Disaster Recovery—

The Plan That Makes You Ready

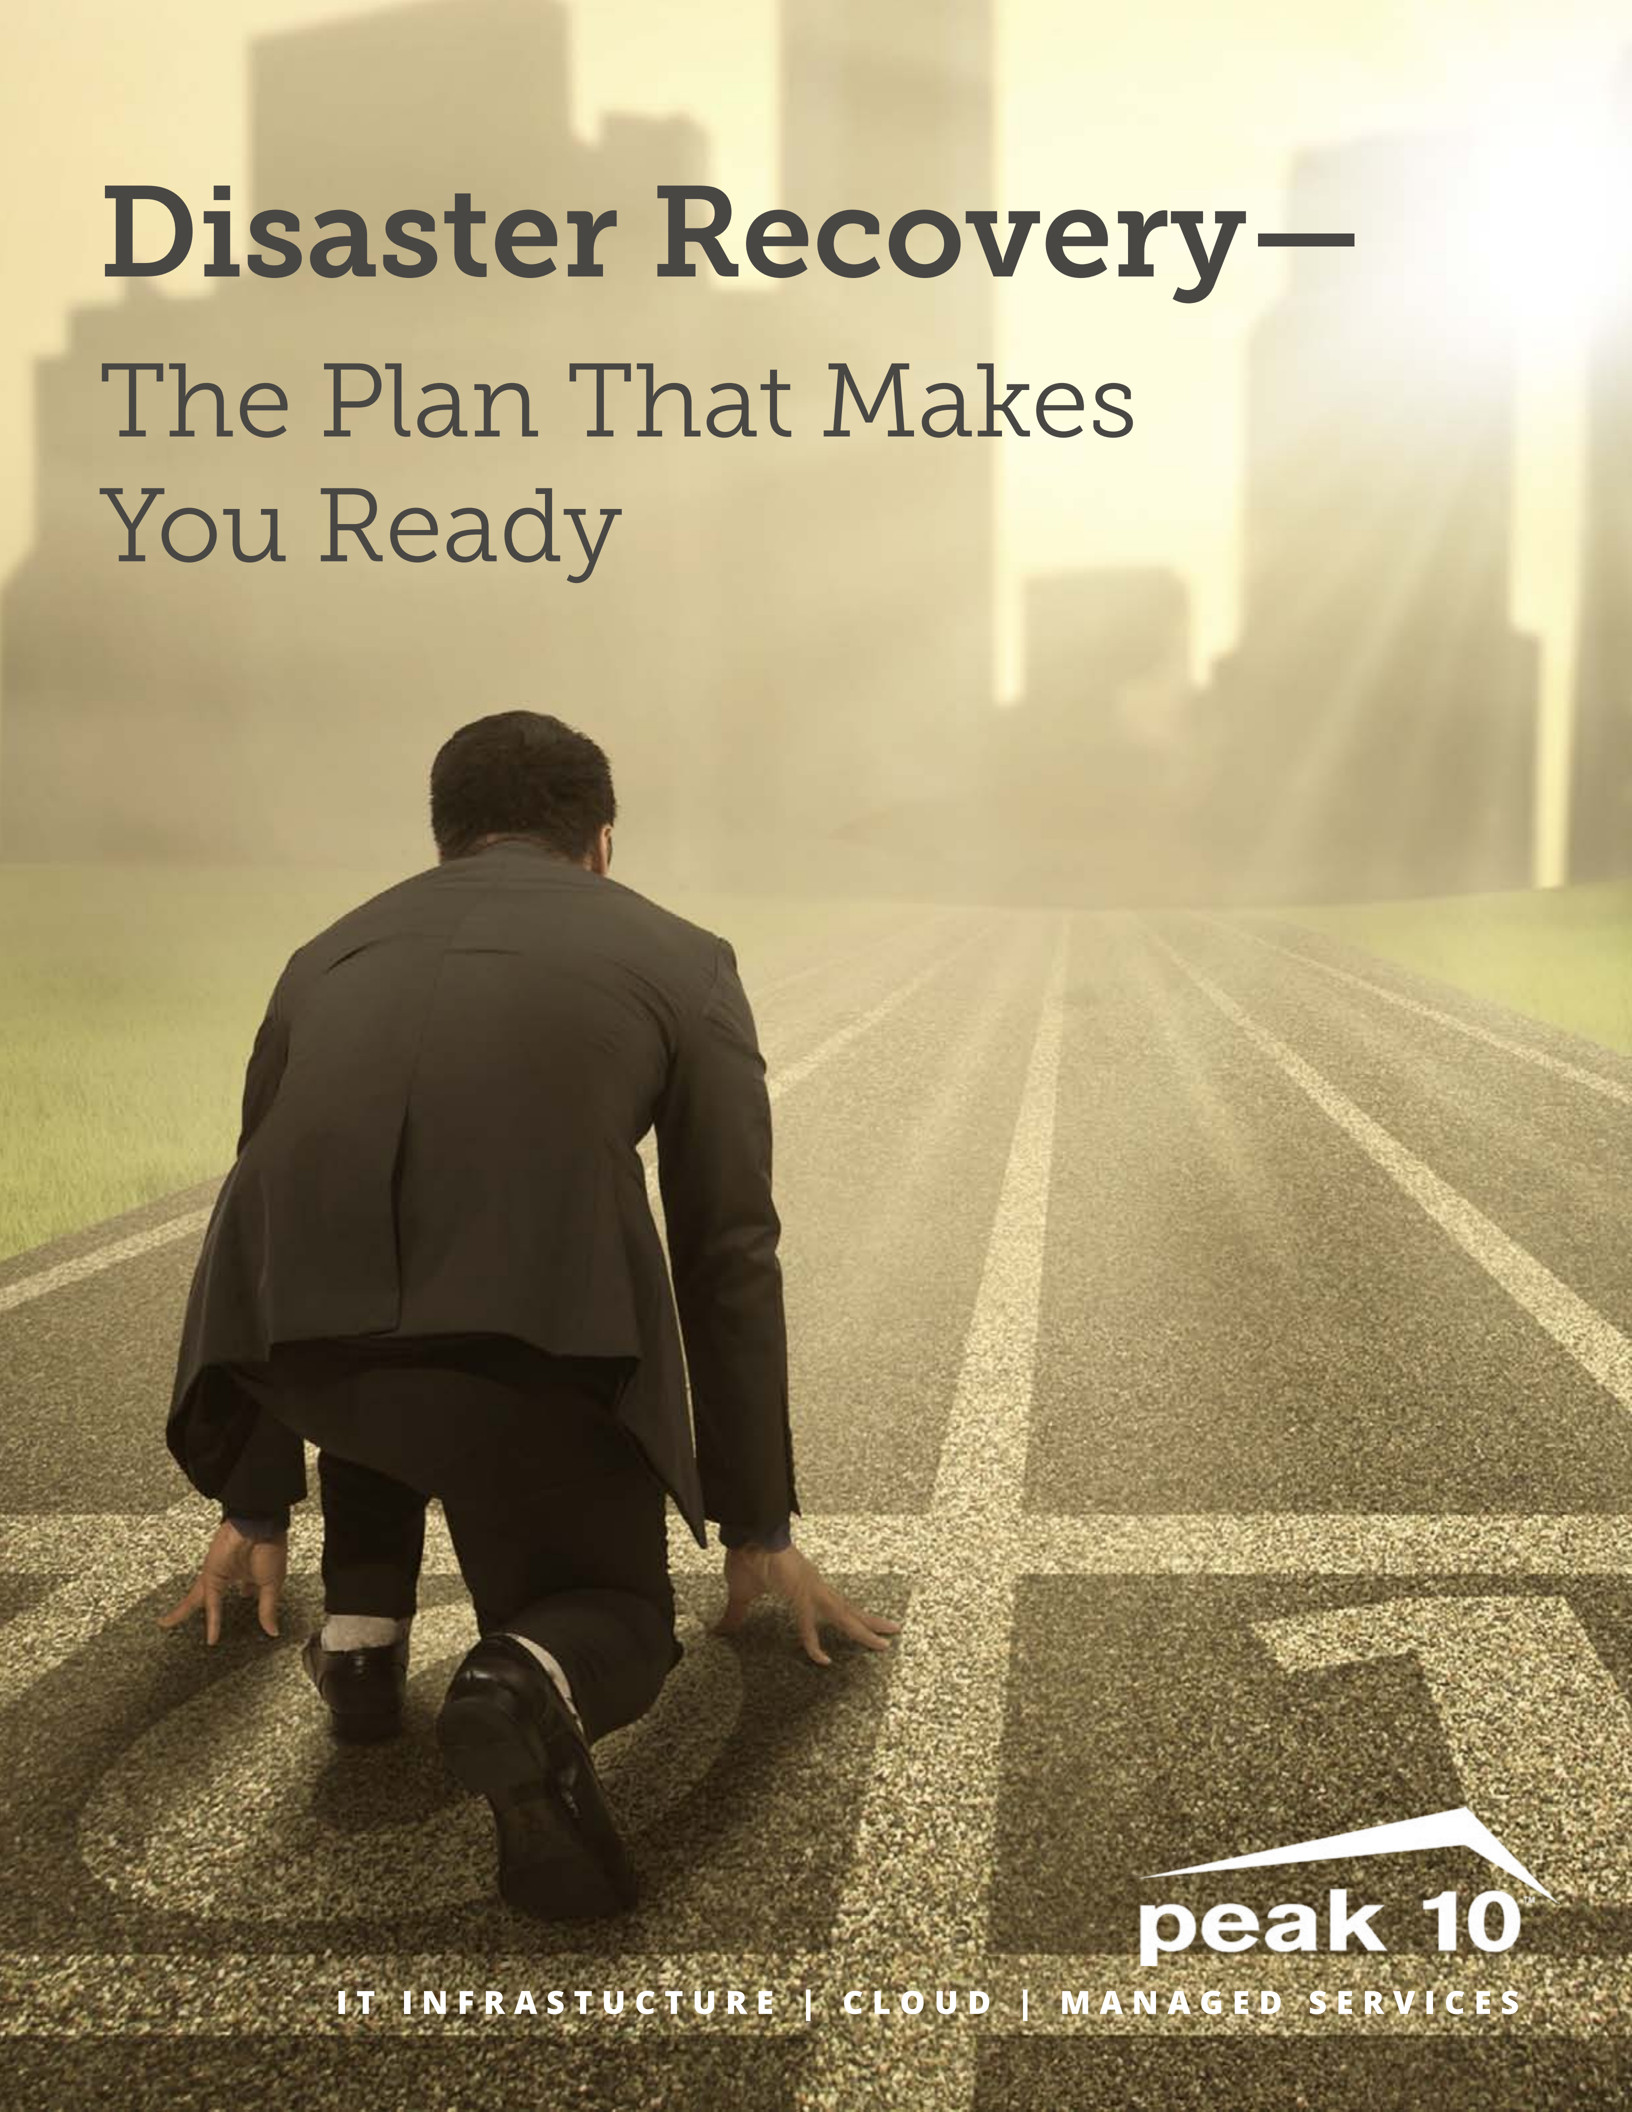

peak 10

IT INFRASTRUCTURE | CLOUD | MANAGED SERVICES

Intro

Businesses’ confidence in disaster preparedness and their ability to meet their disaster recovery (DR) objectives have fallen over the past few years. According to “The State of Business Technology Resilience, Q2 2014,” from Forrester Research, Inc., complexities such as application interdependencies, the increase in the number of mission-critical workloads, and lengthening time intervals from emergency declaration to recovery have infrastructure and operations professionals on edge. (1)page 2

What drives these businesses to want to implement improvements are not the typical reasons, such as application availability, reputation or cost of downtime, although these still are critical indeed. There are legal and regulatory concerns that have come to the fore as primary drivers for more effective, resilient and timely recovery. (2) page 11

This being the case, getting back to the fundamentals of DR planning is more important than ever. DR is a load-bearing wall in business continuity (BC) strategy planning for any enterprise or organization, given our ever-increasing, day-to-day reliance on all things digital.

By necessity, DR planning must be an intra-company, collaborative effort. Lines of business, security and risk management, contract administration, and right on up to the C-suite all have legitimate and vital input. Effective planning may require inter-company considerations, as well, involving suppliers, partners, associates and other outside constituents digitally hooked into your business operations and, hence, the recovery of those operations.

