NotesOn: Risk Management - Disaster Recovery & Business Continuity Definitions

Introduction (V1.5):

As promised in "NotesOn: Risk Management – Disaster Recovery & Business Continuity Essentials" this is the next post in the series to address DR and BC. Its purpose is to define the key terminology used in the business of "DR and BC" so we are all on the same page. Why? Well, Socrates allegedly said "If you would speak with me you must define your terms". Others have attributed this to Voltaire. But. Whomever first said it, they were absolutely, critically correct. To know anything at all about DR and BC you must first know and fully comprehend these terms as an absolute bare minimum.

Background:

For quite a number of years I've been doing Disaster Recovery (DR) as part and parcel of my various IT roles (sometimes specializing in it, sometimes not) and that experience has repeatedly bled over into Business Continuity. One of the most difficult aspects of both subjects that I've run into (after the primary issue of convincing "the powers that be" that their "DR Angels" are not likely to be looking over their shoulders unless they actively *do* something to help stave off the effects of disasters) is locating Standards and Best Practices that guide both DR and BC. Neither subject is of any value to the global Business community if any person can define terms however they wish, set any metrics against measuring success or failure they care to. This particular post is focused on the first element, "The Terms". Before venturing into Disaster Recovery and Business Continuity you must comprehend these, at a gut level. You must know these terms cold before taking on either subject, in any way, from any role.

Fundamental Definitions:

First, you should know that these definitions are ones I've developed and polished over quite a few years of practical, hard won, in the trenches, experience. You won't find them so defined in any dictionary. You may well run across similarly worded ones on other sites and in various books (I fully acknowledge that I have learned the subject by learning from others, as well as by doing), but I have tried to hone, and continue to hone these definitions for maximum clarity, crafting diagrams as appropriate to assist in the learning process.

Second, these are not in an alphabetical listing, I have put them in a "building block" sequence for faster learning so my recommendation is that you read them in the order they are in, at least the first time through.

Finally, feel free to challenge me on any of these definitions via my website. I am still learning, and hopefully always will be, and so I am more than willing to have "the errors of my ways" pointed out. :o)

Special Note: if/when I update these definitions or the included diagrams, and when I add links from definitions to an upcoming, detail level, post, I will update the version number in the "Introduction" header above; so check back from time to time for the latest edition.

Ready?

Here we go:

BC versus DR Diagram:

I'm borrowing a diagram from "NotesOn: Risk Management – Disaster Recovery versus Business Continuity" so that you have it to hand while studying the following definitions:

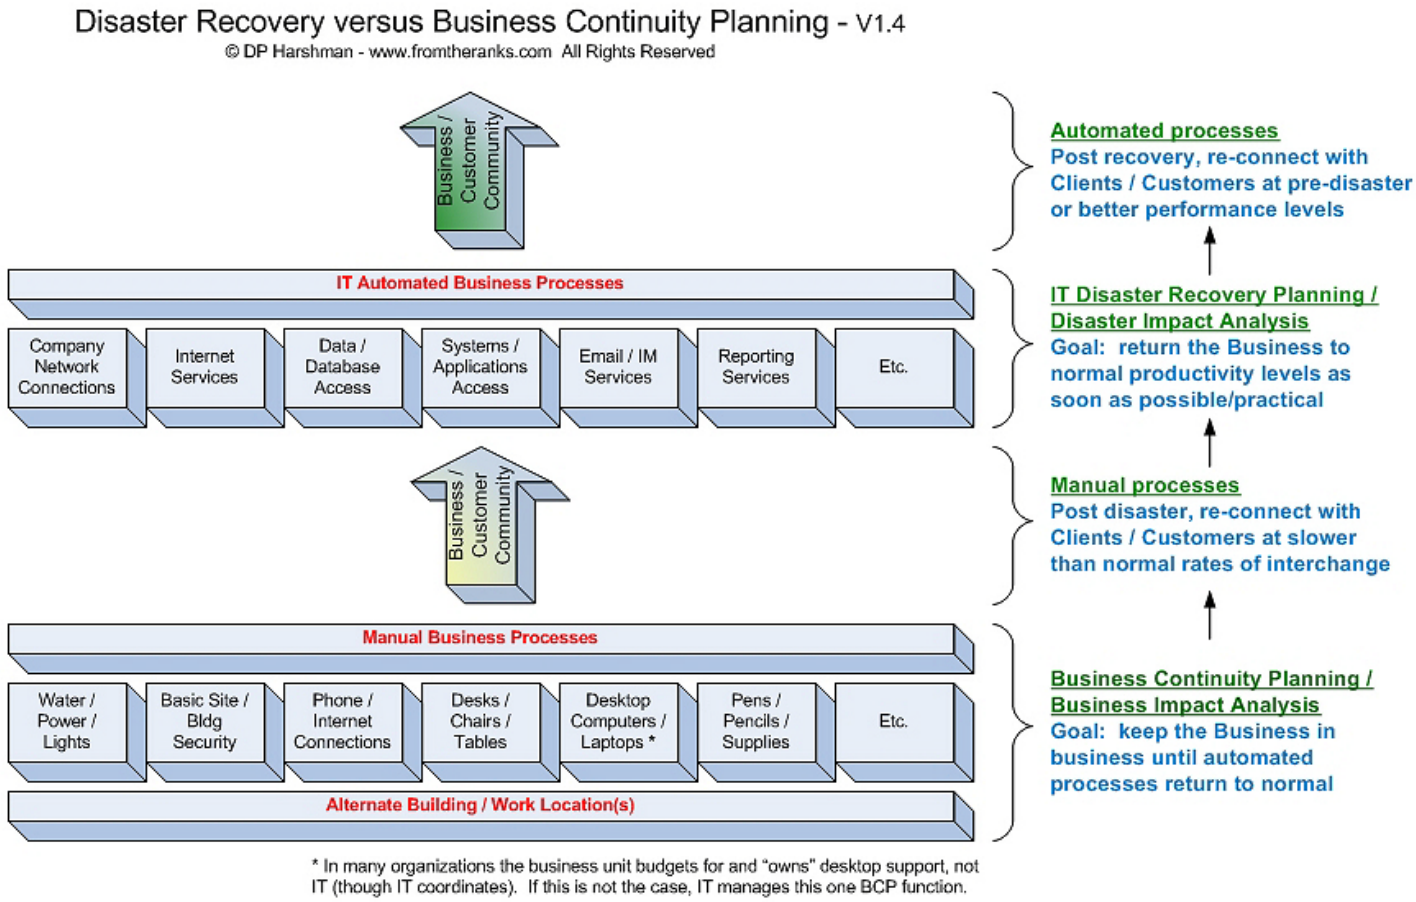

Business Continuity Planning (BCP):

