

The Essentials of SharePoint Disaster Recovery



Sean P. McDonough
Chief SharePoint Evangelist
Idera



About

me

aka



"How I got SharePoint chocolate
in my DR peanut butter"

About me

My background with disaster recovery (DR)

- started before I ever touched SharePoint
- began in the financial services & insurance industry

My background with SharePoint

- began in 2004 with SharePoint Portal Server 2003
- I switch between IT Pro and Developer hats

DR and SharePoint

- co-authored two SharePoint DR books
- regularly speak, blog, and "work" on DR topics

About this talk: why?

Most DR presentations I've seen (and delivered myself) focus on "how to" technical concerns ...

- How to implement backups
- How to establish high-availability

Not enough has been done to discuss the choices and processes that go into DR planning

- aka, the "non-gearhead" stuff



The prerequisites

This is a 100-level talk, so I don't assume much:

- you don't know much about DR (other than "it's a good idea for my organization")
- you are interested in the end-to-end DR process and more than just strictly technical concerns.

Don't take notes unless you really want to

- <http://SharePointInterface.com>

my blog





The Agenda

- Discuss the "big picture"
- Analyze the DR process
- Explore how SharePoint and DR come together

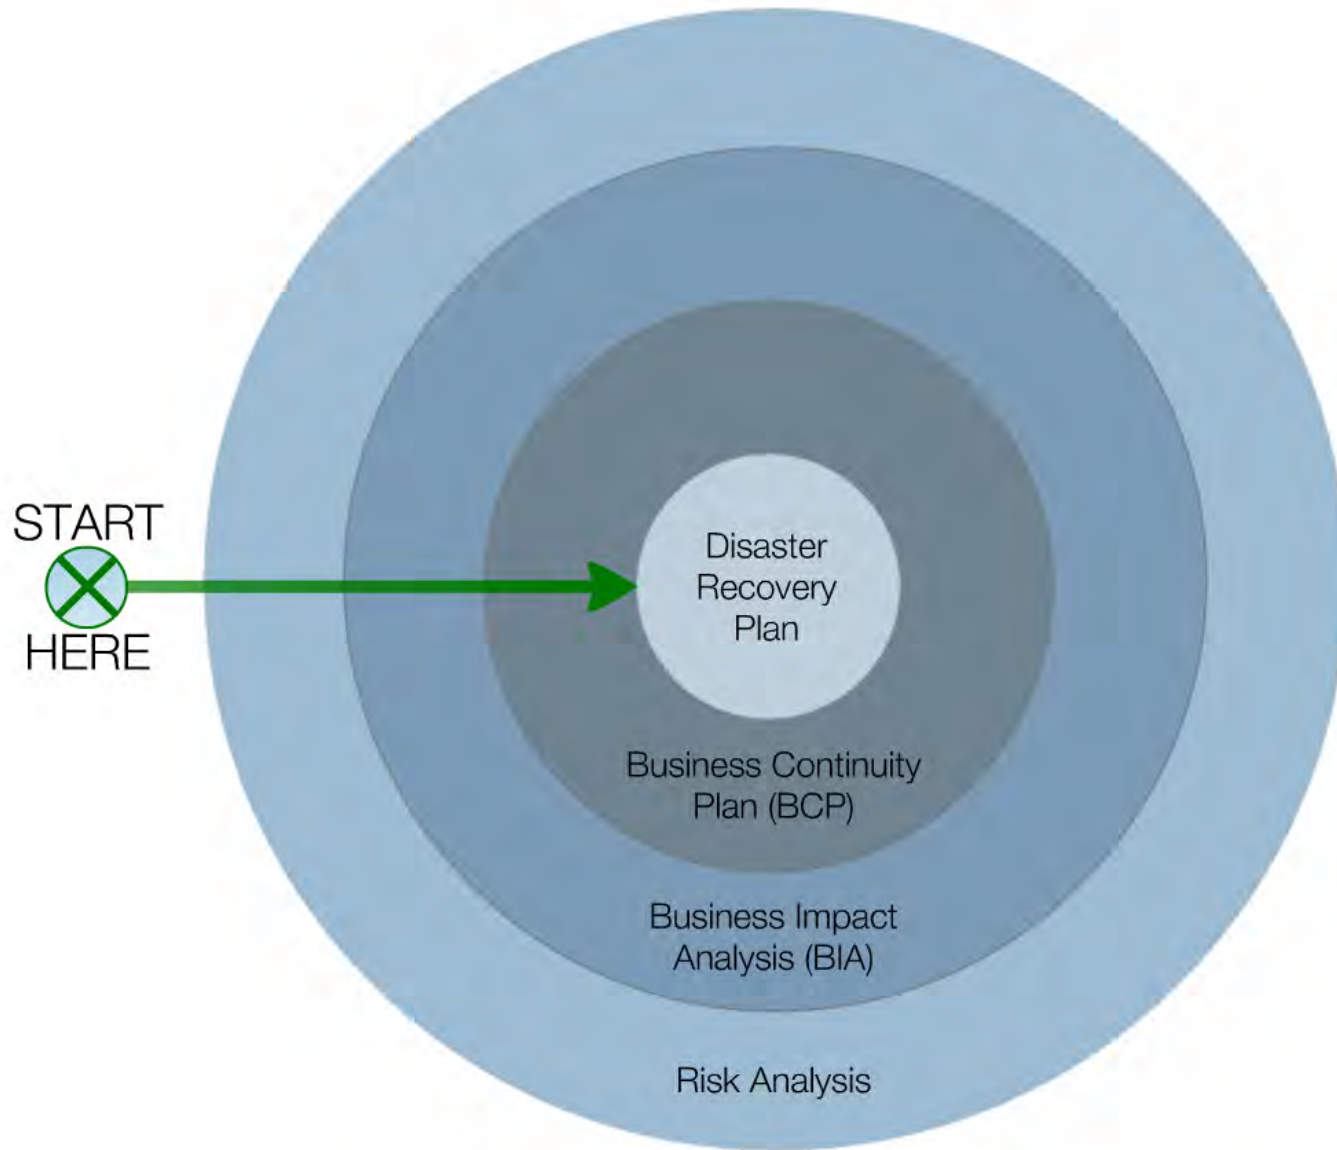


"The Big Picture"

"The Big Picture"



- many layers
- you're probably going to cry as you peel them back



There's a lot that should happen before you ever get to an actual DR plan



Risk Analysis

Risk Analysis

Identifies and quantifies the probable threats to normal business operations and activity.

What could go wrong?

- Primary data center is flooded
- Your network is cyberattacked
- The bulk of employees fall ill
- Power is lost to your location (who kicked the cord?)

Quantify it

- What is the realistic probability of the event?
- If the event occurs, how severe would the impact be?
- Probability x Severity = Overall Risk



BIA

BIA

A business impact analysis maps risks to business processes and systems that would be affected if something were to go wrong.

What comes out of the BIA?

- A document or matrix that maps individual risks to one or more business processes and systems that would be affected
- An estimate of what each interrupted process or disabled system might cost the organization, expressed in dollars per hour (\$/hr)
- Prioritization of processes and systems to protect
- Acceptable loss and downtime windows



BCP

BCP

A business continuity plan addresses the findings of a BIA and defines processes to mitigate and/or minimize interruptions to normal business operations.

