Disaster recovery as a service is growing, but tape won't die

Disaster recovery was once a term associated mostly with the capability to restore data center operations after a physical calamity, but today the loss of IT systems is more often associated with data corruption to include rising malware attacks.

Lucas Mearian (Unknown Publication) | **06 August, 2019 02:05**

Disaster recovery was once associated mostly with the capability to restore data center operations after a physical calamity, but today the loss of IT systems is more often associated with data corruption to include rising malware attacks.

Overall business detections of malware rose 79 percent from 2017 to 2018 -- from about 40 million attack detections worldwide to 72 million -- according to anti-malware security vendor Malwarebytes. While ransomware isn't as significant a threat as it once was, it's still a significant threat and there has been an increase in "focused, sophisticated attacks aimed at businesses," Malwarebyte's 2019 State of Malware report stated.

In fact, ransomware attacks rose by 142 percent from 2017 to 2018 (the latest data available). The only malware with a higher percentage increase year-over-year was backdoor attacks.

"Indeed, the only real spike in numbers has been in the realm of the workplace, with a distinct lack of interest and innovation aimed at consumers," Malwarebytes's report stated. "Malware authors pivoted in the second half of 2018 to target organizations over consumers, recognizing that the bigger payoff was in making victims out of businesses instead of individuals."

Up to half of organizations could not survive if a natural disaster or severe malware attack eliminated or cut access to primary data stores because they don't have a disaster recovery plan in place, according to Al Berman, CEO of the Disaster Recovery Institute (DRI) Foundation. Short of a company standing up a dedicated DR data center – as large financial services companies can afford to do –-- "the low hanging fruit" for most companies large and small is to subscribe to a DRaaS plan, Berman said.

IDC

Unlike backup-only services, which perform replication, DRaaS not only back ups or mirrors data to the cloud, it can also provide standby computing resources to enable rapid systemic recovery by standing up an offsite data center for use while a primary site is restored.

Because there's no need to restore data via the cloud to the primary site – it's simply brought online via remote services – it's possible to achieve near fast recovery point objectives (RPO) and recovery time objectives (RTO). Restoration from a backup service provider can take days or longer depending on the available bandwidth.

"It's just like infrastructure as a service," Berman said. "The idea of mirroring with the cloud is becoming a big opportunity for people because they can reconnect not only to their data but their apps.

"We don't do recovery anymore," Berman said. "Anybody who believes they can be down three or four days while they restore their data is probably a municipality suffering from ransomware."

Essentially, the DRaaS site is a virtual enterprise data center running on a third party's offsite server farms in real time until a natural disaster, malware attack or any other critical interruption of primary data center occurs, at which time the virtual site is activated and becomes an operating data center.

For certain workloads, one of the newest services being offered is DRaaS based on cloud storage from providers such as Amazon AWS, Microsoft Azure or Google Drive. Because a user isn't renting infrastructure from the DRaaS provider, they pay only for the cloud

storage until they need to recover their data, according to Joseph George, vice president of Global Recovery Services at Sungard AS.

Sungard is offering a disaster recovery service that uses AWS's cloud to store data backups for x86 workloads.

"So, you're paying for a base service [from Sungard], but the actual storage and compute you consume is purely based on what you're consuming on top of AWS and their capacity model. And, in some sense, that's nirvana from a DR perspective because… I only pay for what I use," George said. "I don't want to pay for it if it's sitting idle out there throughout the year.

"We're seeing a lot of adoption there and it's where the trend is as well," George continued. "If I look at the next few years, I think there's going to be a lot more x86 workloads recovered via the public cloud."

Sungard is one of the top five DRaaS providers, according to IDC. Others include IBM, Flexential, Carbonite and iLand.

Pat Corcoran, global strategy executive for IBM Business Resiliency Services, said because the core of any disaster recovery solution is "data," BaaS (Backup as a Service) -- "whether tape or electronic or cloud" -- is a key component of DRaaS.

IBM, Corcoran said, is seeing a convergence of BaaS and DRaaS, "where it's more data protection and availability. Every DRaaS solution requires some level of electronic or cloud-based data replication or backup."

"There is no one-size-fits-all solution or price; the cost of a DRaaS solution is all based on the customer's scope for [virtual machines], disk space [storage capacity] and consumption," Corcoran said. "To help reduce the overall cost... to clients, we create solutions utilizing multi-tenancy, capacity-on-demand, syndication, sharing of assets, many-to-one replications, and other software-defined automation tools."

Newer, smaller companies have been able to use a standard, open systems architecture, making them strong candidates for DRaaS on public clouds. Larger, more established companies with a large volume of applications -- some of them old and costly to replace -- require DRaaS solutions that are hybrid, consisting of a combination of multi-public cloud, private cloud and on-prem IT environments, according to Corcoran.

