GIGAOM RESEARCH

Why business-continuity and disaster-recovery processes fail, and how to cure them effectively

Marc Staimer

This report is underwritten by CoreSite.

Table of Contents

Executive summary

Business continuity (BC) and disaster recovery (DR) processes have been around since the 1970s, when IT managers suddenly realized that their IT systems and data were no longer a luxury, but rather a critical necessity for running their businesses. Most large IT organizations have some sort of BC/DR plan. Shockingly, most small and medium businesses (SMB) do not. The primary SMB rationale is that they perceive a high cost versus a relative risk. Many incorrectly assume their file backup, snapshot, or image-replication technologies are adequate.[1] The SMB attitude toward BC/DR can be summed up as "It's a cost center" or "It's pricey insurance."

Large IT organizations also see BC/DR as a cost center or insurance. However, because of the myriad supported systems, IT has to respond to and recover more frequent system outages, failures, and disasters. IT veterans *know* that BC/DR plans and processes are not simply a costly indulgence, but rather an essential job requirement. They must constantly balance cost and the ability to operate in the event of all kinds of disasters, and they realize that failure to prepare is a recipe for catastrophe.

They prepare hoping they never have to exercise their plans. That hope generally springs from a deep-seated anxiety and fear that their BC/DR processes will not work when they need them most. This anxiety's root cause is directly correlated to incomplete BC/DR testing, nonexistent periodic BC/DR systematic testing, or BC/DR test failures. It's also from IT pros' experience with Murphy's Law, or more specifically, from a corollary of Murphy's Law: "Whatever can go wrong will go wrong at the worst possible time, causing the most possible damage."

Why? How complicated can it be? IT professionals have been solving more difficult and convoluted problems for decades. Why do BC/DR processes continue to vex so many IT organizations, especially when there is a wide breadth of new technologies and services — such as "backup as a service" (BaaS) and "disaster recovery as a service" (DRaaS) — that make BC/DR so much simpler? That is what this paper investigates, explaining the ins and outs of BC/DR issues and detailing the methodologies that are proving they can cure this problem.

[1] *(Note: data protection is only part of an effective BC/DR plan, which also includes facilities, people, networks, equipment, access, etc.)*

Rationale behind business continuity and disaster recovery

Disaster recovery is a series of measures and techniques utilized to ensure mission-critical or essential IT-systems operations recover when those systems become unavailable for an unacceptable period – when, for example, a natural or human disaster takes out those essential systems. Business continuity goes much further than DR (and at significant expense), by making sure those critical IT-system business functions keep running without disruption during and after a major incident or disaster.

A natural disaster can be caused by a variety of troubling possibilities, including:

- Weather initiated incidents (hurricanes, tornados, lightning strikes, hail, wild fire, flood, or solar flare)
- Geological incidents (earthquakes, volcanic eruptions, tsunamis, or sink holes)

More commonly, disasters are human derived. They include a broad variety:

- Hardware failures (power, cooling, components, hard-drive head crashes, mother boards, adapters, cables, transceivers, media defects or failure)
- Software failures (data corruption, write errors, read errors, or data loss)
- Network failures (network equipment, network adapters, routers, switches, network providers, or network storms)
- Human error (most common)
- Hazardous chemical or waste spill
- Loss of key/critical personnel (illness, accident, family emergency, even death)
- Human malevolence (malware, theft, vandalism, terrorism, arson, bombing, disgruntled employee, compromised passwords, even espionage)
- Legal or legislative changes

The potential for disaster is statistically high. What makes that potential scary is that disasters are not a matter of if, but rather when, how often, and how bad. Disaster anticipation and planning will keep CIOs awake at night because it is their fiduciary responsibility to get essential IT operations up and running as quickly as possible after a disaster. Failure to do so will end in the very real possibility of bankruptcy, or even liquidation. The United States Federal Emergency Management Agency's (FEMA) statistics reveal disturbing post-disaster business survival information. Up to 40 percent of businesses do not reopen after a disaster, and an additional 25 percent (65 percent total) fail within one year. The United States Small Business Administration's (SBA) records show that more than 90 percent of businesses fail within two years after being felled by a disaster. The main reason for these failures is not being able to bring critical business operations back online in a reasonable amount of time. Historical analysis indicates that those organizations that fail to recover their critical applications, data, and operations within two weeks of a major disaster or disruption have almost a 95 percent probability of being out of business within two years. These are daunting numbers.

