Will you use temporary consultants to configure things and repair them on-demand or train your staff?

Much of your decision-making will depend on how much of a shift these secondary sites represent for your organization. If you already have multiple sites and deal with these problems today, then you likely also have the expertise on hand. If you have always used a single site with simple networking, adding even one connected site can greatly complicate everything. While the staff that you have today can certainly learn the additional functionality, you have no guarantees that they will stay. If you cannot afford to hire that level of talent into perpetuity, then consider hiring a professional networking firm to architect and maintain the inter-site links.

14.3.4. DISASTER RECOVERY HARDWARE

With all of the talk of remote sites, networking tends to dominate the discussion. Don't neglect the systems that will truly make disaster recovery possible. If you use tapes, make certain that you have access to tape drives that can read them. For tapes recorded last week, that's easy. For tapes recorded in 2002, you might have to work harder.

As things transition more to using commodity hardware and online services, this concern shrinks. Look through your backup systems

for anything that might require special handling if you lose the entire facility where the backup was taken. Make sure that you can restore its contents on alternative hardware.

14.4. MAXIMIZING DISASTER RECOVERY ARCHITECTURE

The most difficult question to answer in business continuity planning: "What are we missing?" Even comprehensive guides don't prepare you for everything. Sometimes, after going over a prepared checklist or write-up, we have a hard time thinking beyond it. For help, review your brainstorming sessions from part one. Reach out to colleagues that have a stake but have not seen what you've already come up with. Take a physical walk through your primary site and look for anything that wasn't brought up in meetings.

At no point should you claim that you have "finished" your planning. It is undoubtedly an ongoing process and should be treated as such. Always leave a few blank lines, at least metaphorically, for more information. Add disaster recovery tie-ins to any formalized processes for starting new or updating existing projects of any kind. Start up a system for employees to suggest items that didn't make the initial plans.

" AT NO POINT SHOULD YOU CLAIM THAT YOU HAVE "FINISHED" YOUR PLANNING.

IT IS UNDOUBTABLY AN ONGOING PROCESS AND SHOULD BE TREATED AS SUCH. "



CHAPTER 15: USING REPLICATION TO ENABLE BUSINESS CONTINUITY

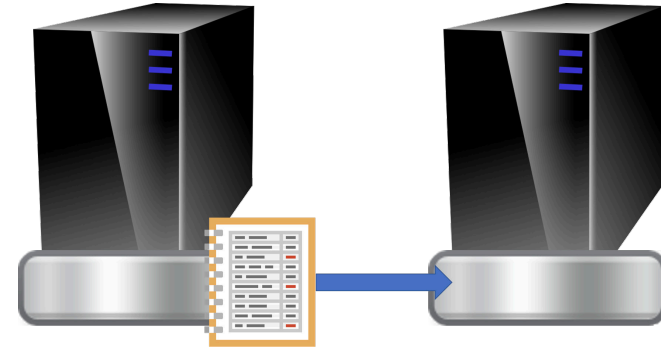


As costs for high-speed networking technology decline, we gain more ways to maintain operations through a catastrophe. Replication has changed disaster recovery more than anything else since the backup tape was first introduced. Tapes once granted us the power to conveniently move data to a safe distance from its origin. Now, we can instantly transmit changes offsite as they occur or after a short delay.

15.1. A SHORT INTRODUCTION TO REPLICATION

Replication was discussed back in part one as part of a backup strategy, but in terms of disaster recovery it requires a bit more exploration. The name says most of it: replication makes a "replica", or "copy". "Copy" invokes the idea of backup, but they have crucial differences. On the one hand, replication makes a unique, independent copy of data, just like backup. However, replicas do not have much of a historical record, nor do they have a long useful life.

Replication involves some sort of software running within the operating system or on a smart storage platform. You start by making an initial copy, called a "seed". The replication software then watches the original for changes and transmits them to another instance of the same software, which incorporates the changes into the replica.



Features of typical replication software:

- Runs continuously or on a short interval schedule
- Functions one way at a time
- May act as a component of another piece of software
- Creates a genuine duplicate of the original, not wrapped in a format proprietary to the replication engine
- Replicates without human intervention; failover to replica requires intervention

You will encounter occasional exceptions, primarily with replication systems such as Active Directory that do not treat any replica as the original. However, even in those systems, a change always occurs in one replica first, then the software transmits it to the others. Also, the product

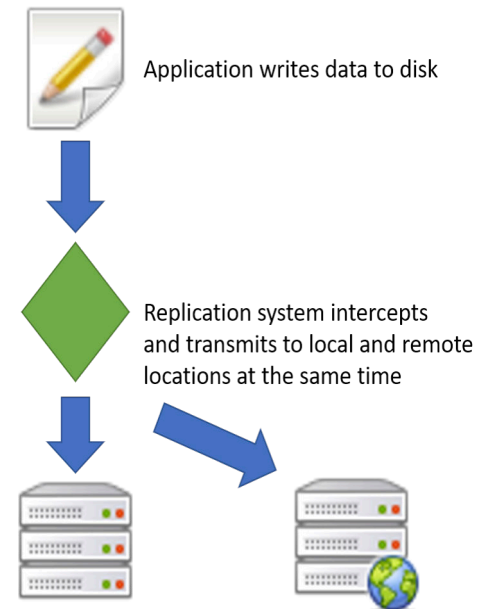
of a replica might be in a proprietary format, but typically only when the replication mechanism belongs to a larger program. As examples, some SQL server software has built-in replication mechanisms, and some backup applications include a replication component. In those cases, the format belongs to the program, not its replication engine.

[Altaro VM Backup](#), for example, supports [WAN-Optimized Replication](#) which enables users to be back up and running in minimal time should disaster strike by switching to a replicated virtual machine that is running offsite.

15.1.1. SYNCHRONOUS REPLICATION

High-end storage systems and some software offer synchronous replication. Details vary between implementations, but all end up transmitting changes from the origin to the replica in real time.

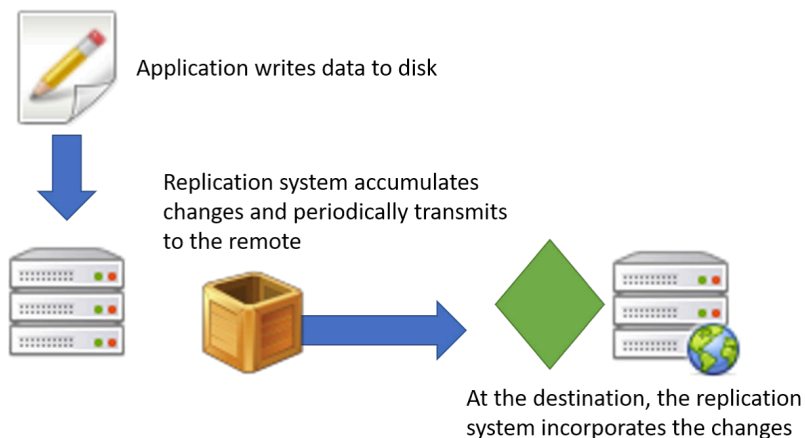
Synchronous replication processes have significant monetary, processing, and transmission costs. They allow for hot sites to pick up right from a failure point.



Some synchronous replication systems allow for geographically distributed, or "stretched" clusters. With these, you can reliably operate resources within the same cluster across distant datacenters. For clustered roles like databases, you can have almost zero-downtime failovers. For items such as virtual machines, a failed datacenter will cause its virtual machines to crash, but remote nodes with synchronous storage can bring the VMs back online immediately. Such protections allow you to architect an active/active design that keeps resources close to their users when all is well, but to continue running in an alternative location when all is not.

15.1.2. ASYNCHRONOUS REPLICATION

You will find a wider offering of asynchronous replication solutions. As the name implies, they operate with a delay. The replication mechanism accumulates changes at the origin for a period ranging from a few seconds to a few minutes. When it reaches a specified volume or time threshold, the system packages the changes and transmits them to the remote point. The receiving replication system unpacks the changes and applies them to the replica.



Asynchronous replication's primary advantage over synchronous is cost. It can transfer, test, and acknowledge paced large data chunks instead of a rapid series of small blocks, so it reduces network load. Also, because of the convenient packaging system, some replication

software will save a bit of history. In case you receive a corrupted data block, you might be able to walk back the recent changes to a good point.

Asynchronous replication can only function in active/passive mode. It does not mix with stretched clusters, although it can create a replica of a cluster at the origin site.

15.1.3. CHOOSING SYNCHRONOUS OR ASYNCHRONOUS REPLICATION

Sometimes, you will have only one viable choice with replication. Software with a high IOPS profile may not function correctly with synchronous replication. You may uncover instances in which a software-agnostic synchronous replication does not work as well as a software package's built-in asynchronous mechanism. A line-of-business vendor may prohibit supporting any installation that sits atop a synchronous replication system. In cases such as those, conditions make the decision for you.

In other cases, you have three primary factors:

- Price differences
- Recovery point objectives (RPO)
- Data value

Synchronous replication usually costs substantially more than asynchronous replication when you only compare the mechanisms. Synchronous replication also demands more from hardware, within the compute layer, the storage subsystems, and the network stack. Cost often sets the parameters before you even consider the other factors.

Recall the [discussion on RPOs](#) from the first part of this book. If a system or data set has a large RPO tolerance, then do not rush to put synchronous replication on it without some other driving force, such as stretched clusters. The shorter the RPO, the more you can justify synchronous replication. Asynchronous replication typically allows for very short delays, down to a few minutes or even a few seconds. If that satisfies your RPO, then prefer an asynchronous solution.

Even with a relatively short desired RPO, low-value data won't justify the higher cost of synchronous replica. As an example, think of a freezer

unit at a food distribution company. The historical record of its temperatures has value, especially if you have outstanding litigation over food storage. However, the temperature of the freezer in the last five minutes before the facility collapsed in an earthquake probably does not matter to anyone. So, the current information only matters operationally, so it only has value when there are current operations. Asynchronous replication can adequately protect this type of data.

Avoid mixing synchronous and asynchronous replication for the same data. It might work without error, but nothing comes without cost. Replication can place a high toll on system resources. Layering replication makes it all worse and may not have any positives.

15.2. CHOOSING REPLICATION SOLUTIONS

You will almost certainly use a mixture of replication technologies in order to achieve the best balance of support, functionality, protection levels, and resource usage. Even before looking at dedicated replication hardware or software, you have access to some replication technologies. A few that you might have right now:

- [Microsoft Active Directory](#)
- [Microsoft SQL Server](#) (the technology named “[Replication](#)” in SQL Server is typically not used for backup or business continuity)
- [Microsoft Exchange Server](#)
- [Microsoft Hyper-V Server](#)
- Backup software application, for example, [Altaro Backup](#)
- Some SAN and NAS devices
- Windows Server Datacenter Edition provides [Storage Replica](#)
- Windows Server 2019 Standard Edition has a limited implementation of Storage Replica
- [Microsoft Azure Stack HCI](#) (Hyper-Converged Infrastructure) uses multiple clustering and replication features and can create geographically-distributed clusters

Look at your major software servers and packages to see if any of them have replication capabilities.

Prefer the most specific replication technology that satisfies your requirements. Follow this decision process:

1. If you have virtual or physical machine running software that has its own replication mechanism (like Active Directory), then use the application's mechanism only.
2. If your hypervisor has a replica function and the software in the virtual machine cannot replicate itself, use the hypervisor's replication tool.
3. If the machine is physical or you can't use the hypervisor's replication (perhaps because you do not have a target system running the same hypervisor), then use operating system replication (like Storage Replica).
4. If you cannot use replication in the operating system, use NAS or SAN replication.

If, like most organizations, you have many virtual machines running a range of server applications (like AD, Exchange, SQL, etc.), then you should decide on replication separately. You will not get the best results by trying to force everything into the same solution. Some things will not work under some replication configurations that other programs can use without trouble.

The decision factors do not directly include backup replication. Most of the replication features in backup applications only make additional copies of backed-up data, not general-purpose data replicas. In that case, they only count as applications themselves (step 1) and only for the archives that they create. If your backup program has a general data replication feature, then you can prioritize it before or after step 4.

This order of preference exists for several reasons:

- If the software manufacturer went to the trouble of building a replication mechanism into their software, then it's probably the best. Many of Microsoft's technologies have been developed over decades. External replication cannot know the inner workings of these programs, so it will not work as effectively.
- Vendors will not always support their software in conjunction with certain replication technologies. For example, Microsoft does not support using Hyper-V Replica on Exchange.
- Replication functions provided by SANs, NAS devices, and hypervisors require the target to run the same or very similar system. If you decide to switch vendors, you'll have

to start replication processes created with their functions all over. If you cannot get sufficient budget to maintain the same services or equipment at all locations, you may run into some last minute or mid-stream problems.

Synchronous replication in software such as Azure Stack HCI or Storage Replica (#3) or on higher-end SAN devices (#4) might present an overriding decision point. You must remain wary of support concerns and other problems. In the absence of such barriers, synchronous replica bumps itself into the #2 preference. You will also need it anytime you intend to use fully functional stretched clusters.

While you can use synchronous replication to supplement in-application replication (#1), never replace it. As an example, a hardware-replicated domain controller will not provide the uptime and redundancy capabilities of two separate domain controllers.

**“ DO NOT REPLACE
BACKUP WITH
REPLICATION. ”**



15.3. WHERE REPLICATION FITS INTO YOUR BACKUP STRATEGY

Do not replace Backup with Replication. Backup and replication have similar features, but you cannot use them interchangeably. If you must choose between them, always choose backup. Replication exists to enable rapid failover. Replication characteristics that preclude its use as a backup tool:

- Little historical information
- Usually only one complete copy
- Limited testing ability
- No capability for regular offline copies
- Replication does not always utilize quiescing technology such as VSS

If undesirable data, such as encrypting ransomware, travels to the replica, then it will probably invalidate the entire thing. You will need to use your standard backup restoration process. Data deleted more than a short time ago will not exist anywhere in the replica's files.

You will always need the long-term and offline protection features of a true backup.

15.4. CONSIDERING REPLICATION LICENSING IMPLICATIONS

Since replication is different from backup, its use may impose some licensing considerations. Microsoft does not consider a replica virtual machine as an "offline" or "cold" copy since the replication mechanism constantly updates it and the replica is a fully functional entity distinct from the original. For that reason, hosts that maintain a replica of a Windows Server virtual machine require a separate license from the source host's license that covers the original. Above, we mentioned that Active Directory replication serves it better than other replication types, such as Hyper-V Replica. If you use Hyper-V Replica to protect a domain controller, you must still license the replica host as though the virtual machine were online. So, running one distinct domain controller in each site gives you the best replication technology and makes no difference to your licensing. Note: this rule applies to any Windows Server instance in a virtual machine, not specifically to Active Directory.

You will need to investigate the licensing rules of your software and consider them in the context of replication. This can become complicated quickly, as it can also depend on the type of replication in use (application, hypervisor, operating system, dedicated software,

or hardware) and other factors. For instance, if you add Software Assurance to a Windows Server host license, you can replicate its virtual machines to other systems without additional licensing costs. For the most comprehensive answers, work with trained licensing specialists at authorized resellers or contact software vendors directly. Altaro Backup solutions include [full 24/7 support](#) as part of its services to help users achieve the perfect configuration. Make use of services like this to ensure your licensing matches your requirements.

15.5. CONFIGURING REPLICATION

Despite the plethora of replication solutions, they share common configuration points. The exact steps will depend on your hardware or software, so we will give a generic overview of the process.

15.5.1. ESTABLISHING REPLICATION SOURCES AND TARGETS

Replication requires at least two endpoints capable of acting as replication partners. That requires a mirroring configuration of hardware and possibly software on each end. To begin, install and configure the hardware and, if you will use a software-based

mechanism, configure that as well. Necessary steps depend on your replication solution. A few examples:

- For generic data replication, configure the hardware or software as an endpoint using the system's directions.
- For Active Directory replication, install the Active Directory Domain Services role on a system in each location. Follow the necessary steps to add them to the same domain and ensure that they have IP connectivity over your inter-site link. Use Active Directory Sites and Services or PowerShell to logically separate the sites. Active Directory will automatically set up its own replication, applying special rules for traffic that crosses sites in the expectation that they have less than gigabit speeds, that the links might have high contention, and that sites may periodically lose connection.
- For Microsoft SQL replication, you first need to fully install SQL on each endpoint. You will then select one of SQL Server's many data synchronization options and configure it accordingly.
- Hyper-V Replica requires you to first configure each participating server to receive replica. For clustered Hyper-V hosts, you must

create the Hyper-V Replica Broker role and configure it instead of working on any individual node. Once you complete that step on all relevant systems, you can then configure individual virtual machines to replicate to specific target replica partners.

- If your backup application includes a replication function, follow its directions for setup and configuration.

As mentioned previously, you will likely have a mixture of replication configurations. Create a checklist of the items to cover in order to ensure that you configure them all.

16.5.2. CREATING AN INITIAL SEED

After enabling replication, the first thing that must happen is a complete build-up of the starting replica. That will probably amount to a very significant chunk of data. Perform some rough calculations based on the data size and the speed of your inter-site connection. If you discover that it might take several days to finish transmission of the beginning replicas, you can create an offline [initial seed](#). The process works like this:

1. Establish the replication partners
2. Define the data or objects that will replicate from the current site and configuration replication
3. Use the application or device's process to create an initial seed on a transportable device (such as a USB hard drive)
4. Physically transfer the seed to the target system
5. Establish the replica from the seed data
6. Start the replication process

Because operations will continue while the seed is in transit and building the replica, the replication system will need some time to catch up. You should not need to do anything else manually.

Most dedicated replication software uses the term "initial seed" or something recognizably similar. Software with built-in replication typically uses other wording. For example, follow the "[Install From Media](#)" (IFM) procedure for Active Directory.

15.6. MAINTAINING A REPLICA

Most replication technologies work unsupervised after setup.

Regardless of your confidence in the tools that you use, you should set up automated monitoring to keep an eye on them. You could also rely on some sort of daily manual verification process. However, your organization probably would not want a 24-hour or weekend period to pass without a viable synchronization. Also, the more a process tends to succeed, the less inclination tech staff will have to check on it.

Your monitoring method depends on the architecture of the replication system. Set up alerts for:

- Windows event logs
- Linux error logs
- Unexpected service halts
- Inter-site connection breaks
- Storage capacity

Some systems may have an option to send notification e-mails. While you should take advantage of those, do not rely on them. If the service fails completely, it will not send anything. It's easier to forget about a message that you never received than to ignore one that you did. Use active external monitoring.

If you don't have much experience with a particular tool, it may take some trial and error to properly balance its monitoring. As with all other things, you want it to tell you what you need to know without overwhelming you.