A significant challenge is managing a fragmented disaster recovery environment. That's why, in 2016, IBM acquired Sanovi, whose software orchestrates the diasster recovery process across the hybrid environment.

"We call this software IBM Cloud Resiliency Orchestration, which can be built into an orchestrated DRaaS solution. In addition, we offer IBM Resiliency Consulting Services to assess, design, build, implement and/or manage any or all of the solution for our clients," Corcoran said.

Recovering from a traditional cloud backup provider can also require an enormous amount of bandwidth, even if it still takes days to recover the data, Berman added.

Sungard and other DRaaS providers have a myriad of pricing models based on SLAs, service-level objectives (SLOs), data capacity needs and how much of the work you want them to perform. A stripped-down service gives users access to compute, software and a web portal but no management; It's basically DR DIY where an company saves on CapEx and OpEx. The basic service costs $50 to $60 per virtual machine (VM). Their top tier SLA will run hundreds of dollars per VM, according to George.

Sungard's fastest RTO offering is five-minute recovery for each VM; recovery times increase with the number of servers/VMs. For example, 100 servers can be recovered in a minimum of two hours, 250 servers will take as little as four hours.

IBM uses SLOs -- an element of an SLA -- when working with more complex hybrid DRaaS environments to establish the specific, measurable characteristics of the recovery as required by a client.

When it comes to recovering servers, Sungard's George said, it isn't as straight forward as it may seem as it needs to be done in the right order; it also includes standing up business applications, ensuring the data isn't corrupted and testing everything to ensure it all works.

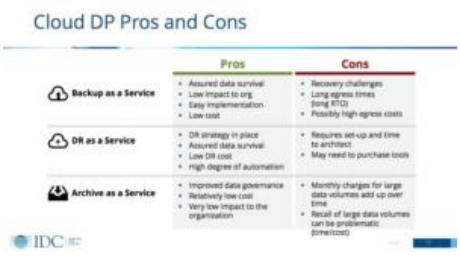

(Click for larger image.)

DRI, a nonprofit that helps enterprises prepare for and recover from disasters, recently conducted a joint survey with IBM on hacking and cybersecurity. The survey of IT managers revealed the mean time to identify if an organization has been hacked is 168 days.

"I knew immediately [from that survey] this is just a different world," Berman said. "It's really something new the fact that your files have corrupted, stolen, or accessed and you find out about it five-and-a-half months later. People used to look at disaster recovery for fire, flood, famine or locust swarm."

For example, in May, the City of Baltimore was hit with a crippling ransomware attack; it took the local government three months and $10 million to recover, not including an additional $8 million in lost or deferred revenue.

"I would have paid the $75,000 [the attackers wanted] instead of the $10 million the taxpayers paid," Berman said.

As awareness increases, the global DRaaS market is expected to grow from $2.19 billion in 2017 to $12.54 billion by 2022, which represents a compound annual growth rate of 41.8%, according to a report from MarketsandMarkets research firm.

Phil Goodwin, an IDC research director, agreed the market flourishing and a lot of it is because of the increase in malware attacks and that DRaaS has fundamentally changed the economics of disaster recovery.

Enterprises are charged for disaster recovery services only when they are used, which makes it cheaper than maintaining a DR warm site or hot site that must be constantly up and running.

"Under the old model, you had to have either two more data centers with more or less redundant hardware or you contracted with one of the big DR [disaster recovery] providers who sold you a subscription at a fairly substantial price that in the event something went wrong you could access their systems," Goodwin said.

"DR as a service really changes that by making those compute resources available on demand as well as now bundling in a lot of additional tools, such as recovery orchestration and workload migration," Goodwin continued. "So, it's not only made it a lot cheaper, but a lot easier to do a disaster recovery."

Additionally, there's a new "race to zero" by enterprises trying to achieve instantaneous RPOs and RTOs.

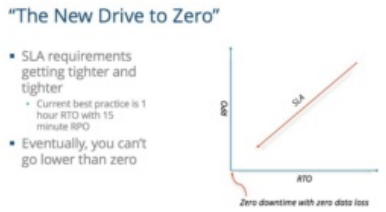

IDC

(Click for larger image.)

"In English, that means zero down time with zero data loss," Goodwin said. "In order to achieve that organizations are implementing immediate replication and continuous data protection that immediately copies anything that comes into the system."

Immediately replicating data to a secondary backup system isn't necessarily a good strategy, though, as a malware infection would also be replicated. There's also been a marked shift by bad actors who are now attacking backup copies first to ensure organizations won't be able to retrieve their data and apps once primary systems are infected.