The only way to not become one of those statistics is to plan, prepare, test, correct, continuously improve, and update the organization's BC/DR plans. The reasoning behind multiple plans is that the processes and procedures will vary for different kinds of disasters, locations, applications, regions, operations, and prioritizations. And yet Dragon Slayer Consulting's interviews with the leading BC/DR vendors revealed some startling information. Most of the user BC/DR tests failed or had problems severe enough to prevent a timely recovery. And even those tests that did succeed had significant problems and errors that would negatively impact operations.

Suggesting a cure to this problem first requires understanding its root cause.

Standard processes for BC/DR

DR and BC are typically entwined. Planning requires knowledge of which business-critical applications and operations cannot go down, which ones must be brought up in the shortest time possible, and which ones have a higher tolerance for downtime. That knowledge is imperative to prioritization and effective BC/DR planning. The planning falls apart when it is not updated as IT changes. IT organizations are not static, they're dynamic. The BC/DR plan must evolve to stay highly correlated to the actual IT ecosystem and priorities as they change. And they do change – frequently. That planning starts with a series of fundamental questions.

BC/DR planning questions

These questions are vital to the BC/DR plan, must be answered upfront, and updated frequently.

- Where will the BC/DR data center be located? In other words, where will these IT systems be brought up?
    - Local data center
    - Remote data center
    - Disaster-recovery-provider data center
    - Mutual-bilateral-cooperative-partner data center
    - Co-location or infrastructure-as-a-service (IaaS) provider such as Amazon Web Services, CoreSite, Google Microsoft Azure, Rackspace, and many others
- What service levels will the equipment (servers, networks, storage, infrastructure, etc.) in the DR data center support?
- How will the applications and data be moved to the BC/DR data center?
    - Backup on tapes shipped
    - Backup on hard disk drives shipped
    - Backup on optical disks shipped

    - Backup to cloud BaaS provider over internet
    - Backup to cloud DRaaS provider over internet
    - Backed up virtual machines (VMs) replicated over internet
        - For physical machines, physical-to-virtual (P2V) conversion to VMs
    - VMs replicated over internet
    - Hybrid-cloud VM infrastructure between primary and BC/DR data centers
    - Storage systems shipped
    - Storage systems replicated or mirrored over internet
    - Hypervisor snapshot with VM movement over internet
    - Hypervisor storage image migrated/replicated over internet
- How much bandwidth will be available for the recovered IT systems?
- How will key personnel get access to the site, all of the necessary systems, and infrastructure?
    - Will the personnel be able to do what they need to without being on site in person?
    - If so, what are the steps they must take to do so?
- Who will be there to help when things don't work correctly?
- What are the procedures for the users to access these systems remotely?
- After everything is up and running and the disaster is over, what are the processes and procedures to failback to the primary or new data center?
    - If disruptive, how can it be minimized?
- What are the testing processes and procedures?
- How often will the BC/DR plan be tested?
    - Partial and full test, including failback?

BC/DR pitfalls range from trivial to crushing

Incomplete RTO and RPO tiering assignments

The most common incomplete RTO and RPO tiering assignments are incomplete recovery time objectives (RTO – the time required to recover the application and its data, and be back up in operation) and recovery point objectives (RPO – the amount of data loss that can be tolerated based on a specific time period) by application and system will wreck any BC/DR plan.