The technology may introduce some concepts that are new to you. For instance, most asynchronous replication systems involve some sort of "logging" and "playback" technique. For instance, [Hyper-V Replica \(HVR\)](#) builds changes at the source into a log file. At the designated time interval, it ships the log file to the target replica host. Once received, the replica host "replays" the contents of the log file into the replica. If something goes wrong with HVR, you will see symptoms in the directory that contains the replica and its log files. HVR keeps the log files for a time, but eventually cleans them up. If you're accumulating log files, that signals a problem. If you have zero log files, you will want to investigate. In the case of HVR, you should have accompanying event

**“ IT'S EASIER TO FORGET ABOUT
A MESSAGE THAT YOU NEVER
RECEIVED THAN TO IGNORE
ONE THAT YOU DID. ”**



log entries that provide detail. However, monitoring the storage location in addition to the logs gives you additional opportunities to detect a problem before it has a permanent effect. It will also set up a best practice for you in case you have a different tool that does not write to the event log or some environment problem that causes logs to roll over more quickly than you can process them.

Create a plan to accommodate resource discontinuation. Most replication systems will not automatically perform cleanup when you stop replicating. Whatever procedures you have developed for decommissioning applications and systems, append a process that describes how to stop and clean up replication. As part of discovery and initial testing, find out how to handle these situations in your replication tools.

Take time to learn if and how you can safely move replicas. If replication is a subsystem of another application, then it typically follows the application's resource moving rules. For example, you can use [Hyper-V's Storage Migration](#) to move a replica virtual machine and HVR will automatically deliver replica files to the new location. If you follow the supported steps to move the NTDS.DIT file on any domain controller, it will not break Active Directory replication. For application-agnostic replication technologies, you may have more work. Research in advance so that you never need to try to figure out how to move items under pressure.

15.7. CORRECTING PROBLEMS WITH REPLICATION

You will encounter three main problem categories with replication:

- Broken connections
- Overwhelmed destinations
- Synchronization collisions

Most mature replication technologies deal with broken connections gracefully. They wait until they can reach the destination again and pick up where they left off. Test new systems before deployment and learn how they cope with and report these events. Use this information to shape your monitoring plan and responses.

With some technologies, the replica system can fall so far behind the primary that it simply gives up and breaks out of the partnership. The exact technique to recover depends entirely on the product. Check its literature for information. Usually, the fix involves a resynchronization, which you can target for a quieter period. Discovering the root cause is equally as important as correcting the condition. If it resulted from a broken inter-site link, then you know

why and you probably have no other recourse than to fix it and move on. However, if the link stayed active, then simple corrective action may only set it up to fail again. Some ways to address repeatedly overwhelmed replication:

- Adjust the delay between transmissions. A natural instinct is to increase the delay to give the target system more time to process log files. However, it sometimes helps to reduce the interval so that the link and secondary system work with smaller files.
- Reduce the load on the inter-site link
- Increase the speed of the inter-site link
- Upgrade the target hardware

The first option is the easiest but involves potentially frustrating trial and error. The last two items will likely involve capital expenditures and contractors. To discover where to focus your efforts, set up monitoring on the resources. Learn if the target becomes overwhelmed because it doesn't receive the data in time to process it ahead of the next package, or if it receives it quickly enough but doesn't have sufficient speed to

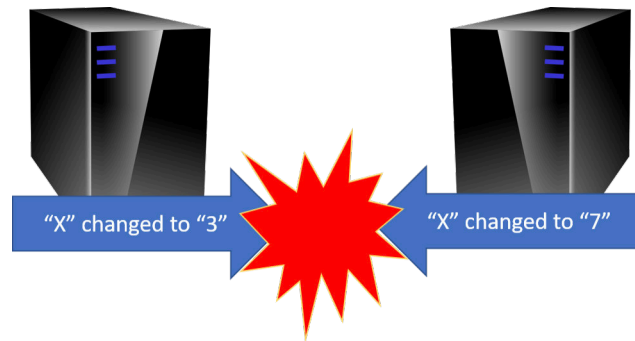
handle it before another arrives. You need to find the bottlenecks before you start trying to fix them.

A common technique for load reduction is removal of non-essential resources from the replication chain. For virtual machines, you can relocate swap data to separate virtual disks and exclude those disks from replication.

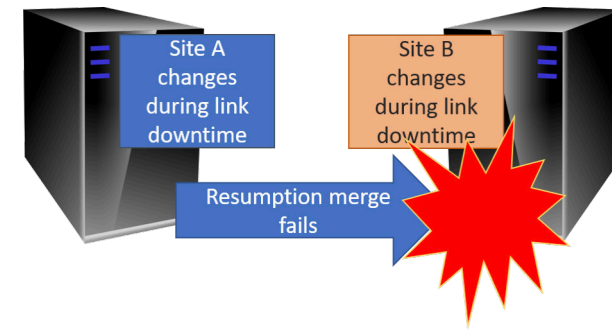
15.7.1. PREVENTING "SPLIT BRAIN" AND SYNCHRONIZATION COLLISIONS

Cluster technologies use some form of external arbiter to prevent access to the same object from multiple locations with the expectation that a completely isolated member will not come online without extraordinary steps. Controls might be complex, like Microsoft's dynamic quorum, or simple, like a lock file. In contrast, replication works with linked but unique objects. Any replication partner must have the freedom to operate on its own replica even when completely isolated. The only arbiter is human operators.

Replication functions properly when one partner processes a change to its object and transmits that change to the other partner(s). When two or more partners in replication receive changes to their local copy of the same item, you have the potential for a collision.



Active/passive replication systems have some capability to minimize these problems. Active Directory uses timestamps and other arbitration techniques to choose the one change that it will keep and records the others as historical changes. Active/passive replication typically does not have such robust protection. Consider a situation in which Site A replicates to Site B. The inter-site link drops. Site A continues operating as normal because it was the original. An operator at Site B has built a script that automatically fails to the local replica when the link drops, on the assumption that such a drop means that Site A has gone down. Unfortunately, that assumption was incorrect. The script runs, resulting in both sites actively making changes to their local replica. We call this condition "split brain". When the link is restored, Site A will try to resume synchronization. If B's replica is not in the condition that Site A expects, synchronization will fail with no automatic way to recover.



Depending on the replication technology in play, you may have a great deal of clean-up work to look forward to. Complete recovery may not be possible. In the case of Hyper-V Replica, you will need to choose one replica as the origin and resynchronize to the other as if it were new. You can copy any data that you want to save out of the replica first, then put it back into the origin. File-by-file replication systems will only have troubles with competing file changes. More complex systems with no viable repair path may suffer permanent data loss.

Even active/active mechanisms like Active Directory have some risks. It should have no problems surviving the above scenario because it was designed with those types of failures in mind. However, you can cause permanent damage to Active Directory in other ways. In the past, rolling a virtualized domain controller back to a previous state could cause irreparable damage to the directory. Research "USN rollback" for more information on that problem. For the purposes of this

discussion, understand that you can break any kind of replication technology by using it in an unsupported fashion. Most such breakdowns require restoring to an earlier backup.

A few best practices can keep you out of split-brain conditions:

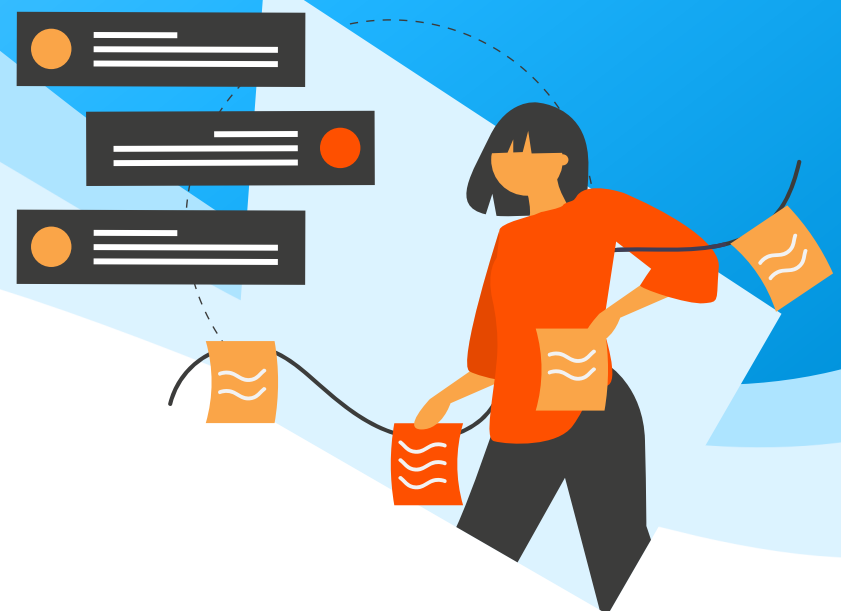
- Do not automate failover for replication systems that have no automated arbitration
- Create a defined process for initiating failover (see the upcoming section on Business Process for Disaster Recovery for more information)
- Do not mix virtual machine snapshot/checkpoint technologies with replication technologies

As a note on the last bullet point, Hyper-V incorporates its checkpoint technology to facilitate backup operations, including Hyper-V Replica. These special-purpose checkpoints pose no risk to replication. Many synchronization collisions occur because a change was made, duplicated to a replica, rolled back at the source, and then the source changed again prior to the next replication interval. The new changes appear to conflict with an earlier change, which throws the replica into an unknown state. Because Hyper-V's backup and replication checkpoint functions never revert, they do not cause collisions.

“ FUNDAMENTALLY, REPLICATION EXISTS TO ENABLE RAPID FAILOVER TO AN ALTERNATIVE SITE.

WHEN USED CORRECTLY, I.T. CAN ALLOW NEARLY UNINTERRUPTED DATA SERVICES EVEN IN A MAJOR CATASTROPHE.

WHEN USED INCORRECTLY, I.T. ADDS A LOT OF OVERHEAD AT BEST AND CAUSES A GREAT DEAL OF DAMAGE AT WORST.



15.8. LEVERAGING REPLICATION IN DISASTER RECOVERY

Fundamentally, replication exists to enable rapid failover to an alternative site. When used correctly, it can allow nearly uninterrupted data services even in a major catastrophe. When used incorrectly, it adds a lot of overhead at best and causes a great deal of damage at worst.

While replication can address the offsite requirements of backup, it does not replace any of its other components. You cannot maintain a series of offline replicas, nor will replication software have a simple way to retrieve historical data (like an e-mail or a single file).

Replication software will overwrite good data with corrupted data without hesitation and then delete its previous state.

Replica supplements backup well, but it will never replace it.

If you have sufficient funding and at least one viable alternative site, replication enhances your business continuity solution.

CHAPTER 16: THE IMPACT AND OPPORTUNITIES OF CLOUD SOLUTIONS IN DISASTER RECOVERY



Cloud technologies have fundamentally altered the way that we leverage computing resources. From software packages that require no maintenance on the part of the customer to machine learning, the cloud gives us new ways to solve challenges.

Unfortunately, it also introduces new problems, sometimes as surprises.

You might expect your cloud providers to automatically provide data protection, but they frequently offer only short-term solutions.

Check your provider's documentation carefully so that you understand what protections they do — and most importantly do not — offer.

Recently, cloud "disaster-recovery-as-a-solution" (DRaaS) offerings have begun to appear. They don't have the same shock and excitement value as other offerings, but they make up for that in convenience and perhaps cost.

16.1. UNPROTECTED CLOUD RESOURCES

It's not always obvious, but many cloud resources have no automatic protection. As an example, your Microsoft 365 e-mail has no long-term backup solution built in. You can work with third-party solutions, such as [Altaro Office 365 Backup](#), to extend the same protection that you would enjoy with on-premises Exchange e-mail, as well as SharePoint Online and OneDrive content.

Gather a full inventory of all your cloud services and determine what sort of protection schemes they employ. Extend your backup and disaster recovery plans to include them.

16.2. CLOUD BACKUP

At the simplest, you can leverage your cloud account as a backup target.

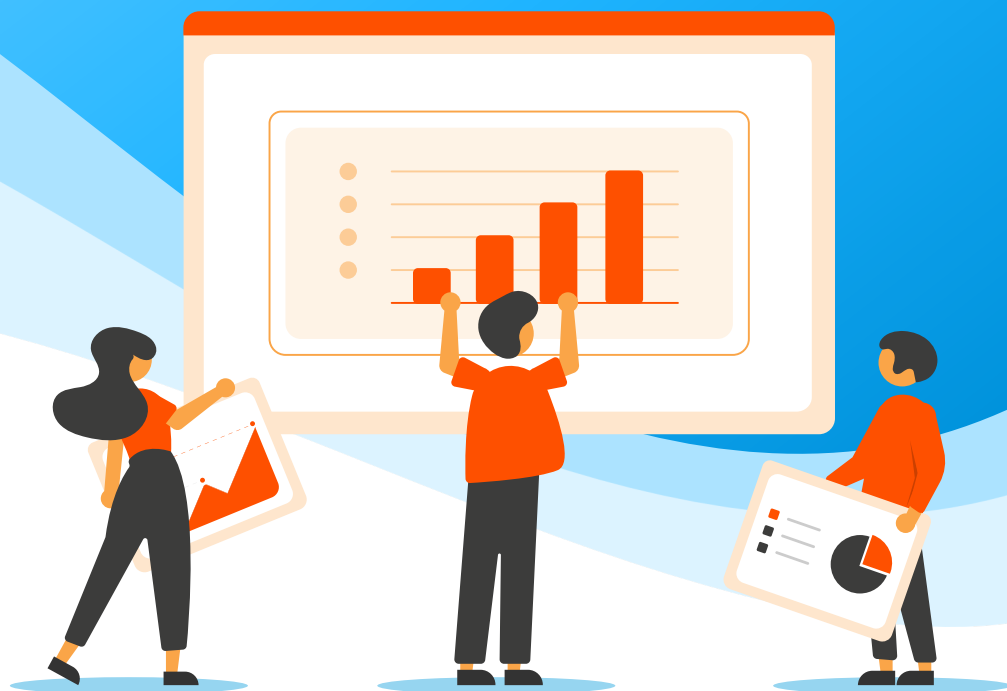
The cloud provider may offer some agent to run in your datacenter.

More probably, you use a cloud-aware backup application, such as [Altaro VM Backup](#), that automatically transfers on-premises backup data into the cloud.

Implementing cloud backup to your backup strategy gives you the protection of offsite backup storage without the need to have your own secondary site. Your staff doesn't need to travel anywhere to drop off or retrieve data. However, it's not flawless and you still need to observe safe rotational practices. If your backup software can reach a particular location without human intervention, then so can any ransomware that hijacks the backup system. Cloud backup provides the best protection in conjunction with a standard offline rotation scheme.

" GATHER A FULL INVENTORY OF ALL YOUR CLOUD SERVICES AND DETERMINE WHAT SORT OF PROTECTION SCHEMES THEY EMPLOY.

EXTEND YOUR BACKUP AND DISASTER RECOVERY PLANS TO INCLUDE THEM.



For higher-tech solutions, you can look into establishing separate storage accounts and designing operations that limit access outside of expected backup windows. For instance, you might disable your backup system's access to a dedicated "Monday" storage account all the other days of the week. If you needed to restore from the "Monday" data, you could manually override the access restriction. Such activities introduce risk and fragility into a critical function, however, and cannot guarantee to keep you safe from a ransomware assault. For the best results, use cloud as a convenience target with standard offline storage as a failsafe.

16.3. ADVANCED DISASTER RECOVERY USING CLOUD SERVICES

Using the cloud storage accounts described above, you can perform a restore operation by connecting your backup application to the account and running the same restore process that you would use with on-premises backup data. Other than data connectivity, this requires no special steps on your part.

However, cloud provides alternative recovery options. With Azure Nested Virtualization, you could use your backup program to restore Hyper-V virtual machines directly into Azure. You would then need to do some work to enable connectivity, but your employees could use Azure as their datacenter.

If you don't have the technology to directly restore a virtual machine right into Azure but you like the idea of using Azure as an alternative datacenter, you have options. You could temporarily restore resources to a lower-powered system onsite, then transfer them into Azure's SaaS and PaaS offerings. That works well for things like SQL Server and file serving. It does take planning and time to execute but saves greatly over building and operating your own secondary datacenter.

Some SaaS providers may offer enhanced services at a premium. They typically market such solutions as "DRaaS" (disaster recovery as a service) or "BCaaS" (business continuity at a service). Features and pricing vary widely and change frequently, so you will need to research them and perform comparisons before deciding on anything. Such services will inevitably cost more than doing it yourself but should not cost as much as running an alternative site.

Keep in mind that "DRaaS" solutions frequently act more like replication technology than proper backup. For instance, Azure Site Recovery (ASR) allows you to mirror physical and virtual systems into Azure, where it may keep them or relay them to another of your facilities. However, it has no historical capabilities. Therefore, ASR can serve as a replication solution, but not as backup.

" FOR NOW, TREAT CLOUD-BASED BACKUP AND BUSINESS CONTINUITY SOLUTIONS AS A CONVENIENCE.

CONTINUE CREATING LOCAL COPIES THAT YOU CAN TAKE COMPLETELY OFFLINE AND OFFSITE.



16.4. CONSIDERING CLOUD-BASED SOLUTIONS

Someday, we might move to pure cloud-based options for backup and business continuity. We still have several problems to address before that can happen. For one, we do not have any foolproof ways to maintain an always-available data storage location in the cloud that can also provide adequate protection against ransomware. Cloud backup and disaster recovery also presents the same challenge as every other cloud service: what do we do when we cannot access our cloud provider(s)? While it may seem archaic to keep your tapes and detachable hard drives around, only fully offline solutions provide any real guarantees.

For now, treat cloud-based backup and business continuity solutions as a convenience. Continue creating local copies that you can take completely offline and offsite.

CHAPTER 17: BUSINESS PROCESS FOR DISASTER RECOVERY



We've talked about cataloging personnel and items, configuring systems to protect against data loss, and setting up sites to accommodate failed over data and dislocated employees. Now, we need to establish the processes that people will follow during and after a disaster.

We covered the topic of downtime procedures earlier. We mainly intend those for times when the system is offline but recoverable. Your disaster recovery business process must accommodate failure of greater magnitude. You can use the downtime procedures that you developed as a starting point and as an idea generator, though.

17.1. EXECUTIVE DECLARATION

Enacting downtime and disaster recovery processes has associated costs. Personnel cease carrying out their normal functions and shift into their alternative emergency roles. Switching from a primary data system to a replica has time and risk implications mentioned in the relevant section. Equipment and inventory inspection and recovery efforts will accrue liabilities and debt, as will calling in contractors for any tasks. To keep things in order, categorize three levels of event response:

1. Define activities that occur during and after a crisis. These should expect minimal or no supervisory guidance. This level includes items such as moving everyone to safety, notifying authorities, and beginning low-impact downtime procedures.
2. Create a "downtime" operational level. Because switching to and from downtime operations incurs time and risk, require that it can only happen when indicated by staff with a particular level of authority. This would not include any low-impact activities that you included in the first level.
3. Specify a "response and recovery" operational level. This involves accounting for all personnel, relocating and failing over to alternative sites, and implementing equipment and data recovery processes.

The indicated names are arbitrary; use anything that makes sense. The important part is defining responses in advance during a calm so that staff have fewer problems to solve during an emergency. Having predefined levels also helps to reduce improper reactions, such as bringing a replica online while the primary site still functions.

17.2. PREPARING AND PLANNING FOR IMPACTED PERSONNEL

Your business continuity plan must cover the human aspect.

It needs to provide actions and guidance, both for the people that enact the plan and for the people impacted by whatever condition caused the plan to go into effect.