What does a BCP cover?

- Manual procedures and work-arounds to keep business moving in the absence of supporting systems
- Key information and logistical plan to address unavailable facilities, equipment, and personnel
- Communications plans
- Disaster recovery links



DR Plan

DR Plan

Disaster recovery plans document requirements and steps for restoring systems to agreed-upon levels of functionality.

What can be found in a plan?

- An overview of what the plan addresses and what it doesn't address (usually reported)
- Recovery protocols: hardware, software, facilities, personnel, etc.
- References to equipment/information/systems/terms
- Procedures for recovery
- Measurable success criteria for recovery

These are a big waste from the point of planning and will be used

Risk Analysis

Identifies and quantifies the probable threats to normal business operations and activity

What could go wrong?

- Primary data center is flooded
- Your network is cyberattacked
- The bulk of employees fall ill
- Power is lost to your location (who kicked the cord?)

Quantify it

- What is the realistic probability of the event?
- If the event occurs, how severe would the impact be?
- Probability x Severity = Overall Risk

Disaster Recovery Journal


<http://www.drj.com/>

Good online reference for disaster recovery articles, whitepapers, AND other resources.

BIA

A business impact analysis maps risks to business processes and systems that would be affected if something were to go wrong

What comes out of the BIA?

- A document or matrix that maps individual risks to one or more business processes and systems that would be affected
- An estimate of what each interrupted process or downed system might cost the organization, oftentimes in dollars per hour (\$/hr)
- Prioritization of processes and systems to protect
- Acceptable loss and downtime windows 

on of processes and systems to p
le loss and downtime windows ✨



These are a key outputs from this phase of planning and will be used extensively in subsequent phases.

BCP

A business continuity plan addresses the findings of a BIA and defines processes to mitigate and/or minimize interruptions to normal business operations

What does a BCP cover?

- Manual procedures and work-arounds to keep business moving in the absence of supporting systems
- Key information and logistical plans to address unavailable facilities, equipment, and personnel
- Communications plans
- Disaster recovery plans

DR Plan

(Disaster) recovery plans document requirements and steps for restoring systems to agreed-upon levels of functionality

What can be found in a plan?