The art and science of evaluating a business from end to end, top to bottom, bottom to top, to determine how best to ensure that a business's key business units and critical processes and procedures are able to continue, successfully, in the aftermath of a disaster of any significant magnitude. The end game of BC Planning is to, at least, maintain an operational relationship with and provide essential, mission critical, services for both existing, and new, customers.

To put it another way, BC Planning is the process of identifying, prioritizing, designing and implementing business processes that temporarily replace failed (typically automated) business processes for an extended period of time during or post a significant business interruption of some form. BC Planning anticipates the potential loss of and therefore includes identification and pre-arrangements for acquisition of all facilities and services necessary to safely, securely house and support the business unit(s) in a new or alternate location until a return to normal, or new normal, operations is obtained. It also plans for alternate communication methods and strategies to keep both BU personnel and appropriate public and public agencies well informed.

At a detail level, the BC team deals with the resourcing of pens, pencils, notebooks, copiers, printers and the like, but it also ensures business unit members have a wall with appropriate connections embedded in it to plug equipment into, *and* an appropriately configured laptop or desktop to plug into that wall's connections. You may choose to disagree that the BCP is also accountable for "computer stuff" in an emergency but it has been my observation that more often than not the BU's are already responsible for their PC, Laptop and PDA budgets during normal operations; typically through a third party provider. IT creates the images to be "burned" onto them but the contractual obligation for personal use equipment normally belongs with the BU. If this is not the case for your BU then IT becomes directly involved in that part, and only that part, of the BCP.

To put a point on it: BC Planning is focused on the business units, all of their critical processes and everything non-DRP related that supports those processes.

> *Editorial Comment:* though it may be obvious to everyone outside your business that a disaster has occurred to your business, if you manage your communications with your customers well, and depending on how your entire organization otherwise handles itself, you need not lose them to competitors during your "DR situation". Well executed BCPs (and DRPs) will help ensure you don't.

Business Continuity Plan (BCP):

The first substantial result of Business Continuity Planning, a BC Plan is a completed and tested living document that provides the organizational unit for which it is written a guide on how to properly respond to a disaster situation; or business interruption event if you prefer. Such a document is normally tied into an organization wide emergency response system (a.k.a. an EOC-ICS organizational structure) but can and should be able to stand on its own. A BCP document can have different styles but in the end it must only contain information and strategies vital to maintaining critical processes and procedures during and after an emergency; i.e. it should not be a three hundred page monograph that no one in their right mind would ever read; especially during a period of inherent chaos. [We'll get into examples of the content of a BCP and the EOC-ICS organizational structure and procedures in future "NotesOn" posts.]

Note: though they aren't often considered to be in need of one, it is entirely conceivable that a manual process and procedure may require a BC Plan as well so don't overlook this possibility. While almost everything is, in some businesses it is 95% or higher, not everything in a company is IT based.

Business Impact Analysis (BIA):

The primary investigative tool of Business Continuity Planning, a BIA is a carefully crafted survey form that is used to identify and examine the critical (and sometimes would-be critical) processes and procedures (PnP's) of each business unit (BU) for the purpose of determining and assigning an impact priority to them should they be interrupted for any significant length of time. The primary focus is first on the critical-to-the-business PnP's because they have the most positive impact on the company under normal conditions and thus represent the greatest potential harm to the company should they be abruptly or otherwise terminated.

As PnP's are being identified, a BIA analyst employs a "risk scale", as the comparative metric tool, to determine which losses pose the greatest threat to the company. To help simplify and accelerate the process,

risks are generalized somewhat and grouped into four or five BIA/DIA impact categories [see definition]. The most important PnP's become the immediate focal point of the BCP team as they want to and need to develop "manual" or other "work arounds" which can be quickly moved into place to cover the gap post a DR event.

Before the "work a round" phase begins however, a key next step is for executive management to evaluate all of the BIA priority requests against a "leveled playing field"; i.e. what may be "number one" for a BU may not be for the company as a whole. Note that this is often an iterative process requiring "back and forth" between the BU's and the executive team. As business processes more often than not utilize at least one IT system a key portion of the end game of this step is a rationalized over-arching strategy that sets the Tier level for, and high level architectural approach for, each mission critical IT system; and often a general set of standards and guidelines for the rest. Building High Availability (HA) systems is expensive so this is an essential step in obtaining the "best bang for the IT DR buck". As Chris Branch * put it in his pre-publication feedback:

> "I would say that one thing BIAs do is help provide an organization the ability to look at all the requests from the BUs and develop an overall enterprise strategy. This strategy should help feed / justify the requirements for what occurs during your IT architecture phase. … The completed BIAs should also serve as a point for an organization to look at what they think their enterprise BC strategy is and if the requests coming from the BIAs verify or require adjustment of the overall BC strategy."

Once the most critical PnP's are addressed the BIA teams take up lower risk ones, though it is rare to address all PnP's used by the BU (ex: though on-boarding a new employee is important, the BCP may not need to address a formalized manual process for it).

