DRaaS for Dummies

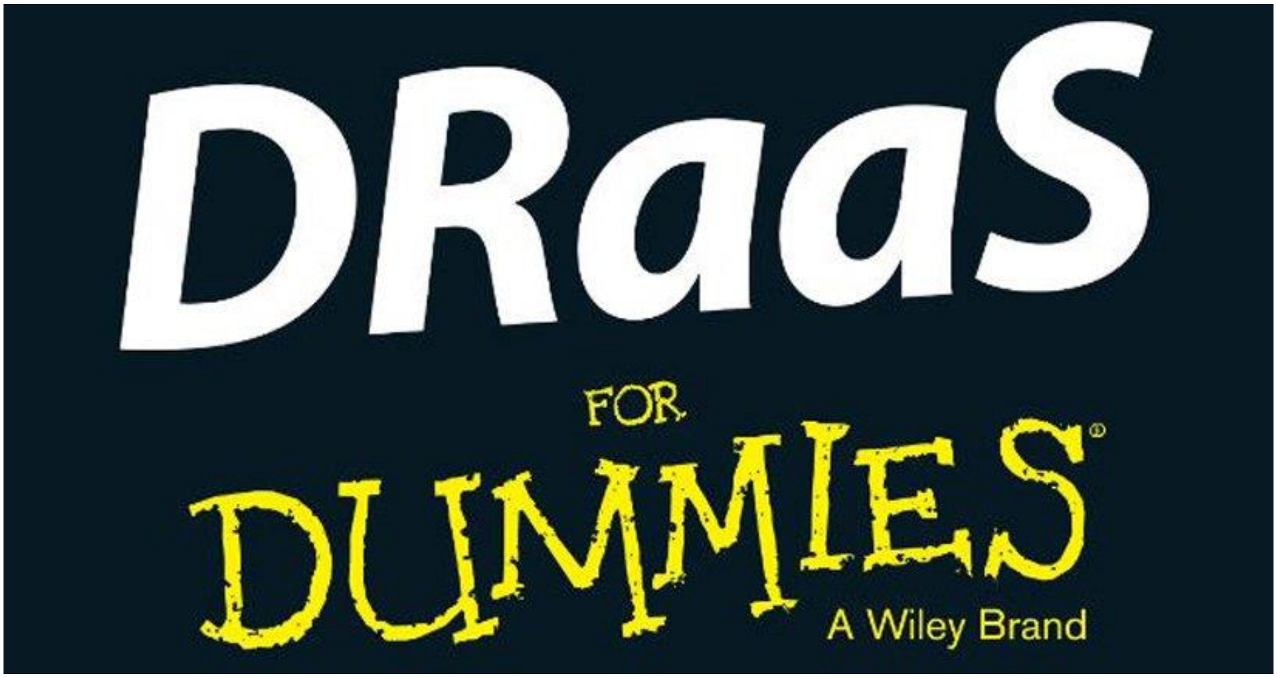

March 11, 2018 / Evolve IP

Disaster recovery (DR) planning has a reputation for being difficult and time-consuming. Setting up alternate processing sites, procuring hardware, establishing data replication, and failover testing has been incredibly expensive undertakings. To top it all off, the need for 24x7x365 business application availability threatens to make disaster recovery planning an exercise in futility.

Disaster Recovery as a Service, or DRaaS, is turning the DR business on its head. The responsibility for all of the gritty details one used to have to juggle in order to ensure that every system, file, database record, and network element was duplicated at an alternate processing site can now be passed onto a trusted service provider. A face — not just an interface.

DRaaS is bringing true DR capabilities to an entirely new pool of organizations — folks who previously considered DR to be out of their reach. Today, DRaaS makes setting up DR almost as easy as setting up a new smartphone — seriously. People who set up DR using DRaaS are amazed at how much knowledge they aren't required to have. This makes DRaaS available and attractive to an even larger audience. *DRaaS For Dummies lays* the foundation for this new approach to DR. After reading this book, you'll have a new appreciation for DR professionals and how difficult it used to be for them.

Chapter 1: Understanding DRaaS

Disaster recovery (DR, for short) is the undertaking whereby an organization invests in computing hardware and software to be used in the event that a disaster renders the primary processing site unavailable. That's about as simply as I can describe it, but in reality, it is far more complex than that.

People in today's always-on, always connected world are far less forgiving of unscheduled downtime that occurs, regardless of the reason. In this world, application availability is king. Not that long ago, people tolerated applications being down for hours or even days at a time (considering the circumstances, of course), but today even a fraction of an hour is considered inexcusable. We want our application and we want it now!

In this chapter, you'll get a chance to take a look at application availability expectations, the mechanisms that comprise DRaaS, and the benefits that organizations can enjoy with DRaaS solutions.