While most application owners and C-level executives might prefer RTO and RPO to be effectively "0" for all applications, that's financially impractical. Data protection follows the 90:10 rule. Cutting RPOs to protect the last 10 percent of the data or cutting the last 10 percent of the time from the RTO equals 90 percent of the cost. Budget constraints require a tiering or matrix of RTOs and RPOs based on mission criticality for the organization. Administrators and department heads will likely argue about where they fit within that matrix.

Inadequate BC/DR data-center hardware, infrastructure, service levels, support, and/or bandwidth

It's not uncommon for a BC/DR data center to have adequate bandwidth to receive replicated VMs, data, backups, etc. But when the disaster hits and operations are switched entirely to the BC/DR data center, will it be enough? Especially during those early peak demands? Some IT organizations will provide less bandwidth at the BC/DR data center than their primary one to save money. The rationale being that since they are operating in BC/DR mode, reduced performance is acceptable. Operating after a disaster hits is a bit more chaotic, and most forget that the much-reduced performance was an effort to save money. That same logic is far too often applied to the BC/DR servers, storage, and infrastructure. The result is highly aggravating when operations resume in the BC/DR data center. Lots of wasted time and cycles are spent attempting to tune the inadequate ecosystem.

This can be especially vexing when utilizing someone else's data center that is providing just a small portion of it for their obligation to provide BC/DR. Universities do this a lot for each other. Although typically perceived as an inexpensive BC/DR, it's also understaffed, undersupported, underequipped, underbandwidthed, and likely to fail. Not a good idea.

A shared BC/DR data center is likely to have several – even many – other users. This can and will have substantial consequences during a regional disaster when everyone is attempting to move to the same BC/DR data center at the same time. Inadequate hardware, infrastructure, services, support, and bandwidth will rear its head quickly and in a catastrophic manner.

Hardware and software at BC/DR data center not kept in sync

Hardware and software that are not kept in sync are more of a problem when recovering physical than virtual environments. Differences in hardware, software, microcode, drivers, even bios can mean failure in bringing up operations at the BC/DR data center.

This problem rears its head when IT shops try to bring up their systems in a bare metal environment. Bare metal recoveries are never clean. Fortunately, P2V conversions have made this less of a necessity.

Doesn't failback cleanly

Assuming the failover to the BC/DR works well (and that's a big assumption), what about failback? Many BC/DR tests focus only on the failover. Yet the failback is equally important. After the disaster is over and the primary data center is restored, operations must shift back to that primary data center. Those who actually test their BC/DR plans frequently don't test failback. Then when it's time to failback, it doesn't. Or it has significant, time-consuming, disruptive problems. Clean failback is essential for minimal operational disruptions.

BC/DR data center not out of region

Out of region means being farther than 250 miles or 400 kilometers distant. Not being out of region greatly increases the risk of a disaster taking out both the primary data center and the BC/DR data center. For example, Hurricane Sandy took out data centers in all five boroughs of New York City, Long Island, and New Jersey. If the BC/DR data center was within the disaster region, there's a good chance it was not functioning when it was needed.

Natural disasters do not care about a BC/DR data center location. Nor do terrorist acts. If the wind on 9/11 had been blowing the Twin Towers' smoke, debris, and human remains west by southwest, the ability of many IT shops to recover would have been a lot lower since many of their BC/DR data centers were just across the river in New Jersey. Most companies choose a close-by BC/DR data center location because of convenience, because making hardware changes to the BC/DR facility is easier and cheaper when it is within driving distance. Too many IT managers do not think about being out of region until too late.

Testing fails or is insufficient

As previously discussed, BC/DR testing rarely produces complete success. There is almost always something that goes wrong. The good news is that testing reveals the problems and they can be fixed. The bad news is that few organizations test frequently enough and many never test at all.

BC/DR tests should take place once a quarter, minimally twice a year. Testing should not be disruptive to operations, should be a good representative sample of an actual major disruptive event, and be relatively simple to execute. It should also be well documented. Every error and failure should be noted and corrected before the next test.