"So, it's going in an finding the .bak files or others associated with backup and its encrypting or corrupting those backup files," Goodwin said. "So, when you discover you've had a ransomware attack or some other attack, you no longer can go back to your backup copy, which is online in order to do recovery."

Tape's role diminished but never gone

DRaaS services aren't always able to ensure mirrored data up isn't corrupted, as anything backed up may also include malware, which more often that not takes time to discover. And, because of that, users may have to rely on older backups from which to recover and will pay for the additional capacity required to store data longer.

Five years ago, companies kept two to three days of replicated data; today, they may keep data online for 30 to 60 days depending on their RPO needs, according to Sungard's George. A faster RTO can be achieved through more data snapshots versus full backups, George noted.

"Fundamentally, though, you're dealing with a situation where you can sustain some data loss," George said. "That's the tradeoff there."

The same doesn't hold true for tape.

Tape backup creates a physical "air gap" between production data and vaulted data, so if a malware attack occurs it's not able to infect tape cartridges stored offline, Goodwin said.

Still, the data storage and archive market has shifted drastically from tape archival to cloud backup and archive services. The tape market today is well below $1 billion and Gartner estimates it is declining at a 5 percent to 9 percent rate annually.

IDC

(Click for larger image.)

Recently, however, there has been some renewed interest in tape archive systems for long-term storage, IDC's Goodwin said, adding the market is likely to see stabilization as customers rediscover the value of tape in terms of the air gap capability to deter malware infections as well as for low-cost storage.

"As the tape market reaches a certain point, it will stabilize like the mainframe," Goodwin said. "It's a similar dynamic tape where the market will stabilize around specific use cases and continue on for foreseeable future."

IDC

(Click for larger image.)

Anecdotally, Goodwin added, he is also seeing interest in tape from hyperscalers – megacompanies like Amazon, Facebook and Google -- who are using tape in the background for "low cost, long-term storage."

So, tape will continue to be used for storage of large file systems that require mass storage are rarely accessed, such as medical images, geophysical data, or maintenance records like those of airlines that must be kept or the life of an aircraft.

And, tape is still one of the most affordable solutions for very high local capacity, according to Jerry Rozeman a Gartner senior research director.

"As extension of a disk architecture it can still be very effective for protection or deep archive," Rozeman said via email. "Tape is portable, it can be transported to any location. Tape does not requires power and [its] lifecycle is longer then disk.

Tips for choosing a disaster recovery plan

Rozeman offered the following advice :

1. **Set the right expectation.** DR to cloud is one of the options, not always the best option
2. **Test, Test, Test on failover, failback, performance, etc.** Initial sync and failback might require a full synchronization which takes time and might be a showstopper
3. **Evaluate the scenarios.** Be preapred if there is a big disaster and multiple customers disaster recovery are linked to the same provider
4. **Understand that the cloud is different.** DR to cloud can mean a switch of hypervisor, and that changes the technology ecosystem (i.e., some stuff might not be working anymore)
5. **Everything is different.** Once you perform a disaster recovery, you are operating a different environment, troubleshooting is different, train on the required skills
6. **Evaluate your Day 2 scenario**. After a disaster you will continue to need DR services and backup functionality must continue in the cloud

While trusting a cloud service for your capability to recover from a physical or man-made disaster can create angst in the most stalwart of IT managers, DRI's Berman said short of a company building out a private cloud infrastructure and mirroring data to it, DRaaS the safest means possible to ensure recovery.

"With a disaster recovery service, you get two versions of security: The security on your system and the security on the cloud system," Berman said. "I think it's difficult to corrupt it because you can actually test the alternate site whereas if you do 'conventional backups,' you have to stop the operations, you have to go to an alternate site, you have to uninstall the operating system, you have to uninstall the apps; it's pain staking.

"And, I'm much more comfortable with an organization whose entire job is to provide clean access to data and networks, and complete monitoring," he said.

Additionally, data streamed to a disaster recovery service doesn't need to simply sit dormant. Enterprises can use cloud data repositories for productive work, such as test development or analytics. Sungard's George said that's where DRaaS services are heading and will eventually offer data analytics, data mapping and discovery services along with recovery.

More important than the medium on which you're storing data for disaster recovery is ensuring the restoration process will work it's needed. That's one of the advantages of using a DRaaS service. Users can test it relatively easily.

"One of the things about using the cloud is you can actually go to the backup site; it's a logical switch that says take me to my recovery site and it's a real-time test," Berman said. "It's like having dual processors. It makes it easier to test as opposed to the standard way people test and they take the backups and they dummy out files and they spend three weeks preparing for it.

"Using the cloud as a backup, you can throw the switch and say I'm now processing from my alternate site," Berman added. "I think it gives you a lot more flexibility."