As the often in-depth interviews proceed, a BIA simultaneously identifies the key roles and players in each critical PnP, any required facilities and services needed to support them (including key IT systems/services) and, as well, succession planning, delegation of authorities, and the integration of the resulting BCP into the organization's overall emergency response procedures and structure, if one exists. The lack of disaster response related policies, procedures and organizational structures (i.e. an Emergency Operations Center built around an Incident Command System), can severely inhibit a company from responding quickly and rationally to a DR event; managing a disaster should not / must not rest solely on the shoulders of a single executive.

BIA/DIA Impact Categories:

No BC and/or DR team can risk assess everything in a group, department, division, business unit, segment, etc., or the company as a whole; it would take far too long. So it is necessary to select key, high level areas, or categories of potential impact, against which to do the assessment ("If you were to lose the <process or system> what would the impact be to <impact category>."). The names of these categories vary from one company to another, and one BCP/DRP team to another, but they tend to fall into the following risk classes:

BIA-DIA Impact Category Rough Equivalents V1-1

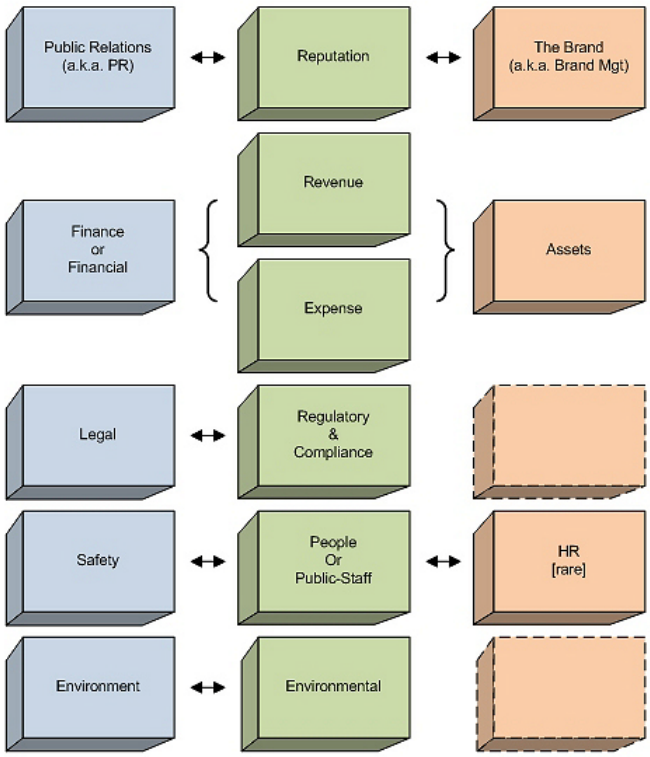

The columns are not fixed, the category names may be and often are interchanged.

Disaster / Disaster Event / Disaster Recovery Event:

A destructive event (also referred to as an 'incident') which in whole, or in part, causes a significant interruption in the ability of all or a portion of a company to provide full continuous service and support to its customers / clients. The event does so by impacting company assets, normal processes and procedures and potentially personnel. Though the event itself may not immediately cause the interruption, the cascading, down-stream effects of it may. An event of disastrous proportions is so declared because such an interruption directly or indirectly affects the survivability of the company.

From an IT perspective a disaster is more narrowly defined. It is a destructive event that significantly reduces the ability of at least one of the company's datacenters and/or the infrastructure surrounding the datacenter and/or the IT infrastructure at the company's facility(s) to provide the normal level of one or more services continuously expected from IT.

An IT disaster may or may not directly affect a business unit's assets, and vice versa, however, the end result is the same, an interruption of service, the disruption of the ability to deliver, for a significant period of time.

Disaster Recovery Planning (DRP):

Disaster Recovery Planning is the process of identifying, prioritizing, designing and implementing "fail over" services and systems that back up important to mission critical IT production services and systems (a disaster may also impact development and testing environments but that is of lesser criticality). To be complete, DR

Planning covers any and all monitoring and support activities and infrastructure necessary to keep vital IT services and systems up, or bring them back up, during and post a Disaster event. The DR team and their plans help ensure that when a business unit member's laptop or desktop is plugged into a wall connection (leading to the network and/or internet) the expected service/system on the other end is available.

To put an extra sharp point on it: DR Planning focuses on monitoring, protecting and provisioning all of the IT services critical to the business units.

Disaster Recovery Plan (DRP):

A Disaster Recovery Plan (only loosely similar to a BCP) is a *very focused* living document that addresses the bottom-to-top recovery of a *single* IT system critical to the company as a whole or a business unit. To be of any value, a DRP must be pared down to the absolute minimum text necessary; thereby allowing IT team members to maintain maximum focus on the technical aspects of system recovery. Think "bare bones". Think "skeleton". And you will be close. [We will go into the contents of a DRP in detail in a future "NotesOn" post.]

Disaster Impact Analysis (DIA):

A DIA is somewhat similar to a BIA except that it is tightly focused on a single IT system or service. During the interview process the DIA survey does not address (other than in passing) each business unit's (BU's) processes and procedures. Instead, it drills down into the usage of the system that has been deemed to be critical (often via a BIA). In brief, the end result of the DIA should include: (a) verification that the system being surveyed is a critical system, (b) identification of the risks to each BU if that critical system is lost for varying periods of time, and (c) the resultant business justifications for investing in a Disaster Recovery solution that will mitigate the risks and help reduce the pain level of losing that critical system.

Note: at this time the term Disaster Impact Analysis, and its acronym DIA, is not commonly used in the BC and DR community. Sometime ago I recognized that it was extremely beneficial to separate the focal points of the two teams via specifically tailored analyses forms: the BC team uses the BIA form to focus on the recovery and protection of business processes and procedures and the DR team uses a specialized DIA form to focus on recovery and protection of each of the critical IT systems and services that support the business units' processes. [There will be at least one upcoming "NotesOn" post detailing the content and use of a DIA.]