- An overview of what the plan addresses and what it doesn't address (equally important!)
- Recovery prerequisites (hardware, software, facilities, personnel, etc)
- References to dependent information/systems/items
- Procedures for recovery
- Measurable success criteria for recovery



Risk Analysis

Risk Analysis
 Identify and quantify the threats to normal business operations and activity.

What could go wrong?

- Primary data center is flooded
- Your network is sub-optimal
- The bulk of employees fall ill
- Power is lost to you (location A and B are not the same)

Quantify it

- What is the relative probability of the event?
- If the event occurs, how severe would the impact be?
- Probability x Severity = Overall Risk



BIA

BIA
 A business impact analysis maps risk to business processes and systems that would be affected if something went wrong.

What comes out of the BIA?

- A document or table that maps individual risks to one or more business processes and systems that would be affected
- An estimate of what each interrupted process or system might cost the organization, often times in dollars per hour (DPIH)
- Prioritization of processes and systems to protect
- Acceptable loss and downtime windows

These are a key output from the phase of planning and will be used extensively in subsequent phases.



BCP

BCP
 A business continuity plan addresses the findings of a BIA and defines processes to mitigate and/or minimize interruptions to normal business operations.

What does a BCP cover?

- Manual procedures and workarounds to keep business moving in the absence of supporting systems
- Key information and logical plans to address unavailable facilities, equipment, and personnel
- Communication plans
- Disaster recovery plans



DR Plan

DR Plan
 Disaster recovery plans document procedures and steps for restoring systems to avoid or limit issues of functionality.

What can be found in a plan?

- An overview of what the plan addresses and what is covered (usually prioritized)
- Recovery procedures (hardware, software, facilities, personnel, etc)
- References to dependent information/systems/data
- Prerequisites for recovery
- Mass e-mail or voice criteria for recovery

More abstract



More concrete

More strategic



More tactical

More "business-y"




More technical

Disclaimer

There are many approaches to quantifying disaster risks and building contingency plans; I'm presenting only one. Form isn't nearly as important as simply ensuring you have a strategy!

that was the

"big
picture"



The focus going forward
is on ...

the DR  Process

... which is driven by RPO
and RTO requirements

This is a good point to define those acronyms

RPO



Recovery
Point
Objective

RTO



Recovery
Time
Objective

MTD



Maximum
Tolerable
Downtime

RPO



Recovery
Point
Objective

RTO



Recovery
Time
Objective

MTD



Maximum
Tolerable
Downtime

That's all great, but what do they really MEAN?

They define operational windows that guide and inform your selection of technologies and strategies for recovery

RPPO



RPO (Recovery Point Objective)

Monday Jul 4 2011

1:00 AM 2:00 AM 3:00 AM 4:00 AM 5:00 AM 6:00 AM 7:00 AM 8:00 AM 9:00 AM 10:00 AM 11:00 AM 12:00 PM 1:00 PM 2:00 PM 3:00 PM 4:00 PM 5:00 PM 6:00 PM 7:00 PM 8:00 PM 9:00 PM 10:00 PM

- "looks backwards"
- defines maximum acceptable data loss

Monday Jul 4 2011

8:00 AM 9:00 AM 10:00 AM 11:00 AM 12:00 PM 1:00 PM 2:00 PM 3:00 PM

wards"



Your data center just
took a mortar ...

RPO (


RPO (Recovery Point Objective)