Today's Disaster Recovery Practices

It's not easy being a CIO today. CIOs are under pressure to make their applications and data available continuously, without regard for the "stuff" that happens: Hardware failures, software bugs, data corruption, and disasters. In today's point-and-click world, business users think it's easy for an IT organization to create a fault-tolerant, disaster-proof environment. CIOs and others in IT know it's anything but.

Enter disaster recovery — DR for those of you who love acronyms. Traditional approaches to DR include hot site, cold site, and warm site, discussed here:

Hot site

In a hot site approach, the organization duplicates its entire environment as the basis of its DR strategy — an approach which, as you'd expect, costs a lot in terms of investment and upkeep. Even with data duplication, keeping hot site servers and other components in sync is time-consuming.

A typical hot site consists of servers, storage systems, and network infrastructure that together comprise a logical duplication of the main processing site. Servers and other components are maintained and kept at the same release and patch level as their primary counterparts. Data at the primary site is usually replicated over a WAN link to the hot site. Failover may be automatic or manual, depending on business requirements and available resources.

Organizations can run their sites in "active-active" or "active-passive" mode. In active-active mode, applications at primary and recovery sites are live all the time, and data is replicated bi-directionally so that all databases are in sync. In active- passive mode, one site acts as primary, and data is replicated to the passive standby sites.

Cold site

Effectively a non-plan, the cold site approach proposes that, after a disaster occurs, the organization sends backup media to an empty facility, in hopes that the new computers they purchase arrive in time and can support their applications and data. This is a desperate effort guaranteed to take days if not weeks.

I don’t want to give you the impression that cold sites are bad for this reason. Based on an organization’s recoverability needs, some applications may appropriately be recovered to cold sites.

Another reason that organizations opt for cold sites is that they are effectively betting that a disaster is not going to occur, and thus investment is unnecessary. I don’t think this is a smart move.

Warm site

With a warm site approach, the organization essentially takes the middle road between the expensive hot site and the empty cold site. Perhaps there are servers in the warm site, but they might not be current. It takes a lot longer (typically a few days or more) to recover an application to a warm site than a hot site, but it’s also a lot less expensive.

Comparing hot, warm, and cold

The trouble with all of these hot-warm-cold approaches is that they do not meet today’s demands for cost effective and agile recovery. Users typically expect applications to be running within a fraction of an hour. Engineered correctly, a hot site can meet this demand, but at spectacular cost. Warm and cold sites don’t even come close.

It should not come as a surprise to you that most organizations “go commando” with regards to their DR plans. They have little or nothing in the way of policies, procedures, or technologies that enable the recovery of critical systems at any speed. This is understandable, as rapid recovery capabilities have historically been so expensive that only the largest organizations could afford them.

Backing Up Your Data

Data backup is an essential part of sound IT management. We all know that things occasionally go wrong in IT, and data loss is a result that no one will tolerate.

Better organizations employ the 3-2-1 rule when it comes to backing up data. Here is how the rule works:

✓ Keep 3 copies of data: 1 primary, 2 backups

✓ Use 2 different types of media

✓ Keep 1 set in the cloud in DRaaS or BaaS (backup as a service)

Introducing Disaster Recovery As A Service

Since the early 2000's, many types of service providers have emerged and built entire industries that reduce the cost and complexity of many classes of technology. For instance, Software as a Service (SaaS), Infrastructure as a Service (IaaS), and Platform as a Service (PaaS) have created entirely new paradigms for businesses' use of technology.

Disaster Recovery as a Service (DRaaS) is a rapidly growing cloud-based service that makes it easy for organizations to set up alternate processing sites for disaster recovery purposes. Like other "as a service" offerings, advanced software enables DRaaS to simplify the entire process for organizations of any size as well as the service providers that offer this service.

DRaaS is important because it represents an innovative and less costly way to back up critical data and quickly recover critical systems after a disaster. DRaaS does this by leveraging cloud-based resources that provide infrastructure that is far less expensive than on-premise systems due to the ability to scale and share cloud resources

To meet the growing demand for software resilience, DraaS has brought simplification and reduced costs to organizations that are serious about implementing DR. With DRaaS, an organization can implement a high-performing DR solution for its critical systems but without any of the complexities. Like other "as a service" providers, DRaaS providers take care of the back-end complexity for their customers and provide a simple user interface for setting up and managing a DR solution.

There are a few different flavors of DRaaS discussed here, related to whether your organization uses public or private cloud:

Public cloud DRaaS

Organizations can implement DRaaS using a public cloud infrastructure. Any public cloud service that meets an organization's security and operational requirements can be used. A typical DRaaS solution will employ customer-managed soft-ware for setting up and controlling cloud-based DR resources.

While their vast scalability provides many cost advantages, going with a public cloud infrastructure means you will likely be foregoing a personal one-on-one relationship. If something goes seriously wrong — get ready to stand in line, if you can find a line to stand in!