Disaster Recovery Architecture:

The art and science of planning, designing and, especially, implementing cost effective IT solutions that physically increase the "disaster event resiliency" of automated business processes and procedures; such solutions are focused first on mission critical systems, i.e. those that would "hurt the most" if lost to the business for an "expensive" period of time. While any solution that increases the resiliency of an IT system *may be* a worthwhile investment (though in these days of deep IT budget cuts "may" would be a hard sell) there are two generally accepted classes of DR solutions: "warm standby" [though "hot standby" is also in use, see its definition] and "cold standby" [though "cool" standby is also in use, see its definition], both of which are implemented in a properly selected and built out DR Datacenter or facility, a.k.a. a DR Site.

Disaster Recovery Site / Disaster Recovery Datacenter:

An alternate datacenter, or equivalent alternate facility, distinctly separate from the site physically housing the Production systems, that is capable of supporting identical IT functions and providing equivalent IT services during and after a disaster event.

Note 1: while distance between the Production and DR datacenters is not a functional part of the definition, it is a critical decision that needs to be made well in advance of a disaster event. As Alan Frawert * put it nicely in his pre-publication review comments:

> "Mission critical DR sites should be well planned … Often, well thought out DR plans fall short because the DR plan did not consider a building collapse or loss of communications for long periods of time. Things often overlooked include having a DR site along the same fault line. (This is actually what I am dealing with.)"

Note 2: DR Datacenters Standards and Best Practices will be discussed in great detail in an upcoming "NotesOn" post, but, for now, know that DR sites for your mission critical systems must be well planned out, and built out, well ahead of a disaster event. For instance, in the upcoming post we'll discuss the location of the DR site using the "risk zone / region" concept.

Technical Definitions:

Once the fundamental definitions above are "firmly under your hat" the next step is to drill down into the following to-be-applied, "rubber meets the road", aspects of DR and BC. Study the following two visuals carefully, come back to them often as you study the detailed definitions below them.

Definitions Diagrams:

This first diagram provides a visual definition of four key terms. If the BIA is done first the RTO and RPO are discovered by the BC group during their data gathering and planning efforts, while all four are, or should be, determined by (or vetted as approved by IT by) the DR team during Disaster Impact Analyses (DIAs). The reason RSL% and DR Readiness are not also defined and refined by BC teams can be found in each teams' objectives [see the four BCP and DRP definitions]:

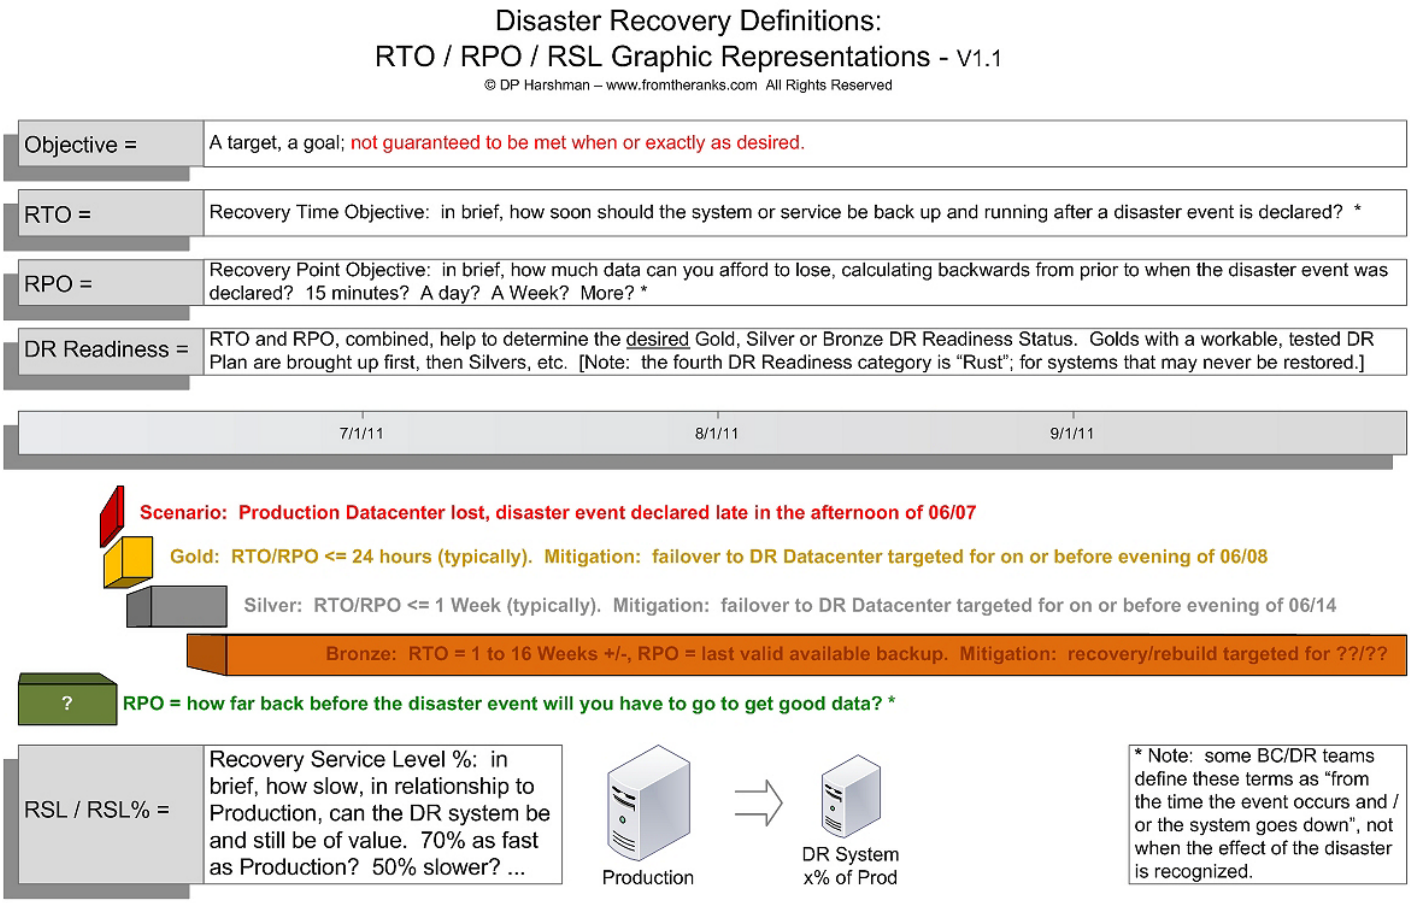

Usage Note: the timing of the declaration of a disaster can be crucial to IT's ability to meet or exceed RTO and RPO requirements. If a disaster event occurs but is unrecognized or unreported in a timely manner, the delay will have an effect on RTO and could have an impact on RPO.