Reality tends to diverge from the "should" column. Most IT organizations test less than once a year. Tests are ordinarily not representative of a major event. Failover and failback are difficult and error prone. Errors and failures are rarely completely resolved.

Cost

BC/DR is and always will be a cost issue because it does not produce revenue. It's only utilized when there is an infrequent disaster or major outage. BC/DR is insurance against partial or complete operations failure in the event of a disaster. It insures that revenue streams are not disrupted and productivity remains high. The strategic issue comes from the organization's risk tolerance, or how much insurance it requires and at what cost it's willing to pay. Insufficient coverage will lead to disaster. Too much coverage that may never be used is a money pit with no apparent return (unless there is a disaster.) BC/DR is a cost center until there is a disaster. Therefore, BC/DR return on investment is incalculable.

Overcoming BC/DR pitfalls

New BC/DR technologies

Several new powerful technologies are transforming BC/DR. Server virtualization has by itself immensely simplified BC/DR. Hypervisors boot VMs off of disk or file system images. Many of the hypervisors, such as VMware vSphere, Microsoft Hyper-V, Citrix XenServer, KVM, and even Solaris VM, make BC/DR easier by permitting VMs to be replicated between data centers with low RPOs and fast RTOs. Server cloud stacks are based on virtualization technologies as well; meaning stacks such as vCloud Director, CloudStack, and OpenStack Nova have similar capabilities. Those hypervisors and cloud variations have enabled efficient VM replication from backup software, VM2VM replication software, and storage2storage replication software (found in most primary storage systems and software-defined storage). Much of this software also has the ability to convert a physical server image into a VM for BC/DR purposes. In the event of a disaster, all VMs (including physical conversions) are mounted and made live in minutes. This near-instantaneous recovery solves the vast RTO problem.

Local and remote VMs are kept in sync via RPO policies. VMs are made live at the BC/DR data center with the flip of a switch. After the disaster is over, the VMs are replicated back to the primary data center. Solving the small RPO issue requires the data-protection software deliver frequent zero-capacity snapshots, or continuous data protection (copies every write or change, then time stamps each copy), or mirrors each write or change.

One deliberately designed side effect (benefit) of these newer BC/DR technologies is their ability to reduce the total cost of BC/DR. It does so by reducing the power, cooling, and number of servers, adapters, switch ports, cables, transceivers, racks, and floor tiles.

These new powerful BC/DR technologies have incubated an entirely new breed of BC/DR service providers, BaaS and DRaaS. And they've done so at price points an order of magnitude lower (1/10th) than just five years ago. BaaS provides local backup and recovery for the most recent data and remote backup and recovery for older data, with knowledgeable and experienced personnel managing the backups. Recoveries are self-serve with expert help when required. DRaaS goes a substantial step further by allowing servers, applications, and data to be recovered in the service provider's data center, allowing operations to be run there until the disaster terminates and operations can be moved back to the primary data center. The DRaaS provides complete BC/DR methodology, local BC/DR recoveries, remote BC/DR recoveries at the BC/DR data center or multiple data centers, negotiated service-level agreements with specific RTOs and RPOs, and a high degree of service and support.

Whether BC/DR is do-it-yourself (DIY) or via a DRaaS provider with these technologies, four indispensable factors must be considered to ensure success.

1. The first is the capability of easy, intuitive, and simple online BC/DR testing of any application and data amount without operational disruption.

2. Second, same criteria for failback. These two capabilities must be in the data-protection software exploited by DIY BC/DR or provided by the service provider with BaaS and DRaaS.

3. Third, application and data health are constantly monitored and corrected if errors are found, to make sure the BC/DR is always working with sound data. This capability must be part of the BC/DR software.

4. The fourth factor is the BC/DR data center. This cannot be emphasized enough and it is repeatedly the least scrutinized part of the BC/DR plan.