Disaster

Monday Jul 4 2011

Tuesday Jul 5 2011

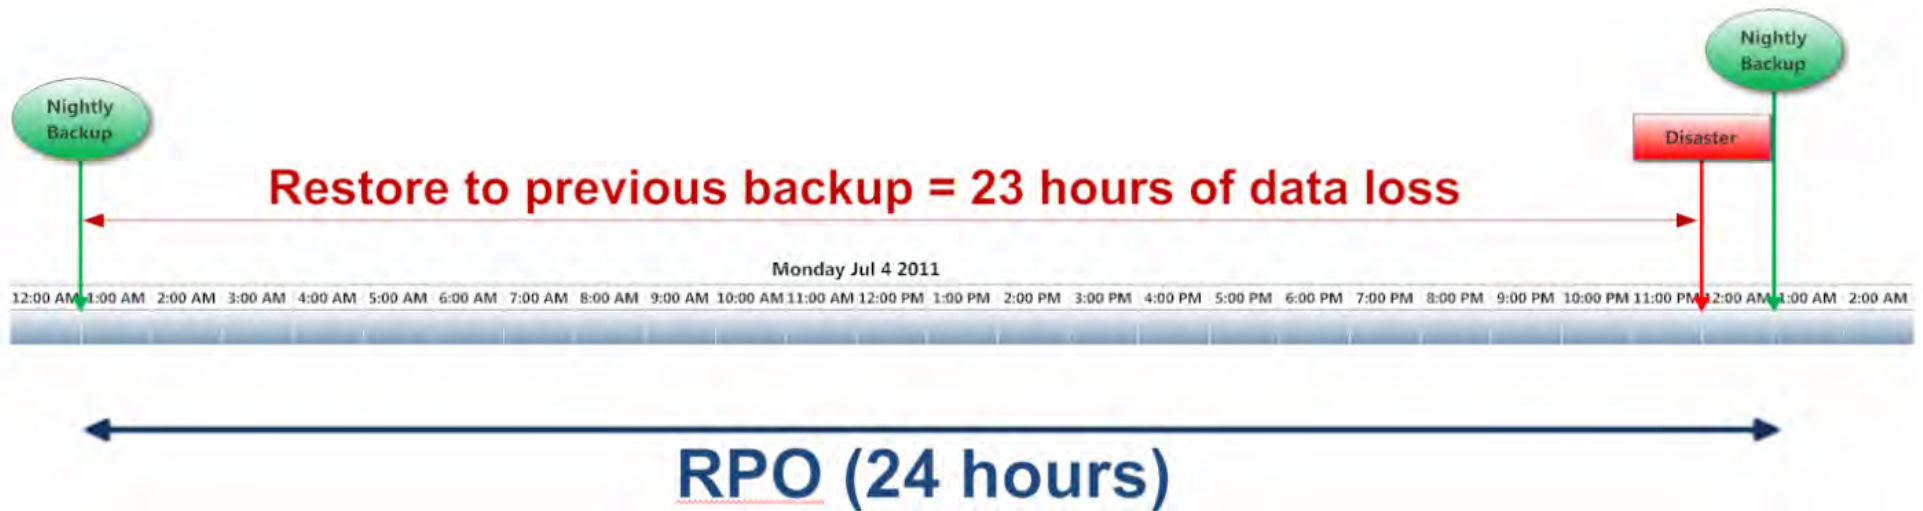
2:00 PM 3:00 PM 4:00 PM 5:00 PM 6:00 PM 7:00 PM 8:00 PM 9:00 PM 10:00 PM 11:00 PM 12:00 AM 1:00 AM 2:00 AM 3:00 AM 4:00 AM 5:00 AM 6:00 AM 7:00 AM 8:00 AM 9:00 AM 10:00 AM 11:00 AM



RPO (Recovery Point Objective)



RPO (Recovery Point Objective)



RTTO



RTO (Recovery Time Objective)



- "looks forward"
- defines how much time you have to get things working again

RTO (Recovery Time Objective)



← RTO (8 hours) →

RTO (Recovery Time Objective)



← RTO (8 hours) →

MTD (Maximum Tolerable Downtime)

- Refers to the maximum time permitted before operations are restored following a disaster
- Different from RTO in that it includes business procedures, human factors, etc. - not just the system restore time (i.e., time to execute a recovery plan)

We'll be focusing on RPO and RTO

The focus going forward
is on ...

the DR  Process

... which is driven by RPO
and RTO requirements

Please allow me a
moment to preach ...





Risk analysis

BIA

BCP



DR Plan

RPO and RTO are determined up here

Implementation takes place down here

Business



Technical

If you're trying to build a DR Plan without business input, you're doing it wrong.

Kind of like ...

DR Plan




Business
Continuity
Strategy



If I haven't beat the horse
to death enough for you ...



<http://sharepointinterface.com/2009/07/08/rpo-and-rto-prerequisites-for-informed-sharepoint-disaster-recovery-planning/>



The DR Process



Assessment



Planning



Maintenance



Implementation

The DR Process





Assessment



Assessment

Building an understanding of

- The SharePoint platform itself
- Your SharePoint environment as it exists today

Accomplished through two "D" words

Discovery

- Logical architecture
- Physical deployment
- Configuration data
- Business data (content)
- Dependencies and interfaces

Before we go too far, we should probably talk about the other "D" word



You're going to have to document your discoveries and SharePoint itself

Believe it or not, there are tools that can help.