The next diagram demonstrates the interaction between all five key DR and/or BC terms. Many more combinations are possible than those shown, ex: it is entirely possible to have a Tier 1 system with a DR Readiness Level of Silver.

System Name	Criticality Level - Tier	DR Readiness Level	RTO	RPO	RSL
System A	1	Gold	1 hour	15 minutes	100%
System B	1	Gold	2 days	4 hours	75%
System C	2	Silver	5 days	72 hours	80%
System D	3	Bronze	10 days	72 hours	70%

DR Readiness Level, Actual:

An Actual DR Readiness Level describes the defined and tested ability of IT to support a system during and after, and the ability of the architecture of that system to remain running during or be brought up after a disaster event within the requested system and data recovery criteria. More precisely, an Actual DR Readiness Level is: the verified as written, routinely maintained, regularly tested processes, procedures and systems necessary for a DR team or team member, or the DR system itself, to respond appropriately to an IT disaster event. Without regular testing of the system's or service's DRP any DR Readiness Level is a "Desired" DR Readiness Level at best. Actual and Desired DR Readiness Levels are typically broken out into three classes – Gold, Silver and Bronze – with a fourth class factually existing though rarely called out: Rust.

DR Readiness Level, Desired:

A DR Readiness Level is and remains a *desired* level when: (a) the business unit has requested a certain level of recoverability but the systems and procedures to sustain that objective have not been implemented and/or fully and properly tested; and/or when (b) the business unit has requested a certain standard (such as Gold), but Senior Management has either not approved it or has determined that a less stringent, and thus less expensive, solution (such as Silver) is more appropriate. Actual and Desired DR Readiness Levels are typically broken out into three classes – Gold, Silver and Bronze – with a fourth class unofficially but factually existing: Rust [see definition].

DR Readiness Level – Gold:

Generally, post a disaster event a Gold DR system must be up within twenty-four hours, maximum, and data loss must typically be limited to no more than twenty-four hours worth. Some critical systems may need more stringent standards. Gold level normally requires at least a "Warm Standby" DR system, e.g. a duplicate or near duplicate DR environment must exist at a remote location and be directly updated on a regular schedule as defined by the Business Owner. [Think of Warm Standby as the DR server(s) is always up and running in parallel to Production, thus is warm to the touch, and merely requires a "switch" to be thrown to activate.]

DR Readiness Level – Silver:

Generally, post a disaster event a Silver DR system must be up within 3-5 days, maximum, but acceptable data loss is still often limited to twenty-four hours worth. Some systems may require more stringent standards

such that, even though designated as Silver, the DR Plan may require a "Warm Standby" DR environment but at times "Cool Standby" is an acceptable option, e.g. a server must be available, or quickly made available, to accept a tape backup (or other restore source) but otherwise it need not be up and running on a 24x7 basis.

DR Readiness Level – Bronze:

This is the default DR Readiness level. If a system is not officially designated as Gold or Silver (which typically requires "certification" before achieving such designation), it is no higher than Bronze. Bronze implies that the system could be down for up to 16 weeks, or more, while it is being rebuilt (depending on hardware, raised floor space, rack space and personnel availability). Further, data restorations are wholly dependent on the availability of any backup tapes (or other restore sources) that may exist plus the resources to do the restore. Bronze roughly equates to: the Datacenter/Infrastructure/Applications teams will get to it when they can.

DR Readiness Level – Rust:

Though an unofficial designation in the DR community it is very real nonetheless. This type of system (typically a "desktop" or other informal application) has little or no probability of post-disaster recovery as the system's state may or may not be monitored, if known about at all, and backup and restore media and methods may or may not exist. Any system designated as "Rust" will likely have to be rebuilt from scratch, if it is ever rebuilt. Rust has the lowest recovery priority level of all systems.

> *Historical Note:* I realized some years ago, while drilling down into a long list of systems that needed "DR'ing", that there was a missing Level. It dawned on me that some of those systems could never been brought back up, post a DR event, and would, essentially, turn into rust. Hence the Level name.

Recovery Time Objective (RTO):

One of the most crucial pieces of information discovered with the DIA and BIA is the RTO value. A strategic value used to help determine both a system's Desired, and future Actual, DR Readiness Level, and a business process's recovery time, it is defined as: 1. [DRP] the desired chronological time from the beginning of a declared disaster event to when the application is "Ready for Use" by the users. 2. [BCP] the desired time goal for the recovery and re-establishment of the specified business function (which may or may not be supported by IT). 3. Most stringently, it is the "maximum allowable outage window", regardless of the disaster declaration time [see Usage Note]. Example: if RTO = 12 hours per both the DRP and BCP then, from the time of the disaster declaration (or its occurrence if the stringent standard is used), if at all possible, either the Production or the DR environment must be acceptably up, running and available to the users to resume their business process within that 12 hour time frame. Compare RPO, RSL and DR Readiness Level.