In this E-book, we are going to discuss DR-related subjects we believe are critical to developing an effective plan. While not in the DR planning business, Peak 10 does provide infrastructure and cloud services that bring a DR plan to life. Every day, literally, we see and hear from businesses that struggle with preparedness planning, solution design and effectiveness testing. Additionally, we are directly engaged with helping to ensure their solutions function as designed, all the time.

1 DR From The Top

The DR plan is obviously a key part of business continuity…

Burdened by data privacy rules and regulatory compliance on multiple fronts, as well as the escalating cost, disruption, and customer impact resulting from unplanned outages, effective DR planning is a business necessity that IT is charged with driving.

As such, the criticality of DR planning, its budgeting and validation must be ordained by the highest levels of the organization. "Without executive level support, DR planning and implementation are destined for failure," said Steve Hasselbach, Peak 10 Solutions Engineer.

DR planning is not an event but, rather, a continual process of forethought, collaboration, testing and adapting. Everyone must be cognizant of the potential effect that their programs and business decisions will have on the plan, and ensure that these influences are factored in. The plan must operate in lock-step with the organization's business strategy. It needs to be a reflection of the dynamism of the entire organization, as close to real-time as is practical; otherwise, you will be attempting to recover things that were, not things that are.

If executive-level support is not forthcoming, the CIO must become the tireless advocate for attaining it. Business is about managing risk for reward … recognizing it, measuring it, deciding on a best course of action, and assessing results of those actions. Presenting DR initiatives, investment options, failure consequences, etc. as a risk/reward business case to directors and executive management allows DR to be evaluated by the same metrics as other investment options. Given the consequences of DR failure, it's a compelling argument to make.

2 A DR Plan Should Be Only About DR

Some companies lessen their chance of developing an effective disaster recovery plan by introducing non-DR aspects into it. Their approach is too broad, or it attempts to include business aspects that are better served by others in the organization.

Mistakenly thinking of DR as a BC plan is where many plans fail. DR is about the recovery of IT services after a declared emergency and a planned return to normalcy after the fact. BC, on the other hand, is a holistic plan meant to ensure business resilience at all times – before, during and after operational disruption.

The DR plan is obviously a key part of business continuity, but every department within an organization contributes its own portion to the business continuity management plan. Who are alternate suppliers? Where will manufacturing move? What if we need temporary workers? What is our communication plan? What happens if we have a chemical spill?

Mistakenly thinking of DR as a BC plan is where many plans fail.

Gartner Inc. offers this guidance for scoping a DR plan:

> **"IT leaders should focus DR planning on the recovery of IT services, and should clearly define the intended use and scope of the plan. This includes developing a concise statement about what's included and what's not, who the intended audience is and how the document should be used. The scope should identify the specific locations, businesses, companies and functions covered by the recovery plan. Additionally, the IT DR management plan should clearly articulate how it fits into the BCM [business continuity management] structure."** (3) page 3

Furthermore, in the heat of a crisis is no time to be reading DR documentation that resembles "War and Peace." DR plan documents need to be logically ordered, modular, simple and declarative. It is important to have printed copies available at multiple locations, in addition to being accessible online.

3 The Hierarchy Of Recovery

You could attempt to recover all applications and workloads following a disaster. However, a risk/reward analysis for doing that would clearly show that it's a poor commitment of resources, even if your pockets were so deep you could afford it.

"Protect everything" and "low cost" are mutually exclusive, at least given the current state of technology. Determining which applications and workloads are most critical, i.e. which ones your business needs first and foremost to resume operations after a disaster, is essential to designing an effective recovery plan.

Inventorying applications is the first step in profiling each one and categorizing its value to the organization. The process will pinpoint which business services are most critical, as well as the most costly when they are offline. Many organization are surprised to learn just how many applications they are supporting, some beyond their useful lives.

IT and business leaders must work together to understand what an application does, who uses it, how it fits, and what its value is to users and the business. This application mapping process also reveals interdependencies with other applications, departments or third parties, as well as ties to regulatory or compliance requirements.