Logical Architecture

- Focuses on the SharePoint's software/service components, what they do, and how they relate to one another
- Particular attention is placed on platform elements you use

Commonly documented

- IIS application pools
- SharePoint Web applications
- Service applications (Search, BCS, Managed Metadata, etc.)
- Zones and alternate access mappings
- Web application policies
- Content databases
- Site collections
- My Sites

Commonly documented

- IIS application pools
- SharePoint Web applications
- Service applications (Search, BCS, Managed Metadata, etc.)
- Zones and alternate access mappings
- Web application policies
- Content databases
- Site collections
- My Sites

Goal: show which pieces of SharePoint are in-use, how they interrelate, and how they work together

Think "birds-eye" view of logical farm components - not physical layout/usage

Physical Architecture

- Focuses on SharePoint's implementation across a set of infrastructure components and hardware

Commonly documented

- Physical servers used by SharePoint
- SQL Servers
- Storage area networks (SANs)
- Switches
- Wide area network (WAN) connections
- Firewalls
- Hardware load balancers
- Active Directory domain controllers
- Email relays and gateways

The modern monkeywrench that makes all of this more complicated:



Virtualization



Logical Architecture



Physical Architecture

- 33% of small and mid-size businesses (SMBs) admitted that they do not back up virtual servers as often as physical servers
- 49% back up virtual machines weekly or monthly
- 37% back up virtual machines each day

Source: Acronis Global Disaster Recovery Index 2012
http://acronisinfo.com/?attachment_id=521

Configuration Data

- Focuses on the data and settings that make SharePoint and its constituent components/pieces operate.

Commonly includes

- Farm configuration database
- Non-content service application databases
- Web.config files
- IIS7 configuration files
- Other configuration stores tied to logical architecture items

Commonly includes

- Farm configuration database
- Non-content service application databases
- Web.config files
- IIS7 configuration files
- Other configuration stores tied to logical architecture items

Initially, it is more important to understand where data resides and the form it takes than to document actual settings

Pay close attention to secure configuration data, configuration data that is stored in a tough-to-reach manner, and distributed configuration

Business Data

- This is data that gets created and exists within SharePoint as a result of day-to-day business

If you remember nothing else,
remember this:

Content
databases

=



as in "most important business
data locations to protect"

Dependencies & interfaces

- These are the points where SharePoint touches other line of business systems - including other SharePoint farms.

Some examples

- HR Data consumed through an external list using BCS
- Search that is supplied through a separate services-only SharePoint farm
- A Page Viewer web part that exposes a non-SharePoint Web application using an iframe
- InfoPath forms that pull data from (or write data to) non-SharePoint systems

Some examples

- HR Data consumed through an external list using BCS
- Search that is supplied through a separate services-only SharePoint farm
- A Page Viewer web part that exposes a non-SharePoint Web application using an iframe
- InfoPath forms that pull data from (or write data to) non-SharePoint systems

These are important to identify for purposes of determining what is ultimately included in (and excluded from) your SharePoint DR plan

Documentation tools



Creating SharePoint diagrams

Technical diagrams (SharePoint Server 2010)

<http://technet.microsoft.com/en-us/library/cc263199.aspx>

Visio stencils for IT Pro posters

<http://www.microsoft.com/download/en/details.aspx?displaylang=en&id=11616>

PowerShell farm documentation

Document farm configuration settings (SharePoint Foundation 2010)

<http://technet.microsoft.com/en-us/library/ff645390.aspx>

Document farm configuration settings (SharePoint Server 2010)

<http://technet.microsoft.com/en-us/library/ff645391.aspx>

Documentation Toolkit for SharePoint

<http://www.spdockit.com/>

note: not free



Assessment



Planning



Maintenance



Implementation

The DR Process





Planning



Planning

Planning

Assessment Results



$$+ \begin{matrix} \text{RTO} \\ \text{RPO} \end{matrix} =$$

Recovery Plan



Well, there is a little more to it than just that

- Define scope and granularity
- Define recovery targets
- Define approach (technology)

Scope and Granularity

Common Questions

- Do you treat your SharePoint farm as one big system or as multiple functional pieces?
- What's not in-scope for your plan? Are one or more separate (but dependent systems) included?
- How do you handle regional disasters such as earthquake, flood, or attack? The choice carries data center implications
- Can you (or do you even want to) leverage the cloud?

- 23% of all businesses don't have an offsite backup strategy in place today
- 42% rely on onsite backups to tape or disk and then take those (physically) offsite each day

Source: Acronis Global Disaster Recovery Index 2012
http://acronisinfo.com/?attachment_id=521

Cloud computing



- Is not a DR "magic bullet"
- Simplifies some aspects of DR (availability) but complicates others (RPO/RTO, security)
- Comes in many different shapes, forms, and hybrids (Office 365, ITaS, private cloud, etc.)

Figure out if the cloud will be part of your DR strategy early on!

Answers to these questions help determine how you ultimately define ...