BC/DR data centers

Some of the BC/DR pitfalls are previously described, such as not being out of region, not having enough server, storage, or network assets, not having enough data center bandwidth, no hardware/software elasticity. For a BC/DR data center to not be the source of the pitfalls and adequately deliver BC/DR, it must provide:

- **Physical security**. Meaning no one can just physically walk off with the customer's data, hard drives, assets, or more. This is commonsense and required for any multitenant environment. Yet as Voltaire so clearly articulated, "Commonsense is not so common."

- **Virtual security.** No one, not even the data center or service provider's employees, can access or even see the customer's data.

- **SSAE-16 Type II[2] in-depth audit of a third-party service organization**. It's imperative that customers have an outside non-biased party that will tell them the services they are buying are the services they are receiving.

[2] *In 2011, SAS-70 was superseded by SSAE-16 but it retains the original purpose for SAS-70 compliance. SSAE-16 Type II audits are conducted by Certified Public Accountants from American Institute of Certified Professional Accountants (AICPAs). The SSAE-16 Type II report and seal is the proof that the service provider's facility, processes, and procedures meet security standards.*

- **Elastic provisioning of servers, storage, networks, and infrastructure**. Charging only for what is utilized when it's utilized keeps BC/DR costs under control. That elasticity does not eliminate up-front costs, however it greatly reduces them by charging a relatively smaller monthly fee for the privilege of access and a much higher fee for actual utilization of the BC/DR assets during a disaster or test. After the disaster, when the assets are no longer being utilized at production rates, the fees drop again.

- **Enough assets and infrastructure to support all clients from a regional disaster**. A natural disaster or terrorist attack will in all likelihood take out more than one customer for a physical BC/DR "availability services" provider.

- **Enough bandwidth to support live production of all clients from a region at once**.

- **Multiple data centers where BC/DR production can be brought up in the event of a disaster, with high bandwidth between those data centers.**

- **Optional high availability to another BC/DR data center out of region**. That means there's at least one data center outside of the local region (more than 250 miles or 400 kilometers). This is the only way to ensure a local natural disaster doesn't take out both data centers. If the primary BC/DR data center is hit by the same disaster as the client there must be an alternative location out of region to recover or the BC/DR event will fail.

- **High bandwidth to/from the BC/DR data centers to other compute-cloud-service-provider data centers.** This enables application workloads and data running in those data centers to be recovered in the event of a catastrophic disaster where they're currently running. For example, there are several large cloud service providers with primary data centers in the Seattle region. Mount Rainier is also in that region. Mount Rainier is a volcano that is geologically overdue for an eruption. When it erupts it is quite likely to cause a major disruption to all data centers in the area because it has the potential to be up to an order of magnitude more explosive than Mount St. Helena was in 1980.

- **Expert on-site service and support for all facilities**. Self-service is a wonderful thing when there is all the time in the world. Reacting to a disaster or major disruption is not one of those times. Help is needed because time is the enemy.

These requirements are essential for DIY BC/DR data centers, but they are just as important for BaaS and DRaaS providers. Many BaaS and DRaaS providers take advantage of third-party data centers for their services. That makes perfect business sense because data-center management is not their business. BC/DR is their business. IaaS, co-locations, and cloud-compute providers specialize in efficient, cost effective, secure, and elastic data centers. That is their business. By leveraging that business the BaaS and DRaaS provider is able to provide an outstanding service with nominal capex investment. The IaaS, co-locations, cloud-compute pay-per-use elastic cost model parallels and correlates closely to the BaaS and DRaaS pay-per-use elastic income model. That reduces risk and obtains profitability earlier and more consistently.

But should a BaaS and/or DRaaS provider become the IT organization's supplier of BC/DR, their data center becomes a mission-critical aspect of the service. Whether the data center(s) are provided by the BaaS and/or DRaaS provider, it should meet the requirements stated above. If they don't, then a failure is much more likely.