17.2.1. PREDICTING USER IMPACT

Disaster recovery plans tend to have a high degree of sterilization and focus on the business, assets, and data. While all of that probably requires the most quantity, none of it is as important as the people. Employee safety needs to top all priority lists. The plan will include a great deal of content on what processes to follow. That will help to keep staff focused, but at all phases, everyone involved in planning needs to remember that crisis conditions look nothing like a typical day at work.

You can make some predictions on the sorts of disasters that our business would be most likely to face, but that has limited value. Most people do not know how they'll react to a catastrophe until they face one. There is no such thing as a "normal" response. Some will focus and work well under pressure; others will not.

**“ YOUR BUSINESS CONTINUITY PLAN
MUST COVER THE HUMAN ASPECT. ”**



People will be scared, in shock, injured, or have any of a number of other adverse responses. Afterward, the effects can linger. The death or serious injury of a coworker can traumatize others. While you have no way to know exactly what will happen, you can plan with the expectation that anyone who needs to put it into action will have a disadvantage.

TIPS FOR EFFECTIVE RESPONSE PLANS:

- Keep all instructions short and clear
- Do not assume that anyone enacting your plan understands corporate or departmental jargon and colloquialisms
- Use acronyms and mnemonics for disaster response training. In documentation to follow during a response, clearly spell out any acronyms or symbols.
- Employ iconography. For instance, if process B depends upon the completion status of process A, use a large icon of a stop sign or similar callout at the end of process A.
- Where iconography does not suitably call attention, use textual clues. For instance, large "Warning" boxes in a bold color and using large fonts.

Research or brainstorm acronyms for problems that are likely to occur and that require uncommon activities. For instance, most people have never used a fire extinguisher. You might create literature on using fire extinguishers that includes the common "PASS" acronym. Then have pictures, or better yet, a video that matches "P" to pulling the extinguisher's pin, "A" to aiming at the base of the fire, the first "S" to squeezing the trigger, and the final "S" to sweeping the nozzle back and forth. If you include directions with extinguishers (highly recommended), you can have a short tag with these items spelled out. Do not assume that anyone remembers (or even attended) the training.

You can create your own acronyms. As an example, you could create a fire protocol and call it "The three Es (EEE): Extinguish, Evacuate, Escape". Your training would expand these to "extinguish the fire if possible", "evacuate others", and "escape yourself". If drilled, people have a better chance of remembering what to do when they have some simple mnemonics to work with.

Do not overuse these memory tools. For instance, if you search the Internet for "emergency response acronyms", you will find lists that contain government agencies, response programs, and common phrases used when response personnel communicate with other.

People who work in disaster response full time might remember these, but no one else will. Have only a few and try to have them on printed literature near any equipment that relates to the situation that they address.

Above all, remember that some catastrophes affect more than just your business. Some of your staff may have had their lives upended. Many will have their own things to recover from. Business continuity planning must include flexibility for employees.

17.2.2. WORKING WITH DISPLACED EMPLOYEES

Catastrophes can render a site unusable for a significant period of time. Plan in advance what the employees will do.

If you will redirect staff to an alternative site, ensure that everyone knows the location. Include a reminder in the notification system. Importantly, have someone verify the viability of the site before sending everyone there. You can use an initial message that informs staff of the situation and instructs them to wait for further notifications. Once someone deems the secondary site usable, send a follow-up notification.

Remember that, just like surviving a disaster, an interrupted work routine causes distress. People will arrive late, get lost, and need to leave at atypical times of the day to reach appointments that normally needed only a few minutes of travel time. Plans should expect erratic attendance patterns while employees adjust.

17.2.3. WORKING WITH OFFSITE EMPLOYEES

Many positions began transitioning to remote work years ago. The need for isolation brought on by the COVID-19 crisis dramatically accelerated that transition. As long as remote employees still have some system to connect to and were otherwise not impacted by the event, little changes for them. Include them in communications about the situation and remember that the conditions will have some effect on them.

You may choose to have some employees begin working from home that would normally commute to a physical location. An effective transition from on-premises to at-home work requires a substantial amount of advance planning, especially if you do not currently have a formal remote work policy. Your organization will need to answer many questions:

- Do employees use their own hardware? Does the company provide equipment? Does the company reimburse?
- How will remote employees maintain communication?
- Will you pay for a premium collaboration service, such as Zoom?
- Will you enforce a requirement of a particular service?
- Will work hours change? Flex?
- Will users connect via a VPN? VDI? Microsoft Remote Desktop sessions? A Citrix solution? Something else?
- Do your systems have sufficient capacity to support the potential number of remote workers?

Some employers worry that productivity will drop from remote workers. Studies have shown that this concern has no real basis. However, if the source event was a major disaster, the psychological effects and any damage to employee's property will impact their work. Even without that, transitioning from the office to the home does take time. Balloon some adjustment flexibility into your plan.

17.2.4. NOTIFYING AND ACCOUNTING FOR EMPLOYEES

From the instructions in the first part of this book, your plan should already include notification trees and contact methods. Response documentation must include an accounting system. Small businesses can do this informally. Medium-sized businesses can require employees to check in with their supervisors who in turn report to a central command structure. Large businesses can do the same in a tree structure or make use of call-in telephone numbers.

Define processes for unreachable employees in the context of a widespread disaster. You can use things like "unknown", but that cannot be a final disposition after attempting a single phone call. Establish a schedule for retries. When multiple attempts to locate an employee fail and you can no longer devote resources to them, report them as missing to authorities. To reiterate, do that only when you have reason to believe that the person might be in danger. Do not call the police if a systems administrator doesn't answer a text message about a server crash. While that might seem obvious, make the conditions very clear in your documentation.

17.3. DESIGN GUIDELINES FOR BUSINESS CONTINUITY PROCESSES

The overall goal of this section is to cover the role and importance of people in a disaster recovery plan. Use this information as a starting point and guidance system for building your own documentation.

The actual processes to include must come entirely from your business experts. Start with these high-level points:

- Guidelines for managers and executives to decide between a short interruption that warrants no major response, an event that justifies switching to full downtime procedures, and a genuine disaster that requires an orchestrated response
- Minimal and full downtime procedures
- Employee immediate response, notification, and accounting procedures
- Relocation activities
- Remote work policies and practices
- Recovery processes

The recovery process portions will need a lot of space. They should only start after the immediate problems have passed. Recovery will include installing replacement systems, restoring data, ordering equipment, organizing contractors, filing insurance claims, notifying customers, and any other activities that staff indicate.

<https://www.forbes.com/sites/larryalton/2017/03/07/are-remote-workers-more-productive-than-in-office-workers/#70bc303231f6>

<https://www.gallup.com/workplace/283985/working-remotely-effective-gallup-research-says-yes.aspx>

CHAPTER 18:

TESTING DISASTER

RECOVERY SYSTEMS



Systems fail. That's an unfortunate fact of computing. If things were otherwise, this book would have been much shorter. However, what happens when the systems that we build as a failsafe against failure suffer failures of their own? Just as you watch over your production systems, you must also put forth the effort to ensure that you always have at least one complete, undamaged, and retrievable copy of your data. You cannot stow away your backup data and hope for the best.

I once took a phone call from a contact whose sole line-of-business application server lost both of its mirrored drives before they knew that they had a problem. Because this business closed at five PM and didn't re-open until 8 AM, we had set them up to perform fully daily backups that automatically ejected the tape upon completion. While trying to help them, I learned that they had hired an onsite administrator for short time. During his stay, he switched them over to weekly full backups with daily incremental jobs. He also disabled the automatic eject feature. When he left, he neglected to train anyone on the backup system. No one knew that they were supposed to change the backup tapes daily. Every night for over a year, their backups overwrote the previous night's data, usually only with a tiny subset of changes. So, when they needed it most, the backup data was not there.

**" SYSTEMS FAIL.
THAT'S AN
UNFORTUNATE
FACT OF COMPUTING. "**



“ I once took a phone call from a contact whose sole line-of-business application server lost both of its mirrored drives before they knew that they had a problem. Because this business closed at five PM and didn't re-open until 8 AM, we had set them up to perform fully daily backups that automatically ejected the tape upon completion. While trying to help them, I learned that they had hired an onsite administrator for short time. During his stay, he switched them over to weekly full backups with daily incremental jobs. He also disabled the automatic eject feature. When he left, he neglected to train anyone on the backup system. No one knew that they were supposed to change the backup tapes daily. Every night for over a year, their backups overwrote the previous night's data, usually only with a tiny subset of changes. So, when they needed it most, the backup data was not there. ”

Case Study



To safeguard yourself against problems such as these (and the attendant horror stories), build testing schedules into your data recovery plan. Assign staff to perform those tests. At the scheduled meetings to update the disaster recovery documentation, require testers to provide a synopsis of their test activities. This will give your organization an external accountability control over the backup process. You will have a chance to discover if no one has performed a backup without the need to encounter an emergency.

18.1. TESTING BACKUP DATA WITH RESTORE OPERATIONS

You have a very straightforward way to uncover problems in your backup strategy: try to restore something. Most modern backup software has some built-in way to help.

Exact steps depend upon your software. Follow these guidelines:

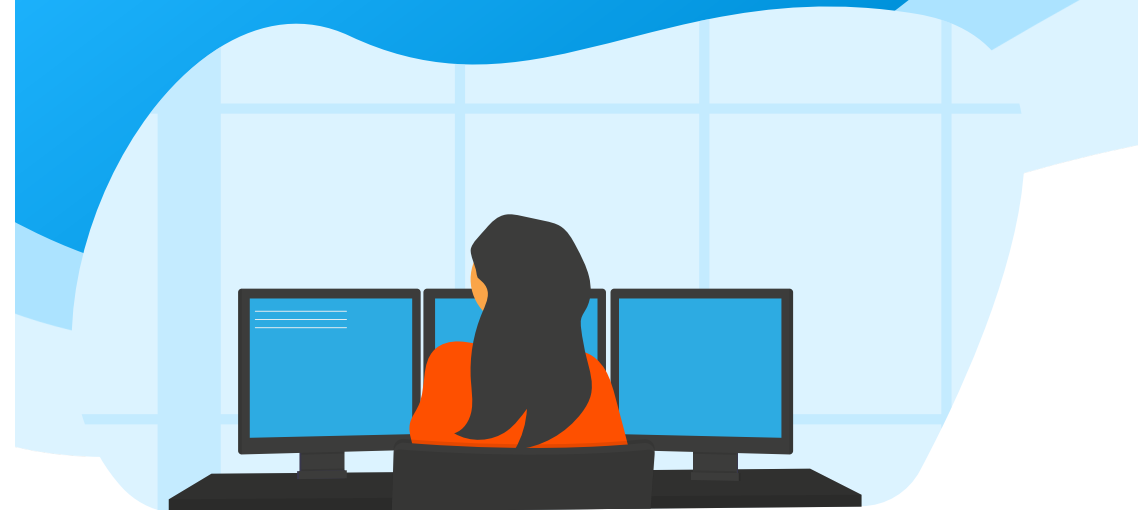
- Redirect to an alternative, non-production, "sandbox" location. If your backup somehow has corrupted data, you don't want to find out with a "test" overwrite of valid production data. If you're ensuring that you can retrieve a virtual machine, you don't want it to collide with the "real" system.

- Test restoring multiple types of data. Bring back individual files, entire SQL databases, domain controllers, virtual machines, and any other type of logical unit that you rely on.
- Rotate the items that you check at each testing interval.
- Test from more than one repository.
- Verify the restored information by accessing it yourself.
Do not interpret a successful restore operation as proof that the data survived.

The major problem with this type of testing is its scope. You will always work with a sample of the backed-up data, not the entire set. You will likely be instructed to test the most important business components at every step. Make sure to test representative items from lower priority systems as well.

Data corruption is sneaky. Unless your equipment suffered a major failure or someone accidentally degaussed the wrong drive pack, the odds are that you will never uncover any degradation or errors. Take heart; since you probably won't find any corruption, you probably won't ever need to restore anything that happened to become corrupted. However, do not take this condition as a reason to skip test restores.

**“ YOU HAVE A VERY
STRAIGHTFORWARD WAY
TO UNCOVER PROBLEMS
IN YOUR BACKUP STRATEGY:
TRY TO RESTORE SOMETHING. ”**



I always insist on multiple full copies of backup data as the primary way to protect against small-scale corruption. Unless the production data is corrupted, there is almost no chance that two distinct copies will have problems in the same place. The purpose of a test restore is not to try to find these minor errors. We are looking for big problems here. A routine test would have caught the problem in the anecdote that opened this section. If someone accidentally (or maliciously) unchecked the option to back up your customer database, you will notice when you attempt a test restore. If a backup drive has a mechanical failure, you will either get nothing or blatantly corrupted data from it.

Use manual test restores to spot-check for corruption, verify that backup covers the data that you need, and that your media contains the information that you expect. To shield against the ever-present threat of ransomware, only use operations that can read from the backup (not write to it) and work within an isolated sandbox. These types of tests are the only true way to verify the validity of offline media.

18.2. TESTING BACKUP DATA WITH AUTOMATED OPERATIONS

Manual tests leave you with the problem of minor data corruption. Backup data sets have only increased in size through the years, compounding the problem. With the fortuitous gradual deprecation of tape, backup application vendors have seized on opportunities to add health-check routines to their software. They can scan through data to ensure that the bit signatures in storage match the bit signatures that they initially recorded.

Features like these call out the importance of a specific distinction in the usage of the word "automation". It certainly applies to a process the computer performs in order to remove the burden from a human. It does not necessarily mean "happens automatically". For that connotation, stick to the word "scheduling". In this context, do not assume that any mention of automated testing in your backup program's interface means that it will handle everything without your help. Some programs have that capability, but this will never be a "set it and forget it" activity.

End-to-end data validation is time-consuming and an intense load on resources. Otherwise, we would happily do it ourselves and not need

the backup program's help. Also, some of them block other backup operations while in progress. So, such processes need three things from you:

- A specific start time
- Sufficient time to complete
- A human-led procedure for verifying and recording the results

In a few cases, especially at smaller organizations, there may not be a major reason to avoid scheduling the start of a validation job.

The concern for its impact on any overlapping backup jobs. If you have a slow backup system that will require multiple days to process everything, then validation is probably not feasible. It is better to capture as many backups as possible and rely on manual spot-checking than to allow an automated verification process to disrupt jobs. At best, these automatic checks can add some peace of mind. But they will never replace the manual work.

You also have the option to create custom checks. You can use scripts or software tools to scan through temporarily restored data. It can look for problems or ensure that it can find expected data. You can potentially interface it with your backup software. For instance, you can restore

data to an alternative location and have the backup create another copy near it. A comparison tool can show where the data differs. Always keep ransomware top-of-mind. If you set up something like this, no process should have write access to production data and the backup location.

Systems administrators tend to be a clever, intelligent group. When we read guides like this, many of us think to ourselves, "I can script that!" That's great, don't let anything here discourage you. A virtualized domain controller that runs dcdiag on itself after a restore operation? A SQL Server that runs through DBCC on restored databases? Your own system for creating and validating checksums on your most important file repository? Things like that are awesome! You can never have too many helping checks. However, you can never rely solely on them, either. In the event of any kind of failure that backup does not recover, management will ask, "Did you verify that yourself?" They will not recognize the value of your scripts. Your skills will not impress them. A solid track record of twenty years without a failure will not make any difference. Worse, if a data loss exposes your company to litigation, judges, attorneys, and jurors will care even less. You cannot fully prevent all mistakes and you cannot personally check every last bit. You must employ the sort of manual processes that non-technical

“ Systems administrators tend to be a clever, intelligent group. When we read guides like this, many of us think to ourselves, “I can script that!” That’s great, don’t let anything here discourage you. A virtualized domain controller that runs dcdiag on itself after a restore operation? A SQL Server that runs through DBCC on restored databases? Your own system for creating and validating checksums on your most important file repository? Things like that are awesome! You can never have too many helping checks. However, you can never rely solely on them, either. In the event of any kind of failure that backup does not recover, management will ask, “Did you verify that yourself?” They will not recognize the value of your scripts. Your skills will not impress them. A solid track record of twenty years without a failure will not make any difference. Worse, if a data loss exposes your company to litigation, judges, attorneys, and jurors will care even less. You cannot fully prevent all mistakes and you cannot personally check every last bit. You must employ the sort of manual processes that non-technical people understand. An answer of, “That particular data set rotates to human validation every three months and the disaster hit at just the wrong time,” would help to pull attention away from you. At the same time, doing the work to properly prepare against such conditions can help to ensure that you never have to face them. ”

people understand. An answer of, "That particular data set rotates to human validation every three months and the disaster hit at just the wrong time," would help to pull attention away from you. At the same time, doing the work to properly prepare against such conditions can help to ensure that you never have to face them.

Most importantly, remember that automated routines can only supplement manual, personal operations. They will never stand on their own.

18.3. GEOGRAPHICALLY DISTRIBUTED CLUSTERS

Due to the logistics involved, few organizations will utilize geographically distributed clusters (sometimes called stretched clusters). Combined with synchronously replicated storage and a very high-speed site interconnect, they offer a high degree of automated protection. Properly configuring one requires many architectural decisions and an intimate understanding of the necessary hardware and software components. This book will not dive that far into the topic.

The basic concepts of a geographically distributed cluster:

- These clusters are built specifically for business continuity. They are not an efficient solution for making resources available in two sites at once.
- Geographically-clusters must use synchronously replicated storage for effectiveness
- Administrators often set resources so that they operate only in location except in the event of a failover. Individual resources are configured to run in the location where they are closest to their users, if possible.
- Each location should be able to run all cluster resources
- Each location should have sufficient capacity to allow for local failures. As an example, if the resources on your cluster require 5 nodes to operate and you want N+1 protection, then all sites require 6 nodes.
- Resources must be prioritized so that, in the event that the cluster does not have enough nodes, that the most important resources remain online

If you have created such a cluster, you must periodically test it to ensure that it can meet your criteria. Because the sudden loss of an inter-site link or storage device will almost certainly trigger a resource crash, it would be best to perform these tests with only non-production resources. The easiest way to accomplish this goal is to schedule downtime for the entire system, take the production resources offline gracefully, and do all of your work with test resources. If your protected resources do not allow that much downtime, then you can use cross-cluster migration tools to evacuate the resources to other clusters during the test.

In many cases, you will not have any good options available.

Alternative options:

- Use test systems with the same fundamental configuration as your production systems and test with those
- Remove a node or two from each site, create a secondary cluster from them, perform your testing, then rejoin the nodes to the production cluster

These alternatives have problems and risks. Test systems let you know how a site failure would theoretically work, but they do not prove that your production cluster will survive. Individual nodes could have undiscovered problems, recently added resources might take you slightly over your minimum required node count, and the cluster configuration may not function as expected in a site failure.

A common problem that's not immediately obvious without testing is that your cluster configuration might take everything offline in all sites because it can't establish a quorum. Worse, it might keep disconnected sites online simultaneously, running on storage units that can longer synchronize, leading to a split-brain situation. You will catch those problems in pre-production testing, but changing conditions can affect them (adding nodes, unplanned outages in multiple sites, etc.).