Recovery Targets

Recovery Targets

- Specify what to restore in SharePoint through mapping of business processes to SharePoint functional area
- Prioritized from most critical to least critical (RPO, RTO, \$\$\$)

Simple example

"I need to restore the HR intranet"

May entail building and/or restoring:

- A SharePoint farm (baseline environment)
- Content database housing the HR site collection
- BCS and associated connections to external line-of-business systems housing HR data
- Secure Store service for required BCS credential sets
- InfoPath Services for HR-related forms



Define Approach

What is the appropriate combination of strategies AND technologies to address your recovery targets?

Common approaches with many variations

- Backup and restore
- High availability (HA)



Your choice will likely be guided by a few key considerations:

- RPO
- RTO
- Resources (Cost)

Factors when selecting

| | Backup & Restore | High Availability |
|-----------|--|--|
| RPO | Typically hours | From minutes down to zero |
| RTO | Typically hours | From minutes down to zero |
| Resources | Less expensive | Significantly more expensive |
| Examples | SharePoint native backups SQL Server backups Enterprise backup systems 3rd party SharePoint backups | Windows clustering Replication products SQL Server mirroring Transaction log shipping |



VIRTUALIZATION



CLOUD

What should you protect?

Technical (recovery) targets you select depend on your strategy, but most plans include the following critical (technical) items at a minimum:

Databases

- Content ✱
- SSP/Service App

Solution packages (WSPs)

Documentation

- Farm configuration
- Server configuration
- Accounts & permissions



2 FREQUENTLY ASKED QUESTION

How do I get a list of
everything I should protect?

Core SharePoint targets aka, "the meat and potatoes"

- Content databases
- Central Admin content DB
- Farm Configuration database
- Service Applications / SSPs
- Search




Related targets "important to remember"

- Solution packages
- SharePoint Root
- IIS Configuration
- Certificates
- IIS web root
- GAC
- Registry
- Bits and bytes



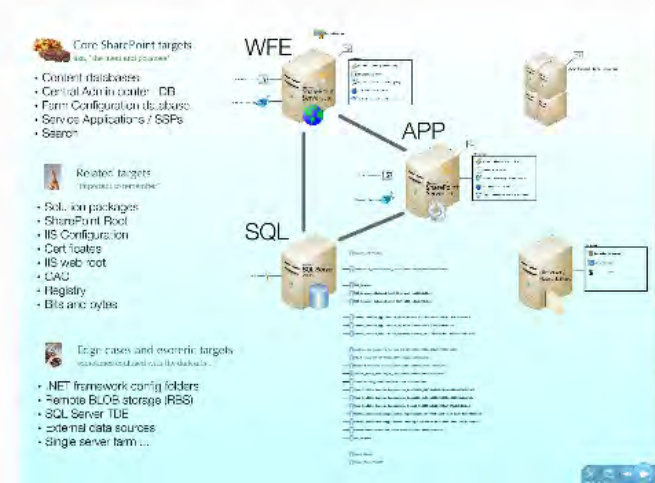
Edge cases and esoteric targets sometimes confused with the dark arts ...

- .NET framework config folders
- Remote BLOB storage (RBS)
- SQL Server TDE
- External data sources
- Single server farm ...



Answer:

There is no definitive "list of everything." No two SharePoint farms (or DR plans) are the same.



You'll have to build your own list based on what you use in your farm

Documentation



Of course you need to document your DR plan!

Documentation of the plan spans this phase and the next phase ...

on plan!

Documentation of the plan spans this phase and the next phase ...

Let's be honest:

documentation spans every phase of the DR process.

There's a focus on it here and in the next phase, though.