> *Usage Note: b*ecause IT's RTO clock does not traditionally start until the disaster is declared (i.e. someone knows about it and reports it), DIA's and BIA's may also determine the "maximum allowable outage window", a value which does not take into consideration the declaration of the disaster but rather the moment that the system goes out of service. What this most stringent value does is force monitoring of all systems related to the business process in question. As an example: the RTO might

be 12 hours from the DR event declaration but the *maximum allowable outage window* may be 24 hours, period. Such an additional constraint could materially affect the architecture of the DR system as it presupposes that recovery must begin before IT (or anyone else) knows about the disaster event.

Recovery Time Objective (RTO) - Cumulative:

Cumulative RTO is the accumulating effect on a target system's Recovery Time Objective (and actual Recovery Time) if one or more layers below the target system also goes down. The layers being considered are: (1) the target system (2) one or more supporting systems -- systems upon which the target system depends (3) the supporting infrastructure – networking, communications, ETL utilities, etc. and (4) the primary datacenter in which the target system resides. Failure of more than the target system will result in an extended outage that will impact the typically expressed RTO value. Example: the requested RTO for System A is 24 hours but, in the absence of any failover systems and environments, if Layers 2, 3 and/or 4 are also down the requested RTO will not be met, unless it knowingly took into account the cumulative recovery time of all layers.

Ref: NotesOn: Risk Management – Cumulative Recovery Time Objective

Recovery Point Objective (RPO):

RPO is discovered during both the BIA and DIA surveys (if during a BIA that data is then fed to IT). RPO is a tactical IT-centric value used to help determine both a system's Desired and Actual DR Readiness Level, and is defined as: 1. The stated permissible amount of data loss as measured in time from the point of the declaration of a disaster event. 2. The desired pre-disaster point in time (usually stated in hours or days) to which data must be restored in order to resume processing without significant impact on the company due to critical data loss. 3. Most stringently, it is the "maximum allowable data loss" period [see Usage Note 1]. 4. Loosely, how much data can the company afford to lose before it hurts. Example: if RPO = 4 hours then, starting from the time of the disaster declaration (or its occurrence if the stringent standard is used), no more than 4 hours of previously entered data can be lost. See RTO, RSL and DR Readiness Level.

Usage Notes:

1. Because IT's RPO clock does not traditionally start until the disaster is declared (i.e. someone knows about it and reports it), DIA's and BIA's may also determine the "maximum allowable data loss", a value which does not take into consideration the declaration of the disaster but rather the moment that the system goes out of service. It represents the total amount of acceptable data loss, period. For example: a Bank may have a maximum allowable data loss RPO of "zero seconds no excuses".

2. RPO and RTO are distinctly separate values. One is not dependent on the other. The time a system is brought up (or a business process is returned to "normal") has no bearing on how much data was lost as a result of the event. They are independent factors, though they do have a mutual bearing on the design of the DR system.

3. Whether discovered during a BIA or a DIA, the RPO value *must be* "accepted" by IT. There may be reasons why the RPO value requested by the BU cannot be met by IT; "It's too expensive!" is the most common reason but there could be technical issues as well.

Recovery Service Level % (RSL or RSL%):

Not yet widely known or employed in the DR community *, RSL% is a tactical value used to help determine a DR system's performance level, and to a degree its Actual DR Readiness Level and is defined as: 1. A measure of the minimum desired performance level of the DR system in comparison to the Production system's current performance level, expressed as a percentage of Production's performance. 2. Loosely, how slow will the business unit(s) let you get away with the DR system being during or post a disaster when compared to Production. For example: RSL = 75% indicates it is desired that the DR environment operate at a level that is at least 75% of the original production system's performance level. This value has a direct effect on hardware and network requirements for the DR system. See RPO, RTO and DR Readiness Level.

* *Historical Note:* RSL% is a value I recognized and a term and acronym I coined quite a few years ago as it had and has extreme value to the DR system design process.

Tier 1 / Criticality Level 1:

In DR planning, Tier 1 systems (sometimes known as Criticality Level 1) are the most mission critical systems for a business unit, a Segment/Division and/or the company as a whole as measured against the potential significant impact of their loss in one or more of the BIA/DIA categories [see definition]. As a rule it is difficult if not impossible to "replace" these with manual procedures and paperwork, and if possible not for long. Tier 1 may be "guessed at" initially but its status should be confirmed via a Business Impact or Disaster Impact Analysis (BIA or DIA). Note: because a system is rated as Tier 1 does not imply it automatically receives a Gold DR Readiness status [see definition]. Research has shown that, due to the costs associated with implementing a Gold standard, some companies have decided that all Tier 1's will be assigned a DR Readiness Level of Silver.

Tier 2 / Criticality Level 2:

In DR planning, Tier 2 systems (sometimes known as Criticality Level 2) are important to the company but not as critical as Tier 1's when measured against the potential impact of their loss in one or more of the BIA/DIA categories [see definition]. One indication of a Tier 2 system is that the work done by it can temporarily be replaced by manual procedures and paperwork, allowing recovery of the on-line system to be delayed for up to a week or so. Tier 2 systems are often assigned Silver DR Readiness Levels (though on rare occasion a Tier 2 may require Gold).

Tier 3 / Criticality Level 3:

In DR planning, Tier 3 systems (sometimes known as Criticality Level 3) are of value to the company but are presumed to able to be "down" for more than a week without significant impact to the company in one or more of the BIA/DIA categories [see definition]. An indication of a Tier 3 is that its services are somewhat easily replaced by manual procedures and paperwork for an extended period of time without serious BIA/DIA

category impact. It is important for business owners and users to understand that post a major disaster Tier 3 systems with a Bronze DR Readiness level could be down for up to 16 weeks, or longer.

Standby Systems/Datacenters Diagram:

There is a great deal of discussion on the internet, and elsewhere, on the various definitions used to describe standby systems and standby datacenters (defined below); one can get into arguments over these terms without too much difficulty at all. In an effort to standardize each of them I offer the following definitions.