18.3.1. TESTING GEOGRAPHICALLY DISTRIBUTED CLUSTERS

When establishing the tests to run, start with probable events. Look specifically at the resources that the cluster operates and poke at their weaknesses. A few ideas:

- Take a storage location offline without notifying the cluster
- Unplug network cables
- Disable the LAN for the cluster nodes in one site
- Reboot the device that connects the sites. Assuming redundant links, try them separately, then all in one site at the same time.

Use your imagination. Also, don't forget to perform the same sorts of tests that you would for a single-site cluster (node removal, etc.).

18.3.2. COPING WITH THE CHALLENGES OF GEOGRAPHICALLY DISTRIBUTED CLUSTERS

In reality, most organizations cannot adequately test production clusters of any type. Do not use that as an excuse to do nothing. You always have some things to try. For instance, if you skip the storage tests, you can perform validation on a Microsoft failover cluster almost any time without impact. Research the clustering technologies that you use. Look to user forums or support groups for ideas.

Make sure not to over-promise on the capabilities of geographically distributed clusters. Take time to understand how to deal with

conditions such as the aforementioned quorum outage. Above all else, take special care to understand how your storage and clustering technologies react to major failures. Do not rely on past experiences or general knowledge. Use strict change tracking and review the build at each disaster recovery update cycle.

Backup will always be your best protection against anything that threatens your cluster. Make certain that clusters have adequate and operational backup.

18.3.3. TESTING REPLICATION

Replication technologies are built specifically to deal with failovers, so they are not as difficult to test as geo-clusters. Testing almost always involves downtime, but usually has manageable impact.

Unlike geographically distributed clustering, testing failover of a resource that shares a common platform with others usually tells you enough that you don't have to test everything.

When you first build a replication system, work through a complete failover scenario. This exercise helps you and your staff more than the technology. Replication and failover do not always work the way that administrators assume. If you see the entire procedure in action,

then you will have a much better understanding of what would happen to you in a real-world situation. Document anything that seemed surprising. If you find a blocking condition, shift the parameters. Continue testing until failover works as seamlessly as possible before going into production.

Many times, small organizations set up Hyper-V Replica without a complete understanding of the technology. They follow the instructions and the prompts, and everything appears to work perfectly. Then, they try to fail over to their secondary site, and nothing works. On investigating these problems, we discovered that many of them were replicating a domain controller virtual machine and had no other DCs online at the secondary site during the failover. When the domain controller went offline, the secondary site could no longer authenticate anything, including the Hyper-V Replica operations.

Incorrectly configuring failover is a common occurrence and why we discussed the importance of using application-specific replication where possible in an earlier section on configuring replica.

“ Many times, small organizations set up Hyper-V Replica without a complete understanding of the technology. They follow the instructions and the prompts, and everything appears to work perfectly. Then, they try to fail over to their secondary site, and nothing works. On investigating these problems, we discovered that many of them were replicating a domain controller virtual machine and had no other DCs online at the secondary site during the failover. When the domain controller went offline, the secondary site could no longer authenticate anything, including the Hyper-V Replica operations. ”



This also stresses the importance of testing; sites that tried to fail over their replicas as a test fared much better than sites that didn't try until catastrophe struck.

For a clear example, consider virtual machines protected by Hyper-V Replica. If you have a test virtual machine that spans the same hosts and same storage locations as your production virtual machines, then start with it. Provided that all conditions match, it will give you a good idea of what would happen to the production virtual machines. If you have any low-priority production virtual machines, or some that you can take offline for a while, test with those next.

When possible, test all resources. Failing over a sample does not uncover problems like data corruption in the targets. Unfortunately, testing with the real systems might not catch it either and failing back to the main site could very well ship corrupted data back to the source. Use your monitoring tools and capture a good backup before testing any production systems. If possible, take the source resource offline, wait for synchronization to complete, and capture a hash of the files at both locations. For file-based resources, you can use [PowerShell](#):

```
Get-FileHash -Path C:\Source\File.txt -Algorithm MD5
```

Use the same command on the destination file; if the results don't match, something is not right. Do not perform a failover until you have found and eliminated the problem.

Note: the MD5 algorithm no longer has value in security applications, but still works well for file comparisons due to its speed and the near-zero likelihood of a hash collision on two slightly different copies of the same file.

Once you have successfully failed a resource to your alternative site, bring it online and make certain that it works as expected. Some configurations require advanced settings and comparable testing. To return to our Hyper-V Replica example, you can set up the replica virtual machines to use a different IP address than the source VMs. If you have done that, ensure that the replicas have the expected connectivity.

After testing at the remote site, fail the resource back to the primary. Depending on the extent of your testing, it may take some time for any changes to cross. Return it to a service state and ensure that it works as expected. Have your backup data ready in case.

18.4. DO NOT NEGLECT TESTING

We all have so much work to do that testing often feels like low-priority busy work. We have lots of monitoring systems to tell us if something failed. If nothing has changed since the last time that we tested, would another test tell us anything? Regardless of our workload and the tedium of testing, we cannot afford to skip it. We humans tend to predict the future based on the past, which means that we naturally expect functioning systems to continue to function. However, the odds of system failure increase over time. The only sure way to find problems is through testing.

You cannot properly prepare for every possibility. Testing helps you to uncover the quirks of programs and the holes in processes. Your organization's data needs you to do the right thing.

During the drafting of this book, I nearly lost almost half of my work. I relied on a synchronization process to keep everything up to date between the different computers that I use for writing. Each maintains their own local backups that update when I save or shutdown my authoring software. Unfortunately, the synchronization process got confused and replaced completed sections with empty ones.

Due to the rapid changes, my backups rolled over before I discovered the problem.

I was fortunate enough that the synchronization system placed the "old" files in the Recycle Bin instead of overwriting them. Had I not taken the time to periodically scan through previous sections "just to double check", a later synchronization might have replaced the Recycle Bin files as well.

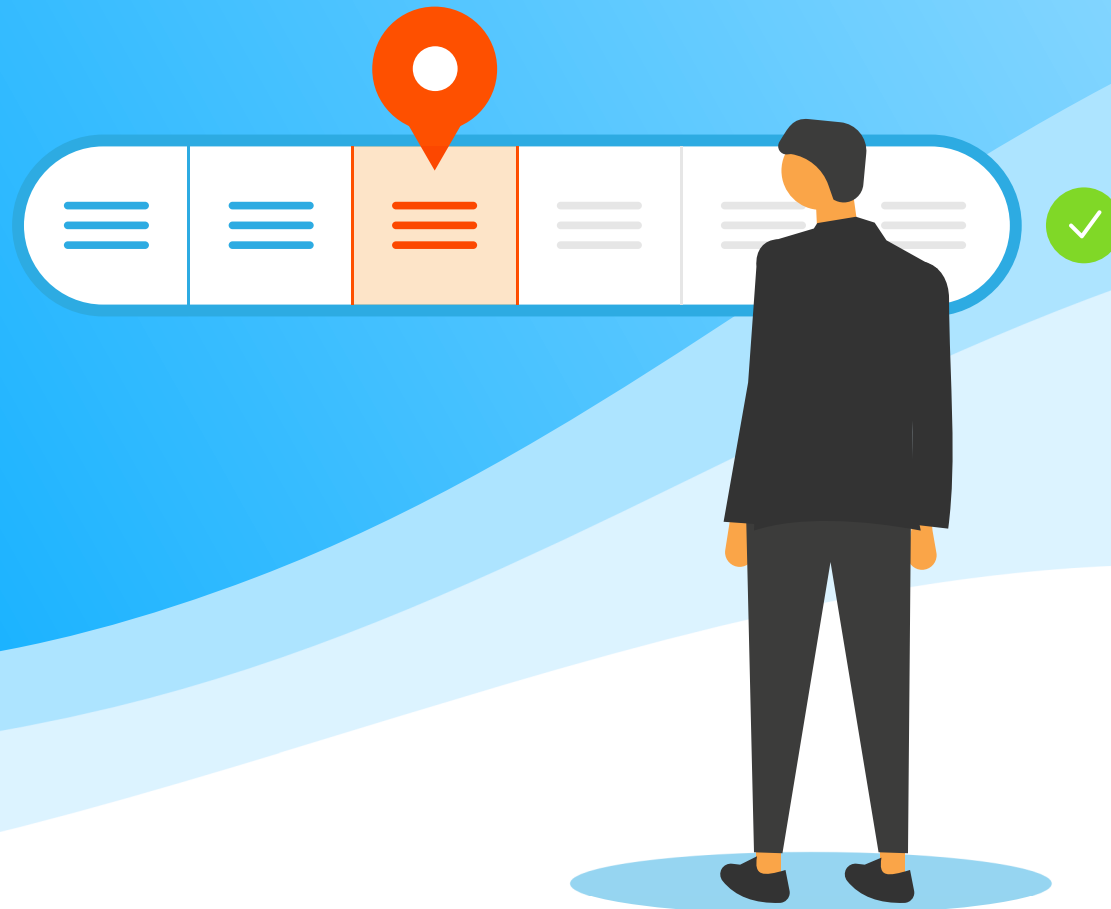
Even people with decades of experience and subject matter expertise encounter unexpected data loss. Systems do not always behave as we expect. Testing can save us.

" During the drafting of this book, I nearly lost almost half of my work. I relied on a synchronization process to keep everything up to date between the different computers that I use for writing. Each maintains their own local backups that update when I save or shutdown my authoring software. Unfortunately, the synchronization process got confused and replaced completed sections with empty ones. Due to the rapid changes, my backups rolled over before I discovered the problem.

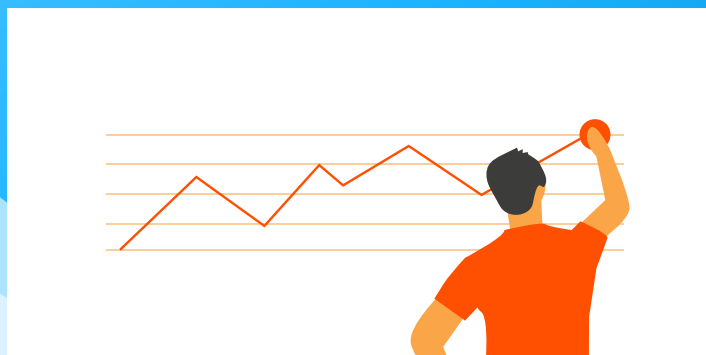
I was fortunate enough that the synchronization system placed the "old" files in the Recycle Bin instead of overwriting them. Had I not taken the time to periodically scan through previous sections "just to double check", a later synchronization might have replaced the Recycle Bin files as well.

Even people with decades of experience and subject matter expertise encounter unexpected data loss. Systems do not always behave as we expect. Testing can save us. "

CHAPTER 19: THE ONGOING PROCESS



“ THE LONGER YOUR ORGANIZATION GOES WITHOUT SOME SORT OF SERIOUS CHALLENGE, THE MORE IT BUILDS A FALSE SENSE OF CONFIDENCE. ”



Back in [chapter one](#), we said that “disaster recovery is an ongoing process, not a one-time event”. Take it to heart. Everything changes. Staff move on to other opportunities, systems reach end-of-life, vendors overhaul their products, buildings find new purpose. Parts of the plan that you built will go stale. If left for too long, all your hard work will eventually degrade into a minimally useful information glob. You must maintain your business continuity plan just like everything else.

19.1 KEEP OTHERS INVESTED

Because of the whole-organization nature of a disaster recovery plan, you need continuing executive support and participation from other business units. To repeat a phrase from chapter two, “success breeds a success mentality”. The longer your organization goes without some sort of serious challenge, the more it builds a false sense of confidence. Sadly, the cliches of “nothing lasts forever” and “all good things must come to an end” will eventually strike. Between delicate technology and human error, it can happen at any time. Don’t get complacent.

19.1.1 THE FEAR FACTOR

For a quick way to spur interest, you can demonstrate looming disaster and near-misses.

A FEW IDEAS:

- Report on failures that IT detected and corrected before disaster struck. Make it clear that failures do occur within your organization even if no one but IT sees them. Reiterate that even with fault tolerance and a highly proactive and reactive team, no system can withstand everything.
- Remind others of the high cost of downtime. All current and new projects should have a metric that estimates the cost of an hour of downtime. That helps to make the risk more tangible.
- When convenient, slip in a mention of disasters that have struck other organizations, especially those in the same industry. Always use these examples as cautionary tales, never as a derogatory commentary on another firm. As an example, “A business like ours in another town was recently shut down by ransomware. We should double-check our safeguards and recovery processes to ensure that we can survive if that happens to us.”

**“ IT SYSTEMS SUPPORT
THE BUSINESS;
THEY DO NOT DEFINE
OR CONTROL IT. ”**



Do not overuse these tactics. The fear approach has limited usefulness. The others in your company only need to remember that no one is immune to disaster and that everyone must remain vigilant.

19.1.2 THE COMMUNITY APPEAL

Anyone in IT knows that all things related to technology eventually find their way under the IT department's umbrella. If you talk about business continuity and disaster recovery from a systems and data perspective, you will quickly lose attention from everyone outside IT. Their jobs do not require them to maintain and protect the blinky lights and Excel files; that's why the company has an IT department.

[Chapter 13](#) explored the concerns of non-data services and the sorts of processes that apply to them. Review that section and find ways that they connect to your non-technical colleagues. Of course, you will take all necessary steps to keep systems online and bring them back if something happens because that's your responsibility. However, they must also take responsibility for the protection and recovery efforts their departments need. Remind them that IT systems support the business; they do not define or control it. Learn how that fact applies in your organization. Find the intersections and points of divergence between IT and non-IT teams. Use them in your appeals.

For instance, "It doesn't help much for IT to bring back the manufacturing floor's computing systems if you have no machines to work with or supplies to feed into them."

Business continuity and disaster recovery affect every part of an institution. "We are all in this together" holds true.

19.1.3 IMPLEMENT THE MANDATE

In many cases, simply taking the lead works well. Business continuity matters and no one can argue against that. Instead of approaching it like some sort of new initiative or bold course correction, simply add it in with all the other things that must happen. The company has quarterly and annual financial meetings, monthly budget meetings, weekly staff meetings, and now it has bi-annual business continuity meetings as well. Like any other meeting, participants are expected to show up on-time and prepared. If a department does not send a delegate, then repeatedly request updates from the department head until you get them.

The internal politics of your organization will control the feasibility and effectiveness of this approach. However, in most cases, when someone identifies an unavoidable risk and volunteers to take charge and organize the solution, others are content to join in and follow along.

19.2 THE ROAD AHEAD

Building a full business continuity plan requires substantial time and effort. It's much easier to keep it going once you've got it started, but it will never be pain-free. For people that have found the need to execute a disaster recovery plan, it's worth every minute.

If you can get your plan ready before disaster strikes, that makes you the hero. We all wish that nothing bad ever happens, so with the best of luck, your deeds will remain uncelebrated. Many of you will eventually take that call.

Be prepared, it's better to be safe than sorry.

This concludes the main body of The Backup Bible. In the final part we will be covering backup and disaster recovery concerns for Managed Service Providers. Although the foundations and basics of backup theory discussed thus far will remain largely the same in an MSP environment, there are several important additional aspects when delivering Backup as a Service (BaaS).

PART 4: PROVIDING BACKUP SERVICES TO MSP CUSTOMERS



As we've already covered in detail, planning, creating and maintaining a robust backup and disaster recovery strategy for a single organization is not a straightforward task but what if you're responsible for the management of backups across multiple clients? Not only are your responsibilities increased due to there being additional environments but conducting backups as a Managed Service Provider (MSP) comes along with numerous challenges. Such as:

- Multiple disparate and unique environments
- Physical separation of each backup domain
- Management of Backup/Recovery operations at-scale
- Customers with varying degrees of complexity
- Regulatory compliance concerns
- More complex SLAs and MSAs
- More training work for your team
- Mistakes can directly impact your bottom line
- And more...

This list isn't designed to scare you necessarily but rather reveal the truth behind providing genuine backup services. Many services providers simply see backup and recovery as a "simple" business to "tack on" as a revenue line-item. This couldn't be further from the truth. Effectively managing backups for your customer base is a very REAL responsibility and due to the sensitive nature of data loss, effective backup services can make or break trust with your customers.

This part of the Backup Bible is designed to guide fledgling MSPs through the process of setting up a Backup as a Service (BaaS) offering and running it at scale. That's not to say that new MSPs are the only audience for this part of the Backup Bible. This text will also serve as a good sanity check for those already established with an existing backup practice that simply want to make sure they're current and doing the right things. It's the wise business owner who will always look to improve, and we'll help you do it!

As we step through this section, we'll cover all the above challenges/concerns. However, before we get into that, let's start with the why.

CHAPTER 20: WHY SHOULD YOU OFFER BACKUP SERVICES TO YOUR CUSTOMERS?



There is one single thing that you do as an MSP that's more important than any other service you offer. More important than security hardening, more important than email and collaboration services, and even more important than patching. It's backup services. Especially in the SMB space, there are always many clients who don't understand the risks they take without proper backups in place.

And as has been mentioned, you don't really want backup, what you want is verified restores so that you know that you can get the data back in case of a disaster.

For example, the scariest situation I've ever seen was arriving at a new client and the secretary proudly proclaiming that she swapped the server's external backup drives every day and kept the others off-site, "because you can never be too careful". This sounded great until we saw that the backup job hadn't run successfully for the last three months! She didn't know how to check that the backups were working. In fact, no-one at this business did. As an MSP your highest priority task is to make sure that your clients never end up in the same boat.

Today, most clients are aware enough that they expect you to deliver backup services for them but most clients still lack the expertise to do it themselves. Understanding the importance and not having the internal

capability means they're willing to pay for it, providing an additional revenue stream for your MSP.

Another story from the trenches concerns a medium sized (450+staff) customer and their (single) on premises Exchange 2013 server. A Cumulative Update (CU) installation failed (which never happens right?), causing services to start failing. "No problem", we said, "let's just restore this verified backup we took before we started the installation". Only problem was, restoring ~750 GB of databases took 30+ hours, halting all email for 2 days. So not only do you need backup and verified restores, but you also need *performant* backups and restores.

There's another big benefit of offering solid backup services as part of your MSP package and that is that it builds trust with clients. There's nothing quite as comforting as a panicking user calming right down as you restore their mistakenly deleted file remotely while they're on the phone.

CHAPTER 21: CONSIDERATIONS FOR BUILDING A BAAS PRACTICE



There are several things to consider when you build or refine your BaaS practice.

Note: The first one is probably to be aware that BaaS can also mean “Backend as a Service” for mobile applications in case you end up in a confusing conversation with a client.

You also need to think through what services you’re offering and be very clear in your MSP contracts exactly what is and is not included. If you aren’t clear, don’t be surprised when the client assumes you’re backing up ALL client devices, servers and ALL cloud services, every five minutes and that you can restore anything immediately. Customers will assume the world unless told otherwise.