Assessment



Planning

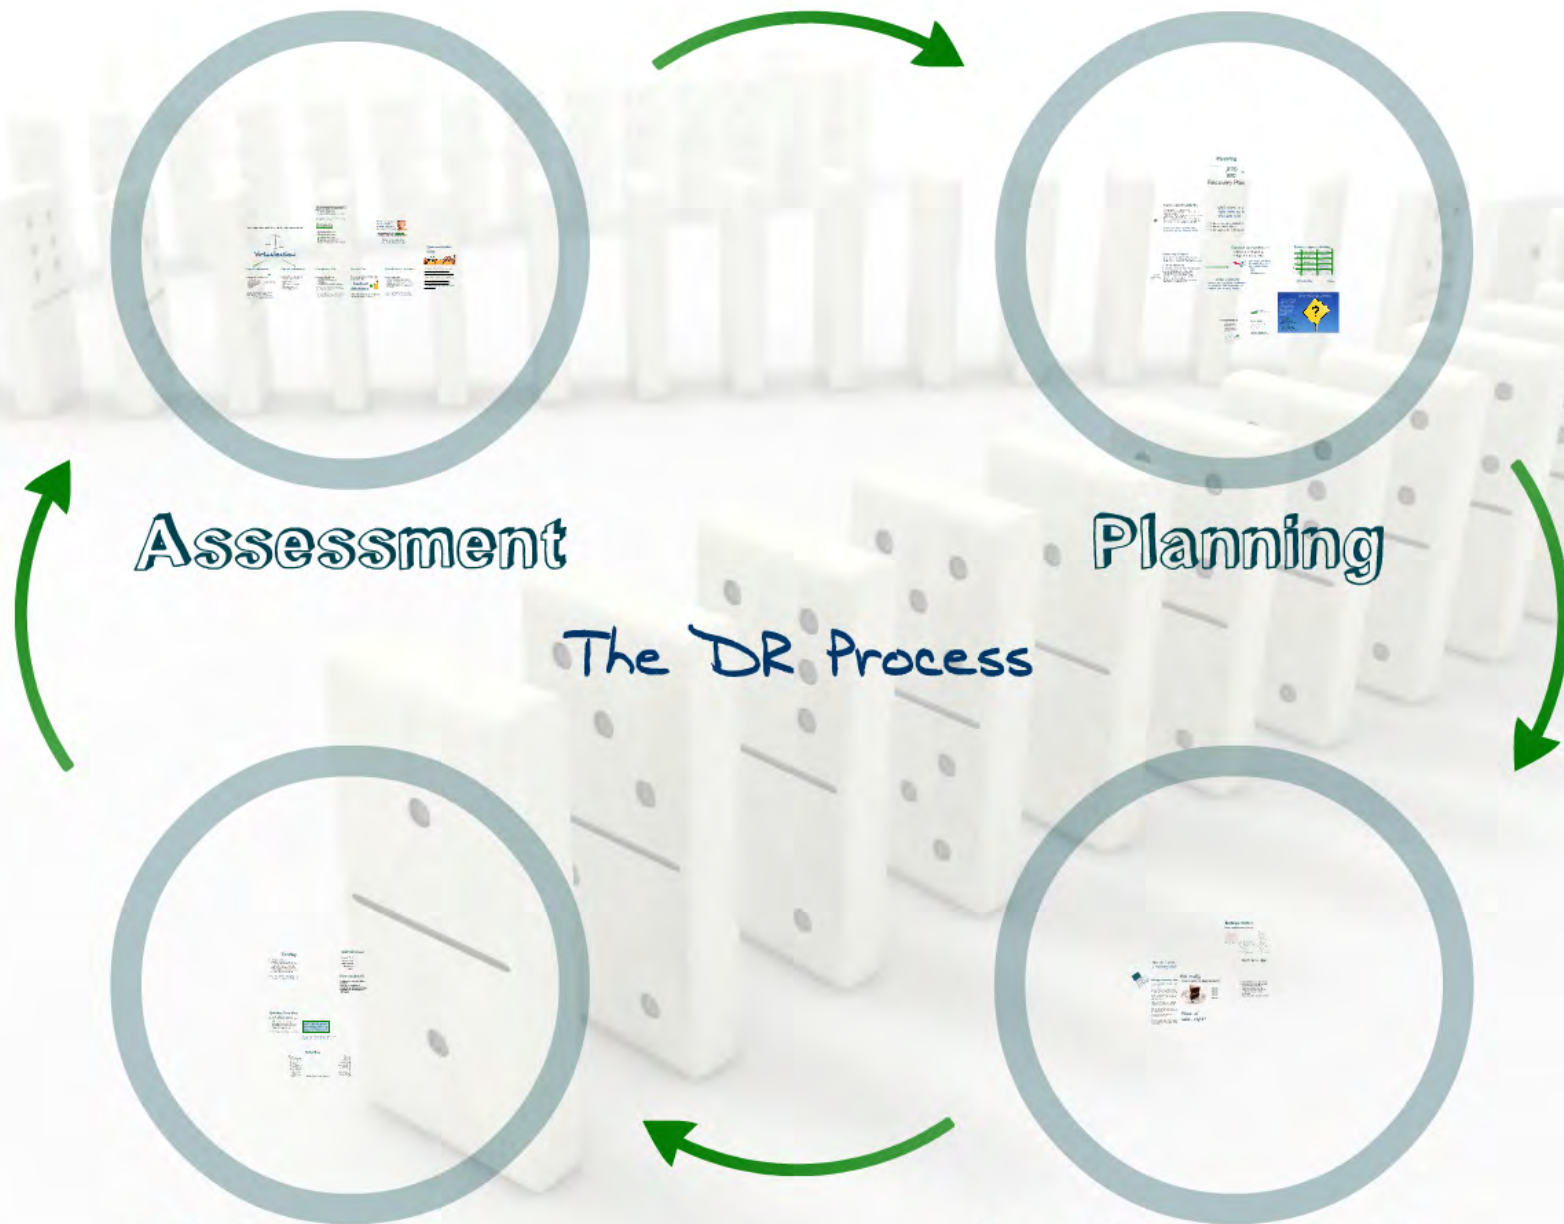


Maintenance



Implementation

The DR Process





Implementation

Implementation

Where the rubber meets the road



Write the plans ...