Private cloud DRaaS

Organizations with their own data centers and private cloud infrastructure can definitely utilize DRaaS solutions. The software components that comprise DRaaS solutions can be installed on an organization's own server infrastructure. In these types of situations, the HQ datacenter in essence takes on the role of the service provider for their different business locations. These solutions will reduce the effort and complexity of data backup and replication mechanisms for organizations that are required to keep data under their direct physical control.

Managed cloud DRaaS

Organizations using managed cloud services can include DRaaS solutions to their service portfolio. Managed cloud service providers can include DRaaS as a part of a standard, hands-free offering that takes care of data backup and data replication details. This permits customers to concentrate on their software applications and other hosted components.

The Role and Need for Secondary Sites

One of the time-honored (and still valid) principles of disaster recovery planning states that a secondary computing location be established. The reasons for this include:

✓ The primary site may be incapacitated because of the effects of a regional disaster. This includes events such as an earthquake, hurricane, or flood.

✓ The primary site may be incapacitated by the effects of a localized event, such as a fire, landslide, power failure, communications outage, or a water main break.

✓ The primary site may have suffered an equipment failure in its IT infrastructure, or an operational error resulting in unexpected and perhaps prolonged downtime.

The best bet for covering all of these scenarios is the use of an alternate processing center some distance away from the primary site, generally 100 miles or greater, depending on the types of disasters that can happen in your part of the world. This helps to ensure that the alternate processing site is not affected by whatever regional event has affected the primary site. This approach is still valid with cloud services. With a cloud-hosting provider, you’ll generally have a choice on where your recovery servers will reside. What you don’t want to end up with is a situation where the DR servers assigned to you are in the same city as your primary site. This would not result in a good recovery scenario, since the hosting provider may be adversely affected by the same disaster that affects your primary site.

Using a cloud-based hosting provider is a cost-effective way to build a secondary site. The main advantage is the preservation of capital. Virtually no investment in recovery systems is required since they are instead leased from the service provider if and when they are needed.

Backup and Replication

An essential part of a disaster recovery plan is some means for transporting copies of mission-critical data away from the primary processing site to another location that will not be affected by whatever event affected the primary site. There are two main ways to copy data:

✓ **Backup.** Data is copied from databases, flat files, and virtual machine images to backup media residing on disc-based storage, but could also include backup to magnetic tape or virtual tape libraries.

✓ **Replication**. As data is being written to databases and flat files, that same data is being transmitted over a net-work to another storage system, usually to an alternate processing center or cloud provider.

The main distinction between backup and replication is this: Backup copies the entirety of a machine image, files, or databases (or the incremental changes since the last backup), in a one-time operation that is then repeated periodically; whereas replication is the continuous or near-continuous transference of updated disk blocks — say, batch updates every five minutes.

Backup was once considered “good enough” for disaster recovery purposes. However, good enough implied that an organization was willing to wait days to recover their systems and get them running again. However, in today’s always-on enterprises, backup is no longer good enough: Backup and replication together are necessary for organizations of all sizes to get its critical applications up and running in 15 minutes or less.

The right strategy for today’s DR needs, then, requires both backup and replication: Frequent backup of virtual machine images and the replication of critical data. Together, these provide system recovery synergy that facilitates rapid resto-ration of critical systems.

DRaaS and Virtualization

Virtualization — the technology that permits multiple operating system instances to run on each physical server — has freed up IT infrastructure and facilitated the revolution that is the mass migration of applications to the cloud. Individual operating systems (which reside within virtual machines) reside in “images”, which are large flat files that can be copied to a DR site for rapid recovery of servers. The power of virtualization and virtual machine management have contributed significantly to the power of DRaaS.

Some DRaaS solutions have the ability to provide advanced, imaged-based virtual machine (VM) replication, which can be used to send VM images to a cloud service provider. Service providers can provide virtual cloud hosts; recovering your server is as easy as booting those hosts from the images sent from your primary site. Better DRaaS solutions include agentless components — meaning there is no software present within individual virtual machines. Instead, you install a module in the virtual environment, which intercepts local disk traffic and sends it to your cloud service provider, where another module receives the traffic and keeps your server VM’s and databases up to date, usually within minutes.

DRaaS Benefits

DRaaS represents the next generation of rapid system data recovery and always-on availability, helping organizations avoid downtime and business disruption without the high costs associated with traditional hot sites.

The low cost and simplicity of DRaaS makes it available to an entirely new class of organizations. The ability to recover applications in the cloud, if and when needed, slashes the cost and complexity of recovery capabilities. Organizations that were on the sidelines, longing for DR capabilities, can now enjoy capabilities that were once reserved for large organizations.

Click here to read the whole DRaaS for Dummies book.