The promise of BC/DR hybrid clouds

A hybrid cloud is a local private cloud that's seamlessly integrated and interconnected with a public cloud or a remote private cloud.[3] Hybrid clouds enable local workloads to generally use the remote cloud for overflow during peak demand. Yet, of all the applications or workloads that make sense for hybrid clouds, none are more suited or more natural than BC/DR.

BC/DR is a simple, straightforward exercise in a hybrid cloud. Replicating VMs and physical machines between clouds is based on well-known, easy data-protection processes with many technology options, including storage replication, server-to-server replication, VM replication, P2V conversion and replication, backup of physical and virtual machines then mounted, and more. All of them achieve comparable levels of BC. If a machine, or multiple machines in the local cloud, or the entire site goes down it merely takes minutes to bring up the VMs or the site live and operational in the remote cloud.

Most BaaS and DRaaS providers take advantage of hybrid clouds in the way in which they deliver their services. They utilize a variety of the previously mentioned data-protection processes. The most recent copies of virtual or physical machines and their data are kept locally. This is to enable BC/DR as fast as possible when a machine fails, data is corrupted, malware hits, etc. That recent data copy and all previous copies (typically deduped and compressed) are also maintained in the remote cloud.

Perhaps the biggest advantage of BC/DR hybrid clouds is their affordable cost. Many costs are deferred until or if a disaster occurs. This reduces the cost of the BC/DR "insurance" without reducing any of the RTO and RPO capabilities unless a disaster occurs. Cost is rarely a factor at that time, however all of the costs are pre-set so there are no surprises, thus providing the best of both worlds.

Not all BC/DR hybrid clouds are created equal

There are a range of BC/DR hybrid clouds with a wide degree of capabilities. Some provide exactly what's needed whereas others fall short of expectations. They'll fall short in simplicity or they'll fall short in meeting RTO or RPO expectations. Some of these shortfalls can be quickly corrected by simply changing to the more modern version of data-protection technology discussed above. The area that is not so easily corrected is when the BC/DR hybrid-cloud remote data center fails to meet requirements.

---

[3] *"Cloud" is the market's shorthand when referring to cloud computing infrastructure. It is a marketing term that is derived from network diagrams utilizing a cloud symbol for wide area networks. The meaning has expanded to include software, platforms, database, backup, DR, infrastructure as a service, and Co-location (Co-Lo). Cloud service must be elastic (expand and contract) on demand, with minimal to zero human intervention, and a pay-per-use subscription structure.*

Just as the private or public BC/DR data centers are a source of many BC/DR pitfalls, so too are the hybrid-cloud BC/DR data centers. Far too many hybrid-cloud BC/DR data centers fail to pass muster, often relying on the confusion or "mystery" of the cloud to get by. Just because the BC/DR is in a remote or public cloud does not obviate the BC/DR data-center requirements.

BC/DR hybrid-cloud data center requirements

The requirements for a BC/DR hybrid cloud include the same requirements as any remote BC/DR cloud. The hybrid-cloud requirements include physical security; virtual security; SSAE-16 Type II in-depth audit of a third-party service organization; elastic provisioning of servers, storage, networks, and infrastructure; enough assets and infrastructure to support all clients from a regional disaster; enough bandwidth to support live production of all clients from a region at once; multiple data centers with high bandwidth between them; optional high availability to another hybrid-cloud BC/DR data center out of region; high bandwidth to and from the hybrid-cloud BC/DR data centers to other compute-cloud-service-provider data centers; and expert onsite service and support for all facilities. BC/DR hybrid clouds in co-locations, IaaS, or public-cloud data centers have non-trivial additional requirements such as:

- **Multi-tenancy:** Multi-tenancy is not simply an option but rather an absolute requirement for any IT organization placing their data in someone else's data center. Multi-tenancy ensures that no other provider customer/client, employee, or manager can access another's data. It demands industrial-strength encryption in flight and at rest with the encryption keys controlled by the customer/client. This level of multi-tenancy is a must for physical machines, bare metal containerized services, and VMs.

