TRON Disaster Recovery Orchestration

TronLabs Romania

Version 1.0

TronLabs Romania is a TRON SR Candidate from Romania (RO) that runs its own private virtual datacenter.
The software used in this whitepaper is **Veeam Backup & Replication version 9.5 Update 3** and **N2WS Cloud Protection Manager**

Introduction

As the global economy continues to further embrace digital transformation and eliminate borders, Information Technology (IT) has become essential to almost every facet of the modern business world. Yet while the advances and increasing adoption of digital technologies have had a profound, positive impact when delivering on strategic priorities, one thing remains true. The event of an IT outage — total or partial — has massive implications to our Ecosystem continuity. In fact, 66% of the big enterprises admit that their digital transformation initiatives are being held back by unplanned downtime[1].

Threats to IT service and business continuity come in all shapes and sizes:

- Both planned and unplanned.
- Both accidental or intentional.
- Both Man-made or natural.

Regardless of the reason for the outage, it is imperative to plan for disaster recovery (DR), whether that disaster results in the outage of a single service, a server or application, or the entire IT infrastructure.

Disaster recovery challenges

Planning for DR is not a new strategy — it's been around for decades. However, the problem with a decades-old approach is that it isn't designed for today's abundance of and reliance on IT infrastructure. Forrester reports that just 18% of surveyed businesses feel very prepared to recover their data center in the event of a site failure or disaster[2]. Legacy solutions just cannot cope with a constantly changing, ever-complex IT environment. There's more to protect and manage, with strict service-level objectives (SLOs) and near-zero tolerance for downtime[3]. Expensive, manual processes are neither scalable, efficient nor cost-effective. A modern DR solution must be as fast and efficient as the applications, services and infrastructure that it's designed to protect.

DR also has compliance implications — Tron being a dPOS with no governance in technical terms, it has no laws, regulations and standards set in place to ensure an organization's responsibility to the reliability, integrity and Availability of its data. If a SR miss blocks production there are no penalties, just some sort of lower rewards. While compliance varies from industry to industry, in an ecosystem without regulations, all but one thing holds true — compliance deficiencies are not an option, and come with significant reputational risks.

This begs the question, if many SR's and SR candidates have DR plans in place, why do the majority feel that their Availability — and compliance — are at risk? We can break this down into four critical components; the plan itself, documenting the plan, testing the plan and the actual execution of failover and failback.

#1 Planning

At the heart of every successful DR plan is knowing what to protect and how to protect it. If we refer to TRON these would be the three types of TRON Nodes:

- Witness Nodes that generate blocks
- Full Nodes that contain a copy of the chain.
- Solidity Nodes that interact with the wallets.

Analyses and assessments are critical when it comes to determining which elements of the IT infrastructure are essential TRON operations. Understand and define the outage risks and their impact, tolerance for downtime and the processes and procedures necessary to restore access.

As a Tron SR you need to ask yourself these questions:

- What are the mission-critical services and applications?
- What are the operational and financial consequences of losing access to such services for me and the rest of the network?
- How often must we be able to recover or return to a point in time?
- How quickly must we be able to recover to meet our needs and ensure that the network is not in any way affected?
- How often do we need to reevaluate?

#2 Documenting

Documenting the DR plan is critical to any successful DR strategy. Documentation clearly defines the DR plan in its entirety and most importantly, the established processes and procedures to follow in the event of an outage. Sometimes when it happens no one knows what to do and how to resolve the problem because the plan was made by someone else.

Manually documenting the DR plan can be an expensive and lengthy process, especially for larger environments and the continual changes occurring. This is evidenced with only 14% or organizations updating their DR plans and documentation continually[2].
Outdated and improper DR documentation doesn't just put the business's Availability at risk though.

In the normal business world, audits are common and may identify failure when complying with regulatory and legal standards, with the inability to prove security, resiliency and recoverability resulting in expensive fines and reputational damage.

In the TRON network there is no such thing. I have yet to see a SR or SR candidate besides us who has this in place.

#3 Testing

Testing is a critical component when it comes to ensuring preparedness, yet only 19% of surveyed organizations run a full test more than once per-year[2].

In TRON there were a couple of tests performed at the main net launch. The tests were just some scripts doing failover for the witness nodes, nothing fancy or uncomplicated. Since then I did not hear anything similar being tested.

This is an unacceptable frequency when considering the reliance on a constantly changing IT environment and the fact that we are considering a valuable dPOS blockchain.

Frequent end-to-end testing can prove challenging, especially when we consider the preparation and resources required to conduct a full test of all systems and their dependencies, without impacting the production environment and interrupting end users, yet it's imperative that this must be done on a regular basis, at least every 1 or 2 month the latest.

#4 Executing

The constant state of change within ever-complex IT environments are a key point of failure when performing failover. Mismatches resulting from uncaptured changes to systems or configurations are the most common technology impediments to a successful recovery[2], particularly when we consider the dependencies one application or service (Full Node) may have on another (Solidity Node).

As one IT service is impacted by an IT outage, it causes a chain reaction and affects the performance and/or Availability of others, cascading those failures into even more systems[3].

Example:
11 Oct 2018 - The network stopped producing blocks and it was a global outage that did affect al elected SR's and the block production. This war resolved hours later and the syncing of the network took its due time.

As you can see in case of an outage, manually powering these on, in perfect sequence is next to impossible, especially when we may be dealing with hundreds, maybe thousands of services, severely impacting what can be recovered, and how long it takes until the service is restored.

Ways of achieving recoverability

I. Via Backup

The Golden 3-2-1 Rule

1. Have at least three copies of data

By three copies, I mean that in addition to your primary data, you should also have at least two more backups. Why isn't one backup enough? Imagine that you keep your original data on device #1 and its backup is on device #2. Both devices have the same characteristics, and their failures are statistically independent (they have no common failure causes). For