Application inventorying and mapping is difficult and time consuming. But, without this hierarchy of criticality, taking the next step – establishing Recovery Point Objectives (RPO) and Recovery Time Objectives (RTO) for individual workloads – becomes a guessing game that can result in flawed data-backup and DR strategies.

...a risk/reward analysis for doing that would clearly show that it's a poor commitment of resources, even if your pockets were so deep you could afford it.

The Hierarchy Of Recovery

How long can you be without a specific application? RTO defines the maximum tolerable or allowable duration of an outage and, consequently, the data replication or disk-versus-tape backup requirements for that application. If your RTO is zero (no tolerance for downtime) then you may decide that having a completely redundant infrastructure with replicated data offsite is cost-justifiable. If your application's RTO is 48 hours or 72 hours then tape backup may be acceptable.

RPO dictates how much data you can afford to lose and, therefore, how frequently data must be backed up. For example, if you do a nightly back up at 12:00 midnight and the unthinkable happens at noontime, then all the data generated in the intervening 12 hours is lost. RPO in this particular context is the previous day's backup. If that is acceptable, then that is your RPO. If it's not acceptable – say, for regulatory compliance reasons – then you'll want to think about another solution for those particular data.

The tighter your RTO and RPO, the higher the cost to design the appropriate solution. But, now you have the business impact analysis you need to wage a convincing argument about investing in disaster recovery. You now know which of the hundreds or thousands of applications in use at your company are the priorities around which to develop your strategy (or the pain of failing to do so). Focusing on which critical business processes must be recovered first and in what order is the difference between a well-orchestrated, successful recovery and interminable chaos.

RPO dictates how much data you can afford to lose and, therefore, how frequently data must be backed up.

4 This Is Not a Test

Setting up a replicated site is relatively straightforward, even if it is costly.

These are words you hope you never hear, "This is not a test," especially when you are unsure how your DR plan will work because you haven't tested it. Seventy-seven percent of companies are not fully confident their DR plan will work.

Not surprisingly, companies that rigorously prepare their organizations for DR are most successful at recovery. For companies that maintain their own DR site and systems, preparation includes a full-blown failover test of their DR plans, including failback, at least once a year. It's costly, disruptive and sometimes looked upon with disfavor by management, but there is no other way to be sure that all your efforts and investments have not been in vain. As disruptive as it may be, it's nothing compared to going into actual recovery mode only to find gaping holes in your plan or infrastructure.

Setting up a replicated site is relatively straightforward, even if it is costly. Maintaining that replication is certainly more difficult, more costly and more resource-consuming. However, making sure everyone in the organization knows his or her roles, what to do and where to be when disaster happens is probably the most overlooked critical success factor in DR plan implementation.

"Best practices dictate that at least once a year employees are told not to come in to work the next day but, instead, to work using the alternate back-up system just as they would in a true disaster," says Steve Hasselbach, Peak 10 solutions engineer and DR specialist. "Only that way can you know people really understand what to do and how to do it."

5 Know What You Are Preparing

DR evokes images of naturally occurring calamities—earthquakes, tornadoes, hurricanes and such. While they seem to be happening more frequently and impacting larger swaths of the country, these weather-related events are not the cause of most outages. And, while something like civil unrest may be problematic, it's probably a "BC management" matter under the responsibility of another entity.

The majority of IT outages result from operational issues such as power failures (43%) and hardware failures (31%). The frequency of outages from human error (13%) surpasses all natural-disaster events. (4) page 10

The majority of IT outages result from operational issues such as power failures (43%) and hardware failures (31%).

Finding the right balance between a plan being too general and overly granular in DR preparation is not easy to do. There is no correct answer. As noted in an earlier chapter, however, a DR plan should be only about DR. Gartner, Inc. offers this advice to its clients:

> **"Organizations should plan for disaster scenarios based on the likelihood of that scenario occurring. Scenarios based on criteria such as notification time (e.g. 'a tornado warning is in effect starting tomorrow at noon'), type of disaster and potential business impact should be established only if there are material differences between the response and the recovery procedures for that type of event. Unlikely (low-probability) events should be identified; however, rather than developing full-scale plans, organizations should focus on general IT DR management principles – a common set of steps that are taken by the recovery team following most types of disruptive events." (5) page 4**