Assemble resources ...

QUESTION FREQUENTLY ASKED QUESTION

How do I write
A recovery plan?



Start by talking to those in your organization who have some responsibility for business continuity.

Write

CHARACTERISTICS of A good recovery plan

- The plan should address a discrete system or subsystem
- Clearly identifies any assumptions made by the plan - hardware, software, dependent system restores, knowledge to execute, etc.
- Identifies participants and the role they play in the plan - typically by group or company rather than individual
- Outlines steps for recovery in explicit detail - written to a lowest common denominator without requiring specialized knowledge
- For each step, potential misteps are spelled-out along with corrective actions that may be taken (if applicable)
- Specifies objective criteria for recovery confirmation and/or success
- Includes any post-recovery notes directing personnel to further actions or plans that are coupled (implicitly or explicitly) to the current recovery plan



Piece of
cake, right?

Not really.

Recovery plans are living documents



Iterate
Iterate
Iterate
Iterate
....



Assemble resources ...

Some things to consider

Software

- Licenses
- Install media

Physical storage

- Documents/plans/lists/etc.
- Secure or sensitive items

Hardware

- SharePoint Servers
- SQL Servers
- Switches
- Storage/SANs
- Firewalls
- AD controllers/appliances
- DNS servers/appliances
- Load balancers

Facilities (if required)

- Rent/buy space
- Data center build-out
- WAN connectivity
- HVAC
- Fire suppression
- (Voice) communications
- Security
- Backup generators/power

Don't lose sight





Your SharePoint recovery plans should be tying into one or more bigger BCPs

- Communications plan tie-ins
- Criteria for DR plan activation
- Clear documentation of (manual) workarounds for non-restored functionality
- Integration points with other DR plans



Assessment



Planning



Maintenance



Implementation

The DR Process





Maintenance

Maintenance

The part that
many of us
wish would
simply go
away.



What's included?

- Exercises that test and validate DR plans
- Updates to your plans as SharePoint environments change
- Budgeting for the changes that will happen

PASS



FAIL



Testing

How testing can help you

- Identify gaps in plans so that you can address them before a disaster
- Validate that you can actually hit RPO and (especially) RTO targets
- With repetition, you can reduce your RTO (practice makes perfect!)

How testing can help you

- Identify gaps in plans so that you can address them before a disaster
- Validate that you can actually hit RPO and (especially) RTO targets
- With repetition, you can reduce your RTO (practice makes perfect!)

Bottom line: without testing you'll never know if your recovery plans actually work

Updating Your Plan

As your SharePoint environments change, so too must your recovery plans

- RPO and RTO may change
- SharePoint farms grow and evolve
- SharePoint used for new purposes
- Offsite DR facilities change

Your DR plans are living documents ...



Don't leave your
recovery plans to
become "undead"

They don't "go away" because you abandon them;
they just take on an un-life of their own ...



Both time AND money

Time

- Carry out DR tests (personnel, facilities time, business downtime)
- Review, maintain and update DR plans
- Review changes to SharePoint farms
- Audit plans with an eye towards compliance with any regulations



Both time AND

Money

- Salary costs associated with dedicating time to DR activities
- Costs associated with offsite facilities
 - Recurring licensing costs*
- Costs associated with independent auditing of systems and DR plans



and money



Assessment



Planning




Maintenance



Implementation

The DR Process



The dirty secret


Nobody gets this
"right" the first time;
that's why it's a
continuous process



Assessm



An important resource



NIST Special Publication 800-34, Rev. 1, Contingency
Planning Guide for Federal Information Systems
http://www.nist.gov/manuscript-publication-search.cfm?pub_id=905266

Wrap-up

- Remember the order of operations:

Risk Analysis → BIA → BCP → DR Plan

- RPO AND RTO drive MANY of the DR planning decisions you'll make
- No two SharePoint environments are alike; no two DR plans are identical
- Recovery plans are living documents that you'll constantly test AND revise



SEAN P. McDONOUGH

Blog: <http://SharePointInterface.com>

Email: sean@SharePointInterface.com

LinkedIn: <http://www.linkedin.com/in/smcdonough>

Twitter: [@spmcdonough](https://twitter.com/spmcdonough)



SharePoint 2007 Disaster Recovery Guide

<http://tinyurl.com/SPDRGuide2007>



SharePoint 2010 Disaster Recovery Guide

<http://tinyurl.com/SPDRGuide2010>