BAAS CHARACTERISTICS INCLUDE:



CLOUD BASED
(ALTHOUGH YOU CAN
AUGMENT IT WITH A LOCAL
APPLIANCE/SOFTWARE)



**SCALABLE
AND MULTI-TENANT
FROM THE GET-GO**



**CONFIGURED
AND MAINTAINED
IN THE CLOUD**



**CAPACITY MANAGEMENT
REDUCED WITH
INCREASED ELASTICITY
IN CLOUD STORAGE**



**SHIFTS BACKUP
EXPENSE FROM
CAPEX TO OPEX**

21.1 WHAT DATA SHOULD BE PROTECTED?

At the most basic level you should offer simple, daily backups of on-premises servers. Again, the devil is in the details, you'll need to make sure your chosen backup product can handle their server applications (Exchange, SQL Server, SharePoint of course but how about Oracle DB or an Apache web server?) and their Operating Systems (Windows is a given, but what about Linux and which distribution?). Their virtualization platform (VMware, Hyper-V etc.) also must be supported. If they run workloads in an IaaS or PaaS in a public cloud, you'll need to account for protecting those workloads as well. There are many times when a single daily backup isn't enough so make sure you make allowances for more frequent backups.

Another decision point here is if you're only offering offsite backup with the data being copied to the cloud immediately or if an on-premises cache is part of the solution. The latter of course improves recovery speeds considerably (for most recoveries which don't involve whole site failures) but adds complexity.

Further you need consider client endpoint (Windows, MacOS, Linux) protection, including work from home devices that may need protection.

There's nuance here - if you have a call center full of PCs that have a standard image on them that you can deploy easily, with all data stored on servers/in SaaS services, backup of endpoints isn't necessary. On the other end of the spectrum, you may have the CEO who's been working from home for the last year and stores sensitive corporate data on his local PC for "security reasons" and who will be very annoyed when his SSD fails.

Apart from defining **what** you'll backup and how often, you also need to define (and test!) how **fast** you can recover files and folders deleted by mistake, not to mention whole systems and server applications as you'll need to manage client expectations. You must also define SLAs for backup and restore operations in your contracts, which we'll talk about a bit later.

21.2 TAKING IT FURTHER THAN JUST BACKUP/RECOVERY

Finally, you'll need to investigate whether you're offering plain backup services or more comprehensive disaster recovery solutions. This involves continuous replication of data and the ability to recover multiple systems quickly in case a widespread ransomware attack

leaves files on many systems encrypted or a complete site failure due to a significant data loss event.

If you're going to offer onsite plus offsite backups or full DR services, are you going to provide the hardware for the on-premises storage? Or is the client going to provide the hardware? Most importantly, if you're offering this as BaaS, is the cost of the hardware built into the monthly price or an upfront cost for the client? If it's the former you need to factor in that you'll be out of pocket for some time as you pay for the backup hardware and then charge the client for it over the months and years to come. Better yet, if this is the route you're going and you have a number of customers lined up already for backup services, you may be able to aggregate the cost of the service across the starting customers. This will of course vary on a case-by-case basis.

21.3 STORAGE CONSIDERATIONS

There are many different aspects of putting together a good BaaS package but at the core should be the decision whether to purchase, maintain and upgrade costly storage hardware or rely on cloud storage entirely. This will of course be somewhat dictated by regulatory requirements of your target customers, but cloud storage (for offsite

backups at the very least) should be considered. [Cloud storage](#) today is very cost effective, highly available and you can pick your storage tier (hot for quick access, cool for cheaper access as long as you don't access it very often, archive for VERY cheap but slow access). To provide DR ready on-premises storage that replicates between datacenters is incredibly expensive and complex - in the cloud it's a single tick box to replicate to another region.

Do your due diligence when investigating different storage services, [you need to factor in egress costs](#) (it's always free to upload data to the cloud, but generally you have to pay to download it), plus replication costs from region to region if you're opting for that level of protection.

If you're going the cloud route you have to pick a platform: Azure, Amazon Web Services (AWS) or Google Cloud Platform (GCP), or smaller players such as Wasabi, IBM or Oracle. Of the big three, Google has recently expressed concern over its profitability and whether to continue so we would probably not rely on them.

AWS and Azure are both valid options and as an MSP it comes down to your familiarity with them, if they cover all the geographical regions your clients have a presence in and if they work with your chosen backup/replication software.

Some backup vendors, like [Altaro Office 365 Backup for MSPs](#), “bake in” the storage / egress and any other costs and abstract this away from you, which provides a simpler experience for you.

To meaningfully make the storage decision you need to know what you’re backing up as that may impact what storage options are easily usable.



Use [the Customer Backup Requirements Template](#) in the Appendix to help identify backup storage needs for your customers.

21.4 SELECTING A BACKUP PRODUCT

Based on all the factors above and considering what most of your clients run and use, it’s time to pick a backup product/service.

Don’t start with price - this is a big mistake. Start with ensuring that the service covers everything you need, including geographical coverage and storage capacity. Then look at the security features on offer - data (including backups) is an attractive target today - just look at ransomware and its impact on the world. As a minimum, the service should offer encryption of all data at rest, with a modern encryption

standard such as AES-256, preferably with keys under your control. Of course, all data in flight needs to be encrypted as well (TLS 1.2+). It should also offer two factor authentication, access to your client’s data must be protected by more than a single phishable or guessable password.

Another very important piece of criteria is support - 24/7 x 365, preferably local to your region, and quick to answer is a must. The last thing you want to hear when your client has a disaster, and you need help from the vendor is “you’re in the queue and your call will be answered in 25 minutes”.

If your clients (or your business) need to adhere to regulations, ensure that the vendor and the service are compliant with the ones you need. Reporting is another feature you’ll need – it’s important you can easily show the client that backups are working through regular, scheduled reports.

Finally, you’ll want to investigate the SLAs and backup / replication / restore times the vendor offers.

21.5 PRICING MODELS

A major factor to a successful MSP business is being clear on exactly what's included and what's not included in your package(s). You need to ensure that you include basic backup in your base offering.

There should never be a time when a “penny pinching” client ends up not having selected backup and now their data is lost. However, for more complex on-premises, extended backup services and DR solutions as outlined above you should be able to charge extra for these as a value-add.

If there's a need for substantial on-premises backup infrastructure you should be able to charge for that as an extra unless you bake in the cost of the hardware in the overall monthly cost.

Be extremely cautious in incorporating full DR solutions etc. into “all you can eat” plans - they can be detrimental to your bottom line (and sanity) as some clients will literally never stop eating.

You'll also need to work out how much data storage each client will need - be careful in making assumptions here. One small SMB client I have processes insurance claims and takes tons of pictures and videos of incidents which need to be stored for seven years meaning they have insane amounts of data to be stored and backed up.

“ A MAJOR FACTOR TO A SUCCESSFUL MSP BUSINESS IS BEING CLEAR ON EXACTLY WHAT'S INCLUDED AND WHAT'S NOT INCLUDED IN YOUR PACKAGE(S) ”



CHAPTER 22:

DEFINING BAAS SERVICE

LEVEL AGREEMENTS



Before we begin, please bear in mind that Service Level Agreements (SLAs) can become a legal issue in some situations. This content of this book is not written to adhere to current legalities, and it is highly recommended that you seek the counsel of a legal professional to assess potential risks and considerations when crafting SLAs for any product or service. That said, let's talk about SLAs!

The technical complexities of offering backup services at scale aside, BaaS can certainly create some challenges on the operational side of the business as well. Unless your business is brand new, you've likely taken the time to define some basic SLAs. If you're asking the question "What if I haven't generally defined SLAs yet?" then I suggest you carve out some time, sit down with your legal counsel, the technical leaders on your team, and a trusted client (if you want outside input) and define some general support SLAs for things such as:

- Ticket Severity Definitions
- SLAs for responding to issues at each severity level
- SLAs for resolution of issues at each severity level

That said what about SLAs for backups? There tends to be two schools of thought here.

1. Match your generally defined SLAs - If you've been realistic, and you can honestly say that your generally defined SLA timelines could also apply to customer access to data, then you could use those for simplicity's sake. This would work for those smaller MSPs that work with customers in the SMB space who do not have strict regulatory requirements for the retention of data. However, as you've likely seen throughout the course of this book there is no one-size-fits-all choice that applies to all cases. This is where the second school of thought comes into play

2. Separately Called-Out SLAs for Backup/Recovery in Your MSA - If any of the below situations apply to your customer base, then I highly suggest you define backup/recovery SLAs separately for your BaaS practice:

- You have customers outside of the SMB space
- You have customers that must adhere to industry regulation such as HIPAA, PCI, ITAR, CMMC...etc.

- You work with organizations where timely access to data could be a safety, health, or legal matter, such as a law office, hospital...etc.
- You've defined multiple tiers in your BaaS practice (bronze/silver/gold) with varying degrees and speeds of recoverability at each level

Creating a separate SLA for backup/recovery in this situation allows you to reference that metric throughout the customer relationship in terms of data protection. In this case, it can provide additional peace of mind for your customers if some of them feel your general SLAs aren't aggressive enough, and used correctly in a tiered model, you can leverage it to generate some additional revenue with the idea that you'll cover the costs of the increased oversight and then a healthy margin on top of that.

22.1 AN EXAMPLE: A TIERED SLA MODEL FOR BAAS

With all that in mind, what would a tiered SLA example look like? We've provided a simple three tier SLA example:

Capabilities	Bronze Level	Silver Level	Gold Level
RTO	24 Hours	8 Hours	30 Minutes
RPO	24 Hours	12 Hours	1 Hour
DR Testing Included	No	No	Yees
Number of Recoveries Included per Month	2	5	Unlimited
VM Backups	Yes	Yes	Yes
Office 365 Backup	No	Yes	Yes
Physical Server Backups	No	No	Yes
Endpoint Backups	No	No	Yes
Alerts and Reporting	Yes	Yes	Yes
Backup Self-Service	No	No	Yes
Cost	\$	\$\$	\$\$\$

As you can see each level moving from Bronze to Silver to Gold, includes additional "features" as well as more aggressive RTO/RPOs.

Obviously, as you move to the right the cost increases to the customer as well as the amount of effort needed on the part of your team.

There is no objective calculation for how your costing would be set here. It all depends on the product choice you make, types of backups included, RTOs/RPOs, competitor pricing, your own profit margins, etc. For example, including DR testing into the gold level would require at least quarterly DR testing and man hours from your team, and it should be priced accordingly. Additionally, things like number of included recoveries per month are important. Recoveries can be time-intensive for your team and as such you want to make sure you're protecting yourself from customers that may abuse the privilege. Or, at the very least, you're costing the package in such a way that you have significant wiggle room.

Again, this is simply an example to get you thinking about how you might define your own BaaS SLAs. Whichever route you take, stick with it, and only deviate if there is a REALLY good reason. It's difficult to be rigid as an MSP, but such situations would include a potentially lucrative client with an open checkbook, special needs in service to a client relationship...etc. Even in those situations, remember a one-off is additional time and effort for your team. Try to always stick with your SLAs if possible.

22.2 A NOTE REGARDING DISASTER RECOVERY TESTING

We have mentioned DR testing in relation to BaaS several times already in this section. Naturally, that's because DR testing is ultimately a part of the backup/recovery strategy of any given organization. It's worth a special mention while we're talking SLAs specifically because DR testing is often a sore spot when it comes to defining SLAs for BaaS. Anytime you define SLAs for your BaaS practice there should be a heap of criteria in the same document that defines specifically what you mean when you say, "DR Testing".

A couple of questions you should be asking yourself to get started:

- 1. What are we testing the recovery of?** - Are you testing the recovery of all protected workloads within the organization, or just select critical components?
- 2. Where are we recovering too?** - The same infrastructure? Is there capacity for a meaningful DR test? A remote location? Is there enough bandwidth to accommodate the test along with production traffic?

3. **Where is the finish line?** - What are the criteria for a successful recovery?
4. **Who defines whether a test is successful or not?** - Is this you or one of your employees? Is it someone within your customers org? CEO? CFO? CTO?
5. **How are you proving success?** - Are you generating reports?
Are you capturing screenshots/videos of workloads online?
Are you testing the recovered workload?
6. **How often are DR tests run?** – Monthly? Quarterly? Annually?
7. **When are DR Tests run?** - During the day? After hours?

As you can see, once you start getting into the weeds on some of the above questions the complexities of DR testing start to emerge. It goes without saying that DR testing is a VERY time-consuming process for both you (the MSP) and the customer as well. As such, you need to make sure your SLAs regarding DR testing are realistic, you understand the manpower needed to conduct proper DR testing, and you've priced any SLA/offering that includes DR testing accordingly.

CHAPTER 23:

REGULATORY CONSIDERATIONS



Doing business with customers that are in a regulated industry can be very lucrative when done properly, but it can certainly come with a lot of red tape. That said, becoming fluent in dealing with a given regulatory body can create a niche specialization for your MSP and set you apart from your competitors. This is a topic that has many opposing schools of thought in the service provider space, and your approach will often come down to the larger business decision of whether to work with regulated clients or not. If you haven't taken a stance one way or the other on this it would be advantageous to do so as it helps you better define your target customer and will help define the knowledge and training that your team must possess to service that defined target customer base.

For example, depending on industry, a regulated client may have any of the following regulations in play:



NEED TO RETAIN DATA FOR A DEFINED LENGTH OF TIME
(TYPICALLY 7 TO 10 YEARS DEPENDING ON INDUSTRY AND REGULATION)



STRICTLY CONTROLLED ACCESS TO CONTROL AND MANAGEMENT PLATFORMS FOR BOTH THE MSP AND THE END CUSTOMER



PHYSICAL SEGREGATION OF DATA FROM OTHER ENTITIES



CONTROLS ON THE GEOGRAPHIC STORAGE OF DATA



CONTROLS ON THE NATIONALITY AND BACKGROUND OF INDIVIDUALS INVOLVED IN THE ACCESS, STORAGE, AND USE OF REGULATED DATA



REQUIREMENTS TO ENSURE A MINIMUM LEVEL OF RTO/RPO



ABILITY TO PROVE ADHERENCE TO ANY/ALL CONTROLS AND REGULATIONS AT ANY GIVEN TIME

All the above line items will certainly change the way you conduct business with the customer in question. More specifically to BaaS however, industry regulations will impact how you conduct backup operations, complete restorations, store data, transport data, manage the solution, etc. This is the main reason why it's a good move to define your target customer when setting up your BaaS practice, if you haven't done so for other business planning already.

If you've decided to take the plunge, you'll need to define what regulations you're comfortable dealing with. Once defined, you should have at least two members of your technical staff up to speed and trained on how to conduct backup operations within the defined regulatory guidelines, with processes in place to make sure any new implementation/change is reviewed to ensure compliance. Additionally, your legal counsel should be comfortable or have contacts that can help provide advice on the given controls should the need arise.

To help you define your choice in this area, we've included some links to the most commonly dealt with industry regulations below:

- **HIPAA** - <https://www.hhs.gov/hipaa/index.html>
- **PCI DSS** - <https://www.pcisecuritystandards.org>
- **SOX** - <https://www.sec.gov/spotlight/soxcomp.htm>
- **GDPR** - <https://gdprinfo.eu>
- **ITAR** - https://www.pmddtc.state.gov/ddtc_public

CHAPTER 24:

DEFINING BAAS SPECIFICS

WITHIN YOUR MSA



So thus far, you've done all this work to define your BaaS services and who you're providing said services too. What's the next step? Some will immediately shout "let's go to market!", and while the enthusiasm is well intended, there is one additional task that is often forgotten about. That's updating your Master Services Agreement (MSA). If you're unaware of what an MSA is, I highly recommend you review the resources below as a properly drafted MSA is beneficial to both you and your customers.

- [Do I Need a Master Services Agreement?](#)
- [How to Create Your First Master Services Agreement](#)

That said, if you're aware of MSAs, then you may be wondering why we should define our BaaS SLAs here. Short answer, if it's in your MSA, the terms are automatically accepted by any entity you do business with. Backup services can be deemed important enough that they belong in your core document, as you'll likely be (and should) provide backup services to all customers you work with. An alternative would be to have a BaaS addendum that contains the SLAs and other legal text related to BaaS Services.

Either way, you need to have the legal ruleset (driven by your SLAs definitions above) in play for any major service that you provide to your customers, so this step is critically important. Check with your legal counsel to see which option will work best for your organization and to see if there are any special considerations for your given situation.

CHAPTER 25: PREPARING FOR BAAS OPERATIONS



So far, we've done LOTS of defining and documenting. But how do you get started when it's time to put your BaaS practice into action?

There are several key areas that require preparation ahead of time and not all of them are straightforward. Certain areas, like team training, will be more effective if you get ahead of them sooner rather than later. Your mileage may vary when it comes to exactly what pieces you need to have in place to go to market. That said, let's briefly cover the areas that you'll need to have figured out, at a minimum.

25.1 TRAINING

Team training goes without saying. You can't go to market unless your team knows the product that you're selling and managing inside out. When it comes to team training, owners often immediately think of the technical training first, and while technical training on your BaaS services is wildly important it's not everything. The point here is don't neglect the other parts of your organization. Your sales staff needs to know how to sell the product; your marketing team needs to know how your chosen solution solves business issues for your target customer; and so on.

Solid advice here would be to check with the vendor of your chosen solution. Chances are they'll have several resources already pre-canned that your team can use to get up to speed on the solution. Some will even offer certification paths with margin kickbacks for getting X number of employees certified within your organization. So not only does training your team help you service the solution (and your customers) more effectively, you may get compensation for doing so!

25.2 DEFINING PROCESSES

As already mentioned at least once here, having clearly defined processes is a VERY important step towards running a successful BaaS practice. Running recovery operations usually happens during an emergency and having clearly defined processes in place will help your team AND your customers know what comes next when your back is against the wall. Like anything, processes are best in moderation. Don't over complicate your processes, but always make sure they cover the necessities. Below are a few of the core processes you'll need to have defined to go to market.

**“ DON’T OVER COMPLICATE
YOUR PROCESSES, BUT ALWAYS
MAKE SURE THEY COVER
THE NECESSITIES”**



25.3 THE ONBOARDING PROCESS

Your sales team has sold the solution! Congrats! Now, how do you get the customer up and running with your solution? There are a few key things here you should make sure are in order before covering the technical side of the onboarding process:

- Are all the appropriate documents signed and filed away?
- Do you have a definitive “Start of Service Date”?
 - If so, is your technical team able to accommodate that start date?
- If needed, is the needed hardware available or on order?
- Where will the data be stored? How?
- Any special regulatory requirements for this customer?
 - Are these requirements documented for your staff to see? Have they been made aware of these requirements?
- Is the customer point of contact being involved in the setup process?

- If the customer has their own IT Team, are they involved in the setup process? If so, how?
- What is the “finish line” for the onboarding?
- What is happening with previous backups from a prior solution?

These are all things that you should be thinking about as you define your onboarding process. Make a list, stick to it, and make sure all the checkboxes are ticked each time you onboard a new customer. Additionally, if you’re leveraging RMM software or a ticketing system, consider creating a template out of this process so it’s the same for each new customer that is onboarding with BaaS services.

25.4 DEFINING BACKUP SCHEDULES

This topic has already been discussed at length throughout this section already. That said however, it’s important enough from a process point of view to warrant another mention here. The point here is to ensure that you’re following some sort of process when applying backup schedules to the customers that you onboard to your BaaS services. Mainly, stick to schedules that enable you to hold to your SLAs and only deviate in special situations when paired with specifically defined SLAs for that one-off with a defined and signed addendum.