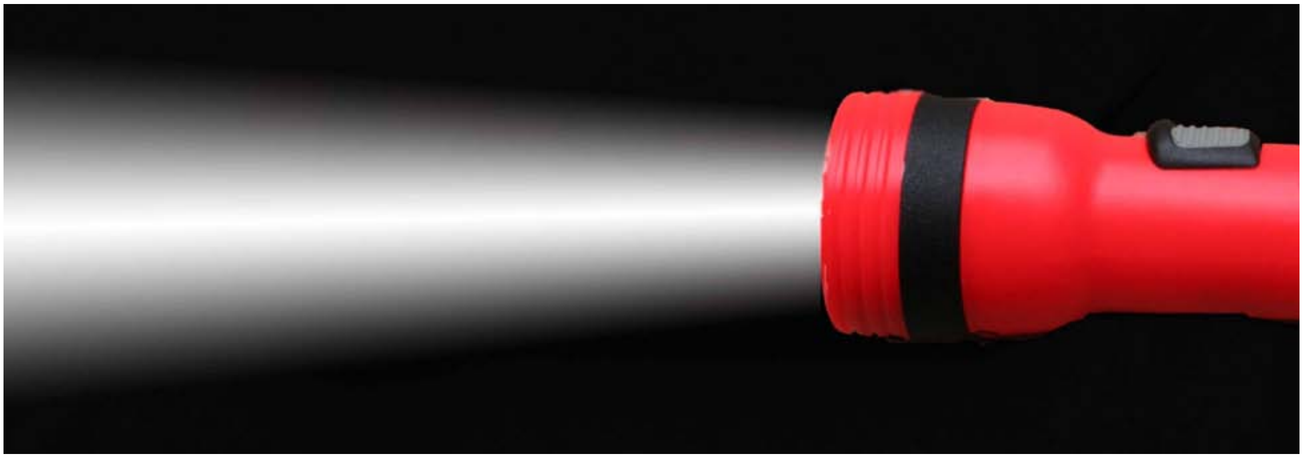

6 Make The Move To The Cloud

Cloud-based DR provisioning allows you to dispense with your hardware altogether, as well as the provisioning of your own power, cooling, redundancy, compliance auditing, security protection, data replication and back-up systems.

Maintaining your own DR site versus using cloud services is a plan implementation choice and not actually a plan factor. Either way, all the preparation, collaboration, workload prioritization, RTO/RPO determination and testing must still be done. There are, however, distinct advantages to be gained by using cloud services, including improving the likelihood of a successful recovery failover and failback at a lower cost. It seems to be catching on.

Organizations are three times more likely to adopt cloud-based provisioning of recovery sites today than three years ago. Also gaining ground is colocation, which entails deploying your DR systems at someone else's site using their infrastructure; that has more than doubled. (6) page 5

Cloud-based DR provisioning allows you to dispense with your hardware altogether, as well as the provisioning of your own power, cooling, redundancy, compliance auditing, security protection, data replication and back-up systems.

Two additional benefits from doing DR in the cloud – using Peak 10's Recovery Cloud, for example – versus doing this yourself are that we test the infrastructure with each customer twice a year, and we are running tests routinely on a weekly basis. We know what to look for and how to respond. All customers collectively benefit from what we do for each individually. It also means that customers know that they are ready when a disaster declaration is made.

A Disaster Recovery-as-a-Service (DRaaS) provider such as Peak 10, assumes complete responsibility for ensuring your computing and networking operations function as intended. Recovery will still be only as good as the quality of your plan and its alignment with your business operations in general. However, relieving yourself of the demands of site, systems and infrastructure management, staffing and testing frees you to focus on the business and human sides of DR/BC implementation.

(1)(2)(4)(6) Forrester Inc., “The State of Business Technology Resiliency, Q2 2014,” 12 May 2014, Stephanie Balaouras

(3)(5) Gartner, Inc., “Improve Your IT Disaster Recovery Plan and Your Ability to Recover From Disaster,” 8 January 2014, Robert Neagle

Contact Peak 10 Today

866.473.2510

www.peak 10.com

peak 10™

IT INFRASTUCTURE | CLOUD | MANAGED SERVICES