What is most important to know is that definition of the terms vary depending on whether you are speaking of datacenters or speaking of DR / failover systems in the datacenter. This is a primary source of confusion so we'll clear it up now. Note: I first ran across a clean definition of Standby Datacenters in a Microsoft TechNet article. I have tweaked it, expanded it, to my liking but credit goes to the author for the original definition.

To aid in your understanding, here is a diagram to study the definitions against:

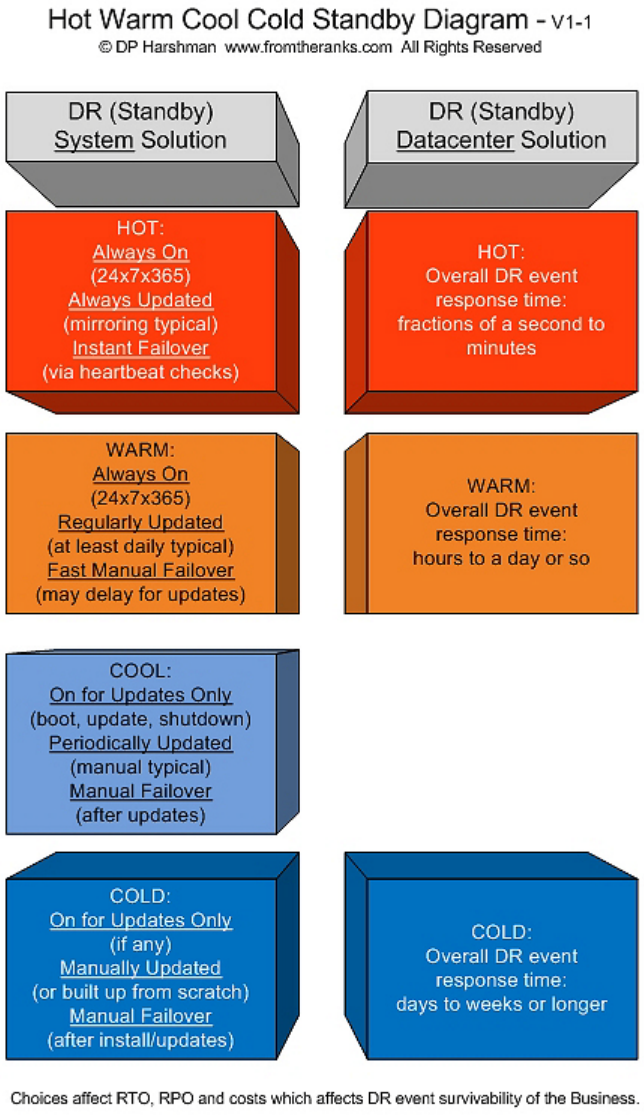

Usage Note: the problem with the terms hot, warm and cold standby in relation to datacenters is that rarely are datacenters entirely one form of standby or another. Often they have a mixture of DR / failover solutions. So, the terms most accurately describe the *overall* not the uniform readiness.

DR Datacenter Solution – Hot Standby

When speaking in reference to backup (DR) datacenters, a hot standby datacenter is one whose external and internal systems are, overall, able to respond to a DR event, an unplanned interruption, in a time period ranging from fractions of a second to a minute or so at most.

DR Datacenter Solution – Warm Standby

When speaking in reference to backup (DR) datacenters, a warm standby datacenter is one whose external and internal systems are, overall, able to respond to a DR event, an unplanned interruption, in a time period ranging from hours to a day or so at most.

DR Datacenter Solution – Cold Standby

When speaking in reference to backup (DR) datacenters, a cold standby datacenter is one whose external and internal systems are, overall, able to respond to a DR event, an unplanned interruption, in a time period ranging from days to possibly weeks.

DR System Solution – Hot Standby:

The term is, as often as not, commonly held to be synonymous with “warm standby”. When referring to DR systems, the technical difference between the two is minimal but existent: 1. “Hot standby” is typically used to differentiate a fully automatic failover solution, wherein if Production fails the DR system engages nearly instantaneously, from a “warm standby” failover solution which often contains a small to moderately long, manually triggered, “switch over” period. 2. In brief, a true “hot standby” is assumed to (a) be always on, (b) be always updated with the latest data and software changes and (c) is able to fail over instantly on a 24x7x365 basis. Compare to Warm, Cool and Cold standby solutions.

DR System Solution – Warm Standby:

The terms “warm” and “hot” are often used interchangeably but unless speaking “generically” there is a difference: 1. In a "warm" DR system solution, the DR system is always up-and-running, per the approved DR Readiness Level, in parallel with the Production system, but, the warm standby has less stringent routine replication schedules than a “hot standby” system; typically ranging from minutes to daily, or even weekly. Though a warm standby system can be and often is brought up manually, the DR architecture may include designs for auto failover, a.k.a. auto switch over, from Production to the DR system; with consequent loss of the data missed since the last replication period. 2. In brief, a true “warm standby” is assumed to (a) be always on, (b) be updated regularly but not constantly (i.e. not mirrored), and (c) is able to be failed over to, usually manually; if any last minute updates are available they would normally be applied before “boot up”.

DR System Solution – Cool Standby:

A "cool standby" DR system varies significantly from "hot" or "warm": 1. A "cool standby" DR system is, as a rule, never "on", except when being updated, i.e. the box's temperature borders between warm and cold. 2. In brief, a true "cool standby" is assumed to (a) always be off, except (b) when being updated, and (c) is able to be failed over to only after the system is manually booted up and updated further, if further updates are available.

> *Usage Note:* this approach has value for low priority, low tier systems where data or software updates are irregular and/or the "pain" from the loss of the system is fairly low. You are still paying for the hardware and software, just not the electricity and cooling. An advantage to "cool" standby is that as it is turned off nearly all the time there is less chance of a power surge or hardware failure bringing the system down; of course the disadvantage is that the server may never be fully "burned in", a process which flushes out weak or failure prone components. By the way, "cool" is considered "cool" and not "cold" because at least it is booted up once in a while.