25.5 PROCESSES FOR HANDLING RESTORATION JOBS

We’ll be talking more at length shortly about the technical and staff procedures for managing multiple backup jobs, but from a procedural perspective what’s the process when a restoration request comes in?

A couple of key questions to consider when determining what this process looks like:

- How is the issue logged?
- What are the criteria for determining severity?
- Who is assigned?
- How are you tracking SLA adherence?
- How are escalations handled?
- How do you prove to a panicked customer that your SLA was met, and all is well?

Answering these questions will help you route and handle incoming restoration requests, even in a difficult situation. Remember, if a customer is asking for a restore, they're often in some sort of emergency depending on the data in question. A defined process will help everyone know what to do when the chips are down.

25.6 AUDITING OPERATIONS

Some will question the meaning of the term “auditing operations”. What does it mean? The term refers to any situation where you or a customer needs to produce data, logs, or other security information for discovery, compliance, or legal reasons. A couple of examples of what this may look like:

- You have a customer that must adhere to ITAR regulations and must prove that their backups are compliant with certain ITAR controls.
- There are legal issues with a customer and a disgruntled former employee
- One of your customers is going through a financial audit.

These are just three of many possible examples, but they're good examples to demonstrate the types of requests that fall within this category. When a request comes in, this process should help you answer the question(s) of who is allowed to request that information, how is that information procured, by whom, and who gets the information at the conclusion of the process? Additionally, how is the process documented and signed off on at the conclusion of the work?

If you need a little assistance in determining the particulars of this process, look to the regulatory bodies you've chosen to work with. For example, if you've chosen to work with ITAR regulated organizations, one of the controls defines how data must be isolated and controlled by only authorized individuals. Knowing this you can tailor your process to make sure that requirement is considered whenever a request of this type comes up.

25.7 THE OFF-BOARDING PROCESS

Finally, onto our last process worth defining, and that is the off-boarding process. While we'd like to think no one will ever want to stop consuming services, the fact is, customer churn is part of the game. At some point you will have a customer cancel services and when that happens,

what's the process for getting them off your backup platform?

This process really boils down to a few key points:

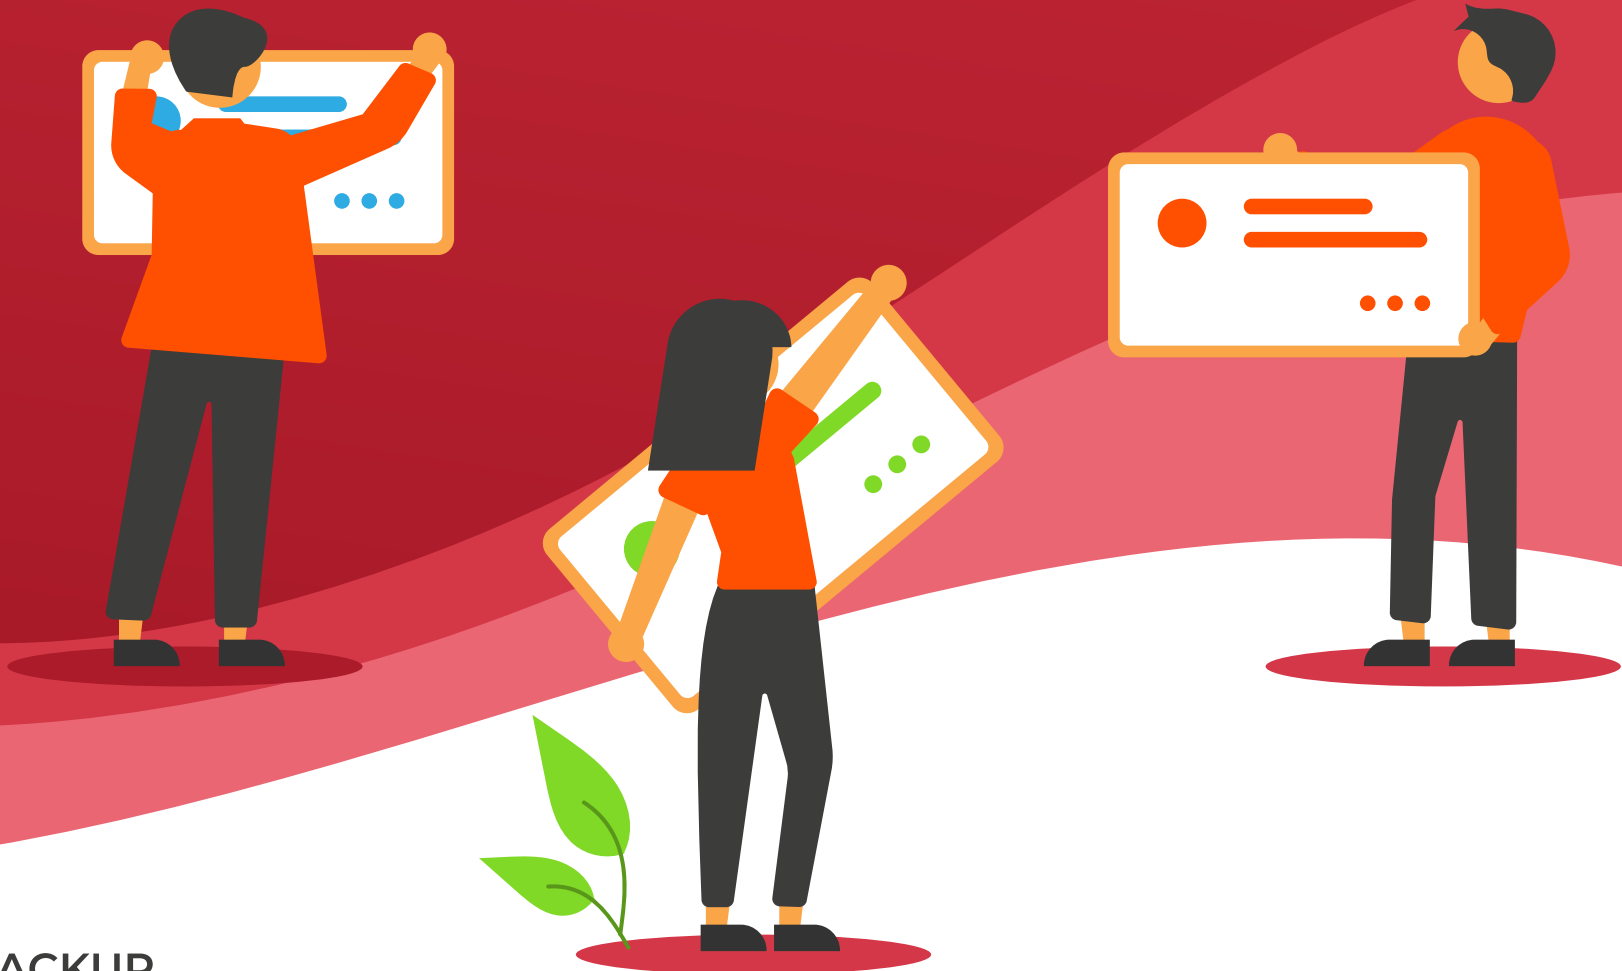
- Is the customer going to retain backups themselves?
 - If so, how? Are they going to use the same backup application and is there a migration path?
- Are you retaining the backups for the customer for a length of time? How Long?
- If the customer doesn't want you holding their data, how are you getting the data back to them?
- Who is responsible for cleanup and removal of the software?
- Is this off-boarding work billable? Is the customer aware of this?

Realistically, each off-boarding process is going to be somewhat different in practice. However, by having a standard process ahead of time, your team and the customer will know what to expect when it's time to part ways. Although admittedly, it's not a very ambitious thing to plan for, failing to have amicable off-boarding can result in resentment and at worst negative reviews and feedback left online. Conversely, excellent off-boarding can be a powerful motivator for attracting returning customers.

CHAPTER 26:

BAAS MSP TECHNICAL

OPERATIONS



While much of what we've talked about thus far has been very procedural and business oriented, there very much is a technical component in play when it comes to providing BaaS service to your customer-base. Sure, the basics discussed in earlier sections of this book still apply, but there are some differences as well. These differences mainly consist of management and deployment at-scale across multiple disparate locations. This section will address those concerns, starting with deployment.

26.1 DEPLOYING BACKUP OPERATIONS AT SCALE

To maintain adequate service levels, you need to develop a process to efficiently deploy backup systems to your clients. You have your choice of several approaches.

MANUAL ROLLOUT

If you have sufficient staff availability, you can deploy backups to clients interactively on an as-needed basis. This works well in a controlled-growth situation where you are onboarding new clients infrequently. If you send technicians to client sites during initial acquisition, then an in-person

software installation can add to a “white glove” experience. If you only use remote tools, then you can achieve a similar effect by installing through a digital meeting application with screen sharing capability.

An addition to the customer experience benefit, a manual rollout also allows for maximum control over the deployment. You can customize the backup application's configuration and respond to the nuances of the client's environment in real time. With properly trained and attentive staff, this technique reduces the error rate to its minimum.

You gain these benefits at the expense of speed and consistency. Clicking through every Windows Installer screen and deciding on options for every installation at every client requires time from a technician that they could spend elsewhere. The friendly phrase “responding to nuances” has a more ominous “making exceptions” interpretation. An exception means that the installation differs from other installations. You can address special cases, but they represent a vulnerability. The installing technician might fail to create proper documents, the documents might go missing, or a future support technician may not check for them. Any of these conditions could cause problems for the client that reflect poorly on your organization.

AUTOMATED ROLLOUT

When manual rollout of a backup application will not suffice, either due to lack of employee resources or a desire to maximize consistency, turn to automated software deployment. “Automation” covers numerous operations:

- **RMM:** You can use a Remote Monitoring and Management tool with software deployment capabilities.
- **Application-based:** Your backup application might offer its own method for automated deployment.
- **Script and policy deployments:** Administrators have needed to solve the large-scale software deployment problem much longer than application vendors have tried to help. You can use built-in or third-party scripting solutions and operating system features such as Windows’ Group Policy to roll out software. Most mature applications with a significant installed base already have published answers that you can utilize or modify to your needs.
- **Combination approaches:** You can mix and match the preceding options in this list until you arrive at an optimal solution.
For instance, you may use your RMM to download software

at the client site and create a policy that deploys the MSI file following a PowerShell script that selects the proper Active Directory Organizational Units.

Automated solutions do have a few notable drawbacks. Primarily, you forfeit the “white glove” experience of a hands-on manual process. However, you will have plenty of alternative ways to physically demonstrate attention to your customers, and many don’t care anyway. Also, your staff can lose full comprehension of the product deployment process. The person(s) that crafted your scripts will eventually leave your organization, and the tools and products that you use and offer to clients will evolve. As always, good documentation provides the best answer to this problem, but it cannot eliminate it.

CLIENT-CONTROLLED ROLLOUT

Beyond the decision of “manual” or “automated” installation, you will also encounter the “responsibility” question. Does your organization handle the deployment, or do you leave it to clients? As with so many things in technology, you have no single best answer.

If you maintain complete control of software deployment, then you also accept nearly all the responsibility. While that statement can seem somewhat foreboding, it usually works out in your favor. You have the most experience with the process and can directly step in if something goes wrong. You will also often have a better relationship with the software manufacturer than your client if the problems prove more difficult than your staff can handle. If a special configuration at a client causes the backup application deployment to fail, then you can take ownership of that as their service provider or you can place the deployment on hold until the client resolves the problem. Make sure that your agreement with the client clearly covers such situations.

At the other end, the client can take charge of installing the backup application on their systems. At first glance, that seems to offload most of the responsibility to them. However, you know that they will bring problems to you with the expectation that you will provide solutions. Even so, with the potential for reduction of the administrative burden on your staff and tools, you might want to accept the risk. To achieve a high success rate, you will need:

- A backup application with a simple and well-designed deployment procedure

- Very clear documentation on the application, potentially augmented by your own. Supply screenshots even for the easiest programs
- A written support statement. As a service provider, you have committed to aiding your client, but you must also set appropriate boundaries when it comes to activities that the client engages in outside of your control

Unfortunately, you can find yourself in an even worse situation than the customer trying to do everything. If both of you have administrative access and can perform tasks without the oversight or knowledge of the other, then you have splashed into the murky waters of “shared governance”. Just like the educational and medical arrangements that spawned this term, you and your client have similar, but not precisely aligned goals. Both of you want the system to operate with minimal fuss. Neither of you wants to expend time and labor fixing problems, especially those that you did not cause and could have avoided. Both of you have the administrative ability to make major problems and hide the evidence. Both of you have incentive to assign such problems to the other entity. Of course, as a service provider, you want to make a name for yourself as reliable. That does not mean that all your employees will always behave with the highest ideals.

You can mitigate the negative impacts of client-driven activity and shared governance by establishing “time and materials” and “best effort” policies for any situation in which you do not have complete control. To some degree, you can sidestep the question of blame by stating in advance that your policies apply regardless of fault. You can also use those policies as negative incentive by offering guarantees on work where your organization maintains full control.

KEEPING DEPLOYMENT PRACTICES IN CHECK

As much as we would like to create a single policy and stick to it in every situation, that rarely happens. Understand the specific needs of your customers and adapt your relationship accordingly. Some clients need a great deal of attention; others only want an emergency contact. No client will suffer with a provider that operates against their wishes while alternative providers exist. You already know that a successful business must adapt to unforeseen circumstances. Absorb the material in this section, decide your default course of action, and prepare for the times when you must take another path. Comprehensible and plentiful documentation gives you the best defense against deployment-related challenges.

26.2 MONITORING AND REPORTING ON BACKUP AT SCALE

So, you’ve covered the deployment process for your clients’ backup software. Congratulations on completing the first step! However, your real responsibility lies in maintaining the ongoing health of your client’s systems. That means that when something does not work as intended, you must find out quickly. Given the nature of the service provider industry, you likely already knew this. Now we need to plan how to accomplish it.

MONITORING WITH A CENTRALIZED SYSTEM

Hopefully, you’ve selected a backup application that provides a multi-tenant centralized management console. If the software can take on the bulk of the rote detail work, that frees up staff time and energy for operations that the computers cannot handle on their own.

Not all programs with a central console will offer all the above features. As with any other software, you will need to make decisions that give you the greatest possible outcome. To that end, you may need to forgo some items on your wish list. However, you won’t necessarily need to discard them entirely.

LOOK FOR THESE CAPABILITIES IN A CENTRALIZED BACKUP MONITORING SYSTEM:



ALERTING



SERVICE TICKET CREATION



SIMPLE PROBLEM RESPONSE

(E.G., AUTOMATICALLY
RESTARTING FAILED SERVICES)



SCHEDULED, MANUAL AND CUSTOMIZABLE REPORTING

WORKING AROUND MISSING CENTRAL CONSOLE SOLUTIONS

You likely have chosen a backup program that does not satisfy everything that you want from a data protection tool. You might have wound up with an application that does not have a centralized component at all. To make up for any software shortcomings, think creatively.

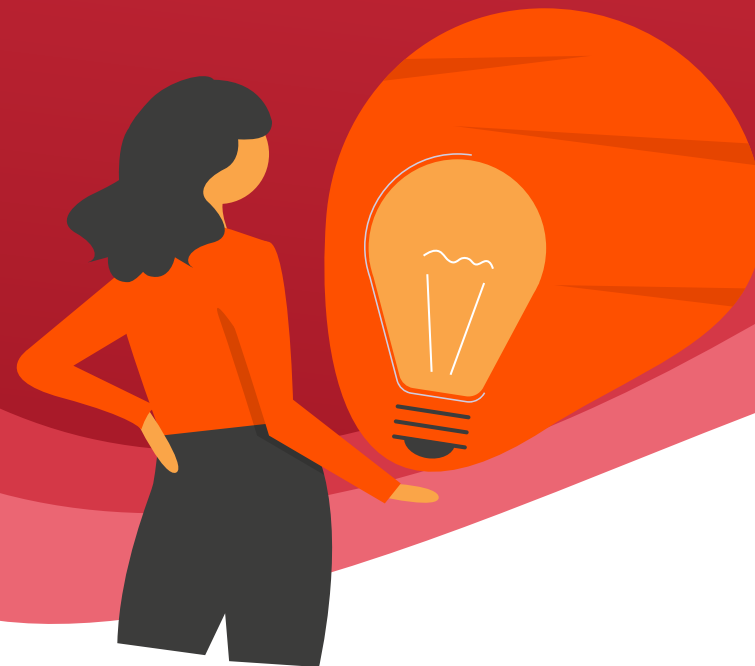
The rapid growth of computer literacy and home automation has introduced people outside the fields of software development and systems administration into the formerly arcane realm of do-it-yourself computing. We have seen a correlated rise in professional computing automation with the widespread adoption of tools such as Python and PowerShell. Administrators and power users have demanded more control planes for their software, and software development companies have responded. Application Programming Interfaces (APIs) have become more common and accessible than ever before. Your backup software probably has some sort of API. If it doesn't natively provide functionality that you want, maybe you can roll your own.

A good monitoring approach always includes log monitoring anyway, but we'll include it here because you might not have any other way to check on the health of a client's backup. Even if the locally installed tool has its own alerting system, you need an external witness as insurance against failure of that system. Regularly check Windows' event logs, application-specific logs, and other systems such as syslog servers.

If you use an Remote Management and Monitoring (RMM) tool to keep an eye on your clients, then it likely has some way to watch backup programs as well. If the backup tool has its own monitoring system, you can still leverage your RMM as an independent watcher.

In the almost-worst case, you have none of the monitoring systems listed above. In the actual worst case, you have some or all the above but none of them work. For your final fallback and lowest priority option, you must have some way to access the client's site. For most providers, that comes in the form of a Virtual Private Network (VPN) connection or something similar. With that, you can join the client's network as though you were on-premises. If even that fails, then you will need to travel to the site or have someone at the client act as your "eyes and ears" while you guide them over the phone.

**"TO MAKE UP FOR AN
SOFTWARE SHORTCOMINGS,
THINK CREATIVELY"**



KEEPING THE CLIENT IN THE LOOP

The topic of customer access and involvement will appear more prominently in an upcoming section. Specifically in the realm of monitoring, you must arrive at an agreement with the customer on their level of involvement. Do they want to receive alerts and reports? Do they want all the information or a summary? Do they want to know when things break but otherwise have no notifications? Get answers to these questions in advance and ensure that you can provide.

DEALING WITH PROBLEMS

Despite your best efforts, something will eventually go wrong. If you have performed due diligence in setting up a monitoring solution, you will know soon afterward. What will you do about it? The customer expects and deserves a calm, prepared response. With planning and documentation, you can give it to them.

Use the capabilities of your tools and define a clear workflow.

As an example:

- On failure, the client's backup system will send an alarm to an e-mail alias at your organization.

- Your e-mail system has a handler that will route e-mails to that alias according to a set of rules. It may handle clients differently according to their service tier. It might direct messages to the help desk e-mail between 8 AM and 5 PM Monday through Friday but route to the on-call emergency line at other times.
- Some human will receive and handle the message within the time frame established by your agreement with the client.
- That person will know how to handle the ticket, either personally or by knowing the appropriate next step in the escalation chain. They will maintain ownership of the ticket until receiving confirmation of hand-off from the next person in the chain.
- Anyone that handles the ticket at any stage will have access to relevant client information written as though the reader has never worked with that client before. For instance, if the owner of the client company has made it clear that he would prefer his building burn down than have someone wake him up in the middle of the night, your responding technician should know that before trying to call him about a failing backup job.