- **Pay per use:** Also known as chargeback or arrears-based subscription pricing. Users are charged for what they use, when they use it, and it's clearly documented. No surprises.

- **Multiple levels of service:** One size does not fit all. Several flexible, customer/client-selectable options are needed. Large organizations will frequently have different requirements than smaller ones. Mission-critical applications will also have different BC/DR service-level requirements than non-mission-critical applications. Multiple service levels allow flexible alignment of requirements with budgetary considerations.

- **Service-level agreements (SLAs):** SLAs that are clearly and contractually defined, easily monitored by the customer/client and the provider, with penalties clearly and specifically spelled out for SLA failures.
- **System and data portability:** If there is a change of circumstance and the customer/client wishes to change service providers or bring their BC/DR hybrid cloud back in house, it should be just as easy as moving to the current service provider.

Summary and key takeaways

BC/DR processes are a requirement. These processes have a jaded history of breaking down or failing at the worst possible time: when a disaster strikes. BC/DR testing has historically rarely been consistently successful. Many of the reasons for these breakdowns are human-error based and can only be resolved by internal policies, discipline, and culture. Common human-error breakdowns include: incomplete RTO and RPO determinations per server or VM, application, data, and storage system; BC/DR plans not kept current as the IT organization is constantly evolving and changing; negotiated SLAs that have no basis in reality and can't be met by implemented data-protection technology; no consistent representational systemic BC/DR testing, documentation, error detections, and corrections.

Other failures can be traced to past generations of technology. New technologies such as server virtualization permit faster, more efficient, simpler, less costly replication, failover, and failback. That in turn has spurred innovation in creative new data-protection software. This data-protection software leverages server virtualization's potential to solve previous BC/DR problems with RTOs, RPOs, SLAs, failover, failback, cost, and even testing. These new data-protection technologies have generated whole new markets in BaaS and DRaaS. BaaS and DRaaS providers can offer BC/DR at much more cost-effective rates than ever before (less than 10 percent of rates as recently as five years ago).

Third-party data centers such as IaaS, co-locations, public-cloud compute, and co-location hybrid-cloud arrangements can solve several other BC/DR issues. For regional disasters they commonly provide out-of-region data centers. For unexpected demands, some provide plenty of space, power, cooling, bandwidth, and assets to cover the most egregious disaster situations, and they are elastic, charging only for what's used on a pay-per-use basis. Most are physically and virtually secure for multi-tenants. The hybrid-cloud data centers hold the most promise at the most efficient cost when it comes to BC/DR.

Key takeaways

Solving BC/DR failures cost effectively requires three things:

1. Internal culture of disciplined processes.
2. Server virtualization or cloud stack with VM-application mobility capabilities, plus the data-protection software/systems that leverage those capabilities for BC/DR.
3. Third-party data centers (IaaS, co-locations, cloud compute, and hybrid-cloud compute) that meet the right requirements.

About Marc Staimer

Marc Staimer is the founder, senior analyst, and CDS of Dragon Slayer Consulting in Beaverton, OR. The consulting practice of more than 16 years has focused in the areas of strategic planning, product development, and market development. With over 34 years of marketing, sales and business experience in infrastructure, storage, server, software, databases, and virtualization, he's considered one of the industry's leading experts. Marc can be reached at marcstaimer@me.com.

About Gigaom Research

Gigaom Research gives you insider access to expert industry insights on emerging markets. Focused on delivering highly relevant and timely research to the people who need it most, our analysis, reports, and original research come from the most respected voices in the industry. Whether you're beginning to learn about a new market or are an industry insider, Gigaom Research addresses the need for relevant, illuminating insights into the industry's most dynamic markets.

Visit us at: research.gigaom.com.