DR System Solution – Cold Standby:

The definition of "cold standby" when compared to "hot", "warm" and "cool" is almost intuitive: 1. A true "cold standby" DR system is, as a rule, little more than bare metal hardware, i.e. it is almost never on, if on at all, and may never have had any software installed on it, or if the basics are installed rarely are any data updates sent to it, or if data updates have been sent they are likely old and out of date. It is safest to think of "cold standby" as a blank slate ready for a DR team to do something with it; depending on resource availability and the recovery priorities of the moment. 2. In brief, a true "cold standby" is assumed to (a) always be off, perhaps never unboxed, (b) updated once to ensure it functions, if updated at all, and (c) able to be failed over to only after being manually built up from scratch, or if somewhat operational the latest available installs, patches and backups loaded onto it are tested before being brought up.

> *Usage Notes:* With a "cold" DR system solution, the DR Datacenter team typically has promised to have the floor space, telecom equipment, power, cooling, and network circuits available to allow them, or you, to build out and hook up a "new box" from scratch. Beyond that service level, the DR datacenter team has made no investment or installed any component in preparation for any DR event that impacts that Production system. They may or may not have the physical hardware on hand, depending on their contractual agreement or your approved budgets (if your company owns the datacenter). They may or may not have *any* of the necessary software on-hand to stand up the server; depending instead on the delivery of backup/restore media to the DR datacenter after a DR event. They may or may not have the personnel available to do the build-out.; depending on the nature of the DR event, their higher priorities and, of course, the Service Level Agreement (SLA). An alternative cold standby solution to having a "bare metal box", or the next thing to it, sitting around waiting for an emergency is the repurposing of an existing (lower priority) server/device, disabling its prior use. Whatever the solution, what is always true is that all cold standby solutions take time to activate,

minimally hours, often days, or possibly weeks; but a cold solution may be entirely acceptable to a business if their tolerance for the down time is sufficient to allow a delayed recovery.

DRP Testing:

DRP, or DR, testing is the process of taking the completed DR Plan and "exercising" it before it has to be used in the real world, i.e. during or after a real disaster event. It is most important to know that a DRP is almost absolutely, utterly, and completely worthless if it has never been tested, tested only once and forgotten or tested only in the most glib, unprofessional, manner. The reason is that both IT systems and organizations change constantly and these changes will almost certainly have a direct effect on the DRP. So, if it is not routinely tested, at least annually, ideally end to end (though there is a "stair step" approach to testing as opposed to "all or nothing"), all of the time and effort that went into the draft DRP is likely to be wasted.

Usage Note: the typical sequence of DR testing is: Table Top, Functional Failover, Planned Failover, Actual Failover.

DRP Testing - Table Top:

The first test for *all* new DR Plans is to, as a team, walk through the completed draft DR Plan from the first page to the last, examining every step and statement. The goal is to catch any potential flaws, so they can be corrected, and to recognize missing data so that it can be added. If errors or omissions are discovered, and there almost always are the first time through, the DRP is re-updated and re-table top tested. At the end of the (often iterative) test everyone involved must be comfortable that there are no *major* omissions or errors in the plan; the fine point errors will be flushed out during the more detailed tests.

DRP Testing – Functional Failover:

This type of DRP test exercises the basic elements of the system's/service's DR Plan and its architecture without actually taking down either the Production or the DR servers. The steps include: (1) turning off the Production database's log shipping long enough to make a clean DR Test copy of it; (2) the DR DB Test copy is then moved to a separate location on the DR server, if possible, or another server if not; (3) the application on the DR system's application server is modified (re-configured) to point to the DR DB Test copy; (4) if Citrix is in use a new icon is created that points to the DR test copy of the application; (5) the tester then clicks on the DR version of the app which should be accessing the DR DB Test copy; and (6) the tester creates a record in the DR DB Test copy that is then verified as not existing in either the Production database or the actual DR database. The changes that were made are then reversed: the test copy of the DR DB is dropped, the DR copy of the application is re-configured to point back to the true DR database, and if used the Citrix/VPN DR icon is made invisible (after being pointed back to the true DR DB); so all is back in place in case of a true DR event. Note: the risk with this type of test is that it is not exercising the entire DR infrastructure from end to end, it only tests the links between the DR App and the DR DB and the immediate supporting infrastructure.

DRP Testing – Planned (Full) Failover:

This test exercises the entire Prod/DR system from end to end and so must be carefully planned ahead of time. A full failover test can be done by: (1) backing up the current DR system; (2) if the DR system's updates are not current, taking a Production system backup (or image) and restoring it to the DR Site location for the system; (3) then making all necessary configuration changes, such as DNS, DFS, .INI file changes, etc.; and (4) once the DR environment is ready to come up, shut down Production and send the users to the DR site. This exercises all of the app, security and infrastructure components of the system, ferreting out permission issues, firewall tunneling issues, etc., so is best done over a weekend or during non-Production hours. Every single actual step taken of the test should be documented and then compared to the DRP during a detailed post-mortem that absolutely should be done afterwards.

DRP Testing – Actual (Auto) Failover:

This is not a planned test. If there is a system failure, due to hardware problems, production datacenter issues or external influences, the Production system should auto fail over (if so designed) or is manually failed over to the DR environment. Ideally, IT then takes advantage of this occurrence by fully documenting all steps taken to bring up the DR system and, later, to fully recover the Production system. Once the event is over the entire DRP team analyzes the failover, from end to end, for DR successes and failures. Any "hitches and glitches" in the DRP are corrected and re-tested at the first opportunity.

Summary:

I sincerely hope that you find this series of definitions of value. I look forward to your feedback and any resulting discussions.

* Before closing, I would like to thank Chris Branch, a top flight Business Continuity professional, and Alan Frawert, a top notch IT professional, for taking their personal time to review this post before publication. Having not one but two "sets of eyes", especially on something as important as this subject, is absolutely invaluable, so: "Thank you."

Hope this helps,

DP Harshman

PDF Link