- Anyone that acquires the ticket will know about the applicable service tier. As an example, no one should feel pressure to provide after-hours support to a client that agreed to a business-hours-only plan.
- If desired or necessary, establish a plan in which a technician will notify others even if the technician will handle the problem to completion. An internal sales representative or a client's CIO may want post-mortem reports but not initial notification.

In most organizations, the worst situations arise when something completely unexpected occurs. You cannot possibly plan for every eventuality, so don't try to. Instead, you can set up a clear information-sharing system to prevent anyone from sitting alone with a frustrating problem without anyone else knowing about the problem. Make certain that your clients understand the limits of their agreement and empower your employees to call upon appropriate resources when they reach their boundaries of their ability.

**"YOU CANNOT POSSIBLY PLAN
FOR EVERY EVENTUALITY,
SO DON'T TRY TO"**



26.3 BACKUP AND MAINTENANCE OPERATIONS AT SCALE

You will likely focus the bulk of your efforts on your clients' backup and maintenance operations. After all the lessons learned from deployment and monitoring, you already have a solid foundation to build upon. Combine that with the normal routines that you would follow for a single location, and you can create a formula that works across multiple clients.

Most critically, you must use automation. Otherwise, you will not have an economical way to manage more than a few clients. Fortunately, all backup applications at least offer job scheduling. Ideally, the one you choose also cleans up its metadata and removes data that passes defined retention period thresholds. The best programs include some way to automatically test the integrity of online backup data. Your staff will need to manually perform any of these tasks if the software cannot.

The need for verifying backups opens an opportunity for you. While the software might check online data without intervention, it needs help with offline backups. You can offer a premium service or tier that includes regular manual checks conducted by your technicians. If you need them to handle any tasks that your backup program cannot, then they can selectively add this chore.

Use the same tools to facilitate these operations as discussed in the previous sections: central consoles, RMM, and commodity remote access tools. You can use scripts and APIs to mold backup and maintenance tasks to fit your operations.

CONSIDERATIONS FOR MANAGING MULTIPLE BACKUPS

While backup and maintenance activities at scale work much like small systems, resource control presents different challenges. When you control the entire infrastructure for a single organization, you can easily track storage and networking trends and participate in infrastructure expansions. As a service provider, you will almost certainly have clients that make drastic changes without warning. You must prepare to accommodate them.

If client activities all occur within their premises or cloud accounts, then you have less to worry about. You will need to help them with resource exhaustion problems, of course, but you can handle each client individually. If clients transmit data to facilities or cloud accounts that you own, then you face a risk that they will impact each other.

A few of the things you can do to mitigate the threat of widespread capacity problems:

- Implement quota systems
- Maintain strict monitoring systems that alert earlier than they would for non-shared resources. As an example, you want to know if your storage reaches 85% capacity, or your ingress or egress network frequently sustains over 80% bandwidth utilization.

Work with a hardware or cloud provider that can ship or provision additional resources quickly.

- Explore cloud-based options for emergency overflow.
- If you operate datacenters that receive client backup data, apply redundancy liberally and maintain plenty of slack capacity.
- Use analytics tools and techniques to plot usage trends and calculate predictions.

In short, if your clients rely on you for storage space and network speeds, ensure that you can satisfy their demands.

26.4 PROVIDING BACKUP ACCESS TO CUSTOMERS

Ideally, all your clients would make backups that they never needed to access. Reality won't allow that. They will need to recover data, usually for relatively small events such as lost e-mails or accidentally overwritten files. You could have your staff handle all incidents, but that could overwhelm them. Instead, you can offer your clients the ability to access their data unaided.

Controlling access to data depends on where it resides. You have one approach for backups that remain at the customer's site and a different methodology to restrict what they store on your systems.

CONTROLLING CUSTOMER ACCESS TO ON-PREMISES BACKUP DATA

You have only one serious concern regarding clients accessing data that resides on their local systems: the "shared governance" concern that we discussed in the deployment section. If the client maintains full control over all their systems, then you have nothing to do here.

In fact, you should avoid getting involved as much as possible to avoid disagreements. Offer documentation. If the customer needs guidance,

provide it by voice and screen-sharing tools, never in any fashion that the client cannot oversee.

If your organization maintains sole administrative access, then you will also have more responsibility for restore operations. However, you can likely grant non-administrative ability to read backup contents. Check documentation for the backup application that you use at the client's site as well as the related storage system. Continue reading the next sections, as all access conditions have similarities.

CONTROLLING CUSTOMER ACCESS TO HOSTED BACKUP DATA

If you host your clients' backup data within your datacenter or cloud account, then you have two concerns: allowing clients to access their own data and not allowing anyone to view anything that belongs to someone else. Take extreme care in your solution architecture. Misconfigured access rules lead directly to data breaches.

If the backup application that you use for your clients has a central console, then it may also have a built-in client portal. Leverage that. Watch the vendor's notifications and public warning sites for

discovered vulnerabilities. Keep the application and its host systems current with all security updates.

BEST PRACTICES FOR BACKUP DATA ACCESS

Many rules apply universally, regardless of the access method or the distance between clients and their data. Adhere to the following:

- You must carefully and clearly document methods and policies. Clients must demonstrate a clear understanding during their onboarding process.
- Document and notify all clients of any changes to methods or policies well in advance.
- Discourage clients from using a single access method for all employees. Require them to use individual accounts. Where possible and valuable, use group-based access control.
- Work with clients to select appropriate personnel and access levels. Consider arrangement such as granting full access for technical staff and read-only powers for managers.

- Do not substitute proper reporting with read access to data. Any kind of access increases the attack surface and represents a potential attack vector.

You will need to find a balance for working with your most demanding customers. The only best practice that can apply to all clients is: “know the client”. One demanding customer may fill up your help request queues if you don’t give them enough power. Another may consume an inordinate amount of your time if you give them too much. As with all such things, clear documentation and communication will protect you better than anything else.

26.5 RETURNING BACKUP DATA TO CUSTOMERS

You will sometimes need to return data that you hold to a client in bulk. This usually only happens for two reasons: the client has decided to part ways with your organization, or the client has suffered a catastrophe. Both problems constitute a high priority and high duress situation for you and the client. You must handle them quickly and professionally.

MASS DATA TRANSPORT

Unless you and your client have uncommonly high-speed Internet connections, the time necessary to transmit all their data over the wire may surpass the limits of patience and deadlines. Your only viable alternative to digital transmission is physical transmission. If your client’s site is close to your datacenter, then you might have a secure way for them to come retrieve their data. In most cases, you’ll have to take responsibility for shipping.

If you must rely on third parties, you introduce another element into the responsibility matrix. Common carriers offer secure shipping options, which will suffice for most organizations. Specialty carriers exist that offer solutions particular to IT.

Establish policies in advance. Some considerations:

- What encryption scheme will you use for in-transit data? Will you share a symmetrical key with your client? Will you require clients to provide you with a public key and rely on them to properly secure and protect their private key?

- Will you delete local copies as soon as you ship or retain a copy of customer data until you receive confirmation of receipt and verification of integrity?
- What verification technique will you use?
- What degree of responsibility will you take for data lost or damaged in transit? What remediation will you offer?
- Will you use special media for transport or standard media with protective packaging?
- Instead of, or in addition to, physical transport, will you offer digital transmission to another provider (such as an Azure storage account)?
- Will you expect the client to assume the role of expert of their industry and provide information on data-safety standards and laws, or will you assume the role of data protection expert and perform the research on their behalf?

Related practices and policies must appear in your client agreements. If you modify the policies, notify clients immediately and require acknowledgment. Failures and miscommunications in data return present serious risk to your company's reputation.

OFF-BOARDING CLIENTS

You built your company to succeed and your processes to match. However, even the best providers lose clients occasionally. You must plan for that just as much as you plan to properly store their information.

Start with an identification of client-related data. Their backup data comes to mind as an obvious concern, but you also keep information about the client – organizational structure, contact and shipping information, billing data, and more. Research and understand any applicable laws regarding retention of that data.

Focus on trustworthiness as your highest imperative during termination of a customer relationship. No matter how things ended, remember that your business must survive independently of any customer, which means that it must always operate in a way that does not deter potential clients.

Devise and document your own general off-boarding plan first. You do not need to make it complex, but it is important to have. It should be at minimum a four-step process:

1. Delete customer backup data on request or upon confirmation of receipt of returned data.
2. Delete customer payment data after settling of final invoice.
3. Delete all customer-related data that would not apply in a legal challenge after 60 days of final invoice payment.
4. Retain legally sensitive information for duration set by regulation or in accordance with legal advice.

Such a plan will cover most situations. Prepare for some clients to insist on different arrangements. You will need to decide whether to accept or reject their conditions, but in all cases, do nothing outside of a written agreement. All activities must fall within the defined parameters of written, mutually agreed upon procedures.

Of course, you have an obligation to carry out your intent. Employ scripts and automated tools to help with the cleanup. Preferably, make these viewable for clients to help assure them that you adhere to the rules. You can never truly prove that you have eradicated all information, but it boosts your credibility when you have something more than an ad hoc approach. Ensure that these tools have meaningful safeguards to protect against accidental or malicious acts.

26.6 TECHNICAL OPERATIONS REVIEW

As a service provider, technical operations define your core business. Even non-technical clients can recognize your level of preparedness and professionalism. Frequently review your processes for any ways to smooth your practices. Sometimes you need more automation; sometimes you need less. Remain vigilant, seek out ways to improve, and carefully heed client feedback. You may find that your best customer experience results from activities that you did not anticipate.

CHAPTER 27:

NEVER THE END



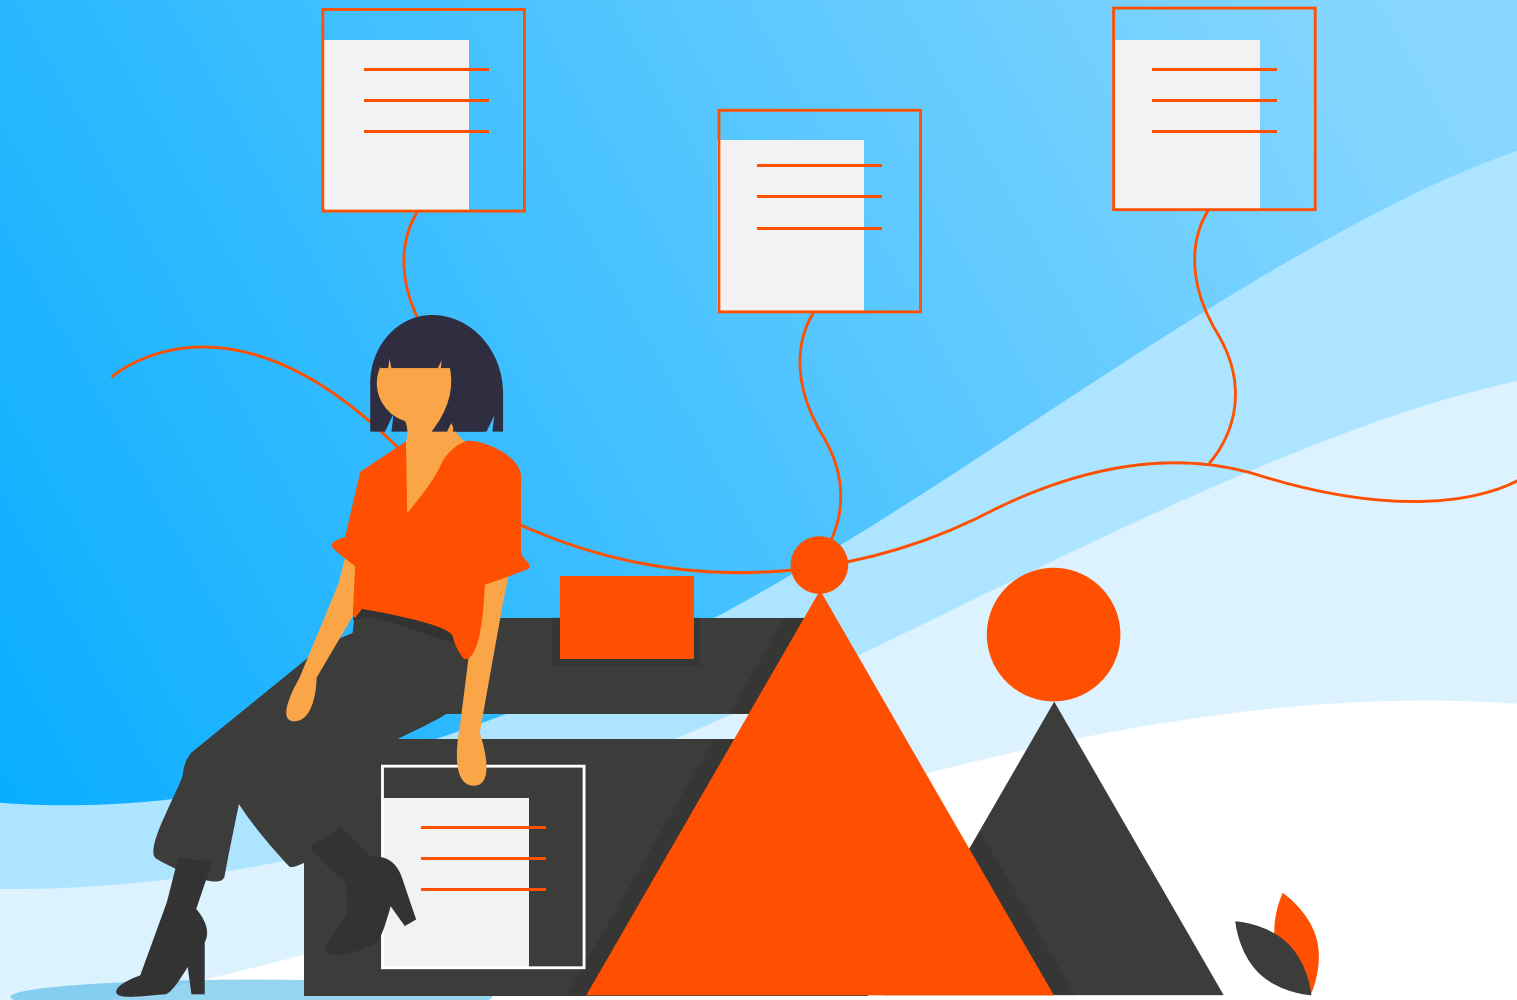
If you've read everything up until now, you'll know that the core message of this book is that backup is an on-going journey, as are backup services. Many people in IT continue to downplay the importance of backup & recovery as a basic service – perhaps due to a perception of being boring and unexciting but hopefully we've demonstrated that those labels are certainly not the case. Delving into the intricates can be fascinating and when the time comes for your backups to save the day, there will be a lot of excitement around your work.

On top of that, this final part will have shown that the people side of the equation is even more important when it comes to delivering backup services as an MSP. Not only is there a customer service aspect to the delivery, but you also must convince prospects of the importance of backup & DR, why you're the one to deliver it, and then do the needed discovery internally for that customer, just like you learned in previous parts of the Backup Bible.

As stated at the beginning of this part, few things in your MSP are as important as backup & DR services for your customers. Here you've learned about all the pieces you'll need in order to sell and deliver those services effectively and consistently, but that's not the end by any means. Customers need change new backup technologies emerge, data storage

in the cloud continues to innovate, etc. Most importantly of all, the risk to your customer's data (and potentially their livelihoods), will continue to grow. By implementing a well thought-out and scalable backup & DR plan, you've certainly helped mitigate that risk, but vigilant monitoring, communication, and an ever-evolving strategy will always be necessary to keep your customer's data safe, and their trust for you in place. Accordingly, this book will be updated regularly to cover significant changes in the backup & DR landscape so that when you feel it's time to update your plan, this guide will be here to help you be successful in protecting you and your customer's data.

GLOSSARY



Where possible, the terms in this book follow industry standards.

In cases of ambiguity or where no generally agreed-upon definition exists, this glossary explains words as used in this text.

Application-consistent backup: Backup tools can notify compatible applications of a pending job. The applications can respond by completing in-memory operations and flushing data to disk. If restored, the application will be in a completely usable state. The backup will capture data for incompatible programs in a crash-consistent state.

Backup: A "cold" instance of duplicated data.

Cold backup: Backup data that requires meaningful effort to retrieve. Tapes in a safe deposit box qualify as "cold" backup.

Cold data: Infrequently accessed data. Compare to "hot" and "warm" data. A "warmer" "temperature" indicates a higher usage rate.

Cold site: Cold sites vary from little more than an empty building to a standby facility that has sufficient hardware to operate as a primary site but no data. Cold sites require more effort to bring online than warm sites, but less effort to maintain.

Business continuity: The ability of an organization to continue performing its desired activities and functions throughout an emergency.

Cluster/clustering: A group of computers that provide rapid, automatic failover protection to resources, usually in an active/passive configuration. Microsoft SQL databases and Hyper-V virtual machines can be protected with a cluster.

Crash-consistent* backup: A backup captured at a precise point in time. This type of backup has its name because if completely restored, the recovered system will behave as though it had crashed at the moment the backup was taken. A crash-consistent backup does not protect any running processes, nor does it give them any opportunity to save active data. However, it captures all files exactly as they were at that moment.

Delta: The difference between a file or a block as it was at the previous backup and its current state. Most vendors use the term "delta" in some way to designate technology that operates at or below the file level.

* Not everyone agrees on the definitions of "crash-consistent" and "application-consistent". Some vendors have introduced their own labels. Ensure that you understand how any given vendor uses these terms when you study their products and talk to their representatives.

Differential backup: An abbreviated backup that only captures data that changed since the most recent full backup. Usually operates at the file level.

Disaster recovery: The process of returning to full functionality after an emergency.

Fault-tolerant: A system or subsystem that can withstand at least one failure without service interruption.

Full backup: A complete, independent duplication of data that you can use to recover all data without any dependency on data.

High availability: An expectation and measurement for the amount of time that a service can perform its primary function within a given period, expressed as a percentage. A system that is online for 357 days and 4 hours within a 365-day period is *available* for about 98% of the time. If its promised availability was 99%, then it did not meet its high availability requirement. Many technologies exist to increase availability.

Hot backup: Data in a backup archive that you can access just by opening your backup software. To support a "hot" status, the data must exist on some sort of always-on storage repository.

Hot data: Frequently accessed data. Compare to "hot" and "warm" data. A "warmer" "temperature" indicates a higher usage rate.

Hyper-converged: Also seen as “hyper-converged infrastructure” (HCI). A high-density computing platform with no physical equipment dedicated to compute or storage. Similar to traditional virtualization but adds distributed storage as a function of standard commodity server equipment. HCI hosts almost always belong to a group of hosts that split and share responsibilities for virtual machines, including their compute, storage, and networking.

Hot site: A geographically distant location from your primary site that your staff can access and that receives regular, automated data updates. A "hot" site can stand in for the main site in a short period of time with little effort.

Incremental backup: An abbreviated backup that only captures data that changed since the most recent backup of any kind. Usually operates at the file level.

Initial seed: Creating a starting replica on physical media to be transported to the replica site. This prevents the need to send an extremely large data set across a relatively slow inter-site link.

Media: Storage for backups. Used as a catch-all word whether you save to disks, tapes, optical discs, or anything else.

Recovery point objective (RPO): The maximum acceptable time span between the latest backup and a data loss event.

Recovery time objective (RTO): The desired maximum amount of time before a system returns to a defined usable state.

Replica: A "warm" or "hot" instance of duplicated data.

Restore: The act of retrieving data from a backup. Restoration can return data to a live system or to a test system. Most tools allow you to choose between complete and partial restores.

Rotation: Re-using backup media, usually by overwriting older backups. Some backup software has intricate rotation options.

Split-brain: A condition in which two separate running systems contain different, incompatible copies of information.

Warm backup: Backup data that you can access easily but not instantly. As an example, data on a USB hard drive in an on-premises storage closet or a cloud vault.

Warm data: Data accessed too often to be considered "cold" but not often enough to be considered "hot". These distinctions are completely arbitrary and used mostly for prioritization and optimization efforts. A "warmer" "temperature" indicates a higher usage rate. No precise boundaries exist between them. In backup and replica contexts, we mostly concern ourselves with the rate of change.

Warm site: Warm sites have the same essential rules as a hot site, but they do not receive data directly. A warm site needs to receive a recent copy of production data from warm or cold backup and restore it before it can stand in for the primary site.

APPENDIX — LIST TEMPLATES AND CHECKLISTS

Use the following checklists and templates in this section
as starting points for your own business continuity documentation.
Adjust according to your own specific business requirements.



CHECKLIST FOR MEETINGS IN THE PLANNING PHASE

Disaster recovery touches every part of your enterprise, so it needs input and participation from all corners. You will need to organize and conduct several meetings.

- ☐ Business case meeting – explain the need for a disaster recovery plan
 - ☐ Invite executives and department heads
 - ☐ Explain the consequences of inaction
 - ☐ Secure a willingness to commit to funding
 - ☐ Secure a commitment of personnel time
 - ☐ Identify key stakeholders
- ☐ Scope meeting – begin discovering the scope of your disaster recovery plan
 - ☐ Invite key stakeholders
 - ☐ Explain the desired data points
 - ☐ Provide a checklist or questionnaire
 - ☐ Allot time for completion with a deadline
- ☐ Scope follow-up meeting
 - ☐ Invite key stakeholders or designated representatives
 - ☐ Collate data points
 - ☐ Discover overlaps
 - ☐ Schedule additional time and meetings as necessary

- ☐ Technology presentation meeting
 - ☐ Invite executives and stakeholders
 - ☐ Show options and possibilities
 - ☐ Begin pricing discussions
- ☐ Additional planning meetings – further topics
 - ☐ Roles and responsibilities
 - ☐ Implementation planning
 - ☐ Procedure design
- ☐ Status report meetings
 - ☐ Invite key stakeholders
 - ☐ Report progress on deployment and implementation

Remember not to make these discussions solely about data and technology. You also need to prepare for losses of physical assets and personnel. More business continuity procedures will deal with employees and activities than with disks and tapes. Use these meetings as opportunities to explore topics such as work-from-home policies during a disaster.

RISK IDENTIFICATION

Use this list as a starting point for identifying the risks that apply to your organization.

Once you have identified risks, create an assessment of their potential impacts.

- ☐ Data theft
- ☐ Physical theft
- ☐ Malicious digital attacks (ransomware, viruses, etc.)
- ☐ Rogue insiders
- ☐ Social instability
- ☐ Power failures
- ☐ Arson
- ☐ Sabotage
- ☐ Natural disaster
- ☐ Departure of critical staff

KEY STAKEHOLDERS

Create a list of the individuals (by name or by title) that have an interest in any phase of the business continuity/disaster recovery planning process.

Name / Title	Organizational Role	Department	In Planning Phase?	In Deployment Phase?	In Recovery Process?	In Routine Update Processes?	Contract Information

DATA PROTECTION QUESTIONNAIRE

As a technology professional, you know how to protect data. As company experts, your key stakeholders know what data to protect. You need to combine your knowledge into an actionable plan. You can use questionnaires as a way to gather the necessary information. Use the following questionnaire as a starting point to design your own.

- Which employees have the most knowledge of the computer and data systems that your department relies upon?
 - Where does your critical data reside?
 - Network servers
 - Centralized desktops
 - Employee laptops
 - Cloud providers
- What computer and software systems do you require for full operational functionality?
 - Line-of-business applications
 - ◆ Workstation-based
 - ◆ Server-based
 - ◆ Cloud-based
 - Commodity applications
 - Internally developed applications
- What technologies do you require for full operational productivity?
 - Communications devices
 - Specialty hardware
 - Commodity hardware (desktops, laptops, printers)
- What personnel does your department rely on? What if those individuals are not available?
- Do you currently protect your data? Technology? Assets?
- Who supports your technologies?

THE BACKUP BIBLE

- What technologies do you require for minimal functionality?
- How long can you operate at minimal functionality?
- What technologies do you require for acceptable functionality?
- How long can you operate at acceptable functionality?
- Estimate the departmental cost of an hour of complete downtime.
- In the event of a system failure, how far back could your department recreate data?
- How long must you keep copies of your various data points?
- Prioritize your technology

Make certain that everyone with relevant knowledge receives a copy of the questionnaire.

Even if an individual cannot provide an answer to every question, gather as much information as possible.

INFORMATION TECHNOLOGY DEPARTMENT CHECKLIST

The IT department will have its own special data protection questionnaire including, but not limited to, the following:

- List of key stakeholders
- List of application experts
- List of data experts
- Organizational contacts
- Support contacts
- Mission-critical categorized systems and data
- Important categorized systems and data
- Low priority categorized systems and data
- Available technology solutions, capabilities, and costs

It will also be your responsibility to process all of the information as it comes from the departments. Create a usable plan, share it with the stakeholders, and acquire funding for the project.

DESIRED PROTECTION TECHNOLOGIES

To fully investigate and test the available backup and disaster recovery solutions, you need to know the sorts of technologies that you need to protect. Use this list as a starting point for your own list of cases that your tools must accommodate.

- Windows Server and Windows desktop
- UNIX/Linux systems
- Database servers
- Mail servers
- Virtual machines
- Cloud-based resources
- Physical hardware configurations
- Active Directory
- Log-based SQL recovery
- Mail servers
- Backup synchronization for multi-tier systems
- Cluster nodes

THE BACKUP BIBLE

SAMPLE TABLE OF BACKUP APPLICATION TEST RESULTS

This table comes from Phase One at the end of the "[Exploring Disaster Recovery Technologies](#)" chapter. Use it as a template for vetting the backup applications that you test.

Product	Version	Within Budget	Physical Systems	Hyper-V Virtual Machines	Backup to Cloud	Active Support
Product A	7.0	✓	✓	X	X	X
Product B: Advanced Edition	2.1	X				
Product B: Starter Edition	2.1	✓	✓		✓	✓
Product C	4.9	✓	✓	✓	✓	✓

BACKUP SYSTEM DEPLOYMENT CHECKLIST

Refer to the "[Deploying Backup](#)" chapter for an in-context discussion of the items in this list.

1. Acquire software and hardware
2. Place backup hardware
3. Install backup software into your test environment
4. If your software uses agents, push to test systems
5. Fully test your hardware and software. Verify all functionality, even the portions that you might not expect to use.
6. Document the test environment setup and installation process. Take special note of anything that did not go as planned and how you remedied the problem.
7. Install backup software into your production environment
8. If your software uses agents, push to a representative sampling of systems
9. Test expected functionality on the sample systems
10. Document the deployment into production, including fixes and workarounds
11. Continue deploying agents, if necessary, until you have covered your entire environment
12. Capture your initial full backup to store offline and transport it to the offsite location
13. Train staff on usage and have them practice
14. Document backup and restoration processes

SAMPLE BACKUP DOCUMENTATION FOR A SMALL ORGANIZATION

You can use this sample as a starting point for building your own documentation. It was designed with smaller organizations in mind, particularly those without a dedicated IT department and not enough staff to ensure that someone that has experience with the backup system will be available in a disaster recovery scenario. Everything in this sample is fictional; it does not use any real-world devices or programs.

BACKUP AND RESTORE PROCEDURE FOR ABC, LLC

Last update: *January 2021*

This document outlines the organization's configuration and restore procedures.

Hardware Information

- Manufacturer:
- Device type:
- Model name:
- Support telephone:
- Support e-mail:
- Sales/representative name:
- Sales/representative contact:
- Website:
- Site login information in company key vault
- Original device warranty expiration date:

THE BACKUP BIBLE

Connection steps:

1. Choose a physical Windows Server system that will operate the backup
2. Download the latest device driver from:
3. Install the driver (include screenshots)
4. Use the included cable to plug the device into a USB slot

Software Information

- Software manufacturer:
- Program name:
- Support telephone:
- Support e-mail:
- Sales/representative name:
- Sales/representative contact:
- Website:
- Site login information in company key vault
- License expiration date:

Installation steps:

1. Download latest version from:
2. Select a physical Windows Server system to operate the software
3. Run downloadfile.exe
4. Click Next on the first page



Add Screenshot here

THE BACKUP BIBLE

5. Install to the default location

 Add Screenshot here

6. Click Finish to install

 Add Screenshot here

Notes on problems encountered during installation and workarounds:

.....

Steps to connect to cloud account:

1. Open the program from Start menu with the icon

 Add Screenshot here

2. Enter the login information for a local administrator account

3. Go to the Cloud Account section

 Add Screenshot here

4. Enter our login information from the company key vault

Data restoration steps:

1. Open the program from the Start menu with the icon

 Add Screenshot here

2. Enter the login information for a local administrator account

3. Go to the Restore section

4. Navigate through the tree to find the data that you want to restore,
or click the check box at the top left to restore everything

 Add Screenshot here

5. Choose where to restore following the screenshot below

 Add Screenshot here

6. Click the Restore button

THE BACKUP BIBLE

Backup configuration steps:

1. Open the program from the Start menu with the icon



Add Screenshot here

2. Enter the login information for a local administrator account

3. Go to the Backup section

4. Click the checkboxes next to C: and D: to select everything



Add Screenshot here

5. Click the Schedule button

6. Set Full backup to Saturdays, 9 PM



Add Screenshot here

7. Set Incremental backup to Sunday-Friday, 9 PM



Add Screenshot here

8. Set the primary destination to the hardware device and the secondary destination to our cloud account, set up during installation



Add Screenshot here

As you look through the template, take note of all the places that might stop a non-technical user.

What is a Windows Server? Where do they get one? What hardware devices do we need in functioning order to perform disaster recovery? Who can they call for procedural assistance?

The chapters in [part three](#) address these items in more detail.

While creating your documentation, use screenshots liberally. Even technical people find them helpful when working with unfamiliar software.

SAMPLE BACKUP DOCUMENTATION FOR A LARGER ORGANIZATION

Larger organizations typically have their own templates for project documentation. Follow those guidelines first if they exist. If not, you can start with the above template for smaller organizations. You need to preserve the same fundamental information.

Large organizations typically have differences from smaller companies in their documentation:

- A greater likelihood of available technical staff during a disaster recovery operation
- Single-process recovery documentation usually does not suffice. Categorize documentation along department, application, environment, or any other delineations that make sense
- Ability to safely keep copies in multiple on-premises locations

You can use this template to help you build or augment your documentation.

Everything in this sample is fictional; it does not use any real-world devices or programs.

BACKUP AND RESTORE PROCEDURE FOR ZYX CORP

Last update: *August 1, 2020*

This document outlines the organization's foundational configuration and restore procedures.

Hardware Information

- Manufacturer:
- Device type:
- Model name:
- Number deployed:
- Support telephone:
- Support e-mail:

THE BACKUP BIBLE


- Sales/representative name:
- Sales/representative contact:
- Website:
- Site login information in company key vault
- Original device warranty expiration date:

Departments/applications/environments/etc. that use this device

Location	Purpose

Connection steps:

1. Choose a physical Windows Server system that will control the device
2. Download the latest device driver from:
3. Install the driver

 Add Screenshot here
4. Use the included cable to plug the device into a USB slot

Backup Software Information

- Software manufacturer:
- Program name:
- Support telephone:
- Support e-mail:
- Sales/representative name:

THE BACKUP BIBLE

- Sales/representative contact:
- Website:
- Site login information in company key vault
- License expiration date:

Hardware targets used by this program:

.....

Network targets used by this program:

.....

Cloud targets used by this program:

.....

Installation steps:

1. Download latest version from:
2. Select a physical Windows Server system to operate the software
3. Run `downloadfile.exe`
4. Click **Next** on the first page



Add Screenshot here

5. Install to the default location



Add Screenshot here

6. Click **Finish** to install



Add Screenshot here

Notes on problems encountered during installation and workarounds:

.....

Steps to connect to device:

1. Open the program from Start menu with the icon

 Add Screenshot here

2. Enter the login information for a local administrator account

3. Go to the **Devices** section

 Add Screenshot here

4. Select the device

Steps to connect to cloud account:

1. Open the program from Start menu with the icon

 Add Screenshot here

2. Enter the login information for a local administrator account

3. Go to the **Cloud Account** section

 Add Screenshot here

4. Enter our login information from the company key vault

General data restoration steps:

1. Open the program from the Start menu with the icon

 Add Screenshot here

2. Enter the login information for a local administrator account

3. Go to the **Restore** section

4. Navigate through the tree to find the data that you want to restore, or click the check box at the top left to restore everything

 Add Screenshot here

THE BACKUP BIBLE

5. Choose where to restore following the screenshot below

 Add Screenshot here

6. Click the **Restore** button

Backup configuration steps:

1. Open the program from the Start menu with the icon

 Add Screenshot here

2. Enter the login information for a local administrator account

3. Go to the **Backup** section

4. Click the checkboxes next to C: and D: to select everything

 Add Screenshot here

5. Click the **Schedule** button

6. Set **Full backup** to **Saturdays, 9 PM**

 Add Screenshot here

7. Set **Incremental backup** to **Sunday-Friday, 9 PM**

 Add Screenshot here

8. Set the primary destination to the hardware device and the secondary destination to our cloud account, set up during installation

 Add Screenshot here

Application 1 Restore Information

Use these steps to restore Application 1

1. Install Microsoft SQL Server 2017 using our baseline configuration
2. Have a new Windows Server 2019 host ready
3. Follow the general data restoration steps for backup program 1

THE BACKUP BIBLE

4. Pick the Application 1 database as a source



Add Screenshot here

5. In the restore target dialog box, choose the database server from step 1



Add Screenshot here

6. Restore the database

7. Restart the general data restoration steps

8. Pick the name of the original Application 1 application server as a source



Add Screenshot here

9. Pick the name of the new server from step 2 as the target



Add Screenshot here

10. Restore the application

11. Start the "Application 1" service on the restored server

12. Test functionality and connectivity

You will need to duplicate portions of this template to cover as many applications as necessary.

You might need to create site-specific

SITE DESCRIPTION

Create a form similar to the following to record details about your different sites.

Site name:

Site address

.....
.....

Site phone numbers

.....
.....

Site Contacts

Name	Role	Work Phone	Cell Phone

Primary site function:

Additional site functions

.....
.....

Business continuity functions

.....
.....

THE BACKUP BIBLE

CUSTOMER BACKUP REQUIREMENTS TEMPLATE

Use this template as a starting point to create your own customized customer backup requirements list.

This checklist can be used to help ensure that you capture the requirements that a client has on a case-by-case basis.

Client requirement	Yes / No	Notes
Backup of Office 365		
Backup of Google Workspace		
Backup of (Insert Workload Name Here)		
Backup of (Insert Workload Name Here)		
Backup of on-premises servers		OS versions?
Backup of on-premises application servers		Applications?
Backup of “in the office” client computers		Numbers? OSs?
Backup of “work from home” client computers		
Backup of BYOD client computers		
Cloud backup storage only		
On-premises cached backup storage		
Frequency of backups		Might differ per application / service

ABOUT THE AUTHOR



ERIC SIRON

Eric is a four-time awardee of the Microsoft Most Valuable Professional award in Cloud and Datacenter Management.

He has worked in IT since 1998, designing, deploying, and maintaining server, desktop, network, storage, and backup systems.

Eric has provided all levels of support for businesses ranging from single user through enterprises with thousands of seats and developed a particular affinity for small organizations.

He has achieved numerous Microsoft certifications and was a Microsoft Certified Trainer for four years. Eric is also a seasoned technology blogger and has amassed a significant following through his top-class work on the Altaro DOJO.

PART 4: PROVIDING BACKUP SERVICES TO MSP CUSTOMERS

WAS WRITTEN BY ANDY SYREWICZE AND PAUL SCHNACKENBURG



ANDY SYREWICZE

Andy is a 15+ year IT pro specializing in Virtualization, Storage, Cloud, and Infrastructure. By day he works as a Technical Evangelist and Content Creator for Hornetsecurity/Altaro developing and by night he shares his IT knowledge online or over a cold beer. He holds the Microsoft MVP award in Cloud and Datacenter Management, and one of few who is also a VMware vExpert.



PAUL SCHNACKENBURG

Paul started in IT when DOS and 286 processors were the cutting edge. He runs Expert IT Solutions, a small business IT consultancy on the Sunshine Coast, Australia. He also works as an IT teacher at a Microsoft IT Academy. Paul is a well-respected technology author and active in the community, writing in-depth technical articles, focused on Hyper-V, System Center, private and hybrid cloud and Office 365 and Azure public cloud technologies. He holds MCSE, MCSA, MCT certifications.

ALTARO BACKUP



Solutions for Companies and Organizations



Hyper-V & VMware Backup & Replication

Award-winning virtual machine (VM) backup and replication solution for Hyper-V and VMware environments

[Learn more](#)

Microsoft 365/Office 365 Backup

Backup solution for Microsoft 365 mailboxes and files stored in OneDrive and SharePoint, with unlimited storage

[Learn more](#)

Windows Server Backup

Physical to virtual (P to V) backup solution to back up physical Windows servers and restore them to a virtual environment

[Learn more](#)

Solutions for Managed Service Providers (MSPs)



Hyper-V & VMware Backup & Replication

Monthly subscription program enabling MSPs to offer Hyper-V, VMware and physical Windows server backup services

[Learn more](#)

Microsoft 365/Office 365 Backup

Monthly subscription program enabling MSPs to back up customers' Microsoft 365 mailboxes and OneDrive/SharePoint files

[Learn more](#)

EndPoint Backup

Monthly subscription program enabling MSPs to provide backup services for on-premise and roaming Windows desktop and laptop

[Learn more](#)

50,000+ Customers

10,000+ Partners

2,000+ MSPs

121 Countries

Start your free 30-day trial today



Continue your learning on the Altaro DOJO, our dedicated learning platform for IT professionals



Backup & DR

[Go to Section](#)



MSP

[Go to Section](#)



Hyper-V

[Go to Section](#)



VMware

[Go to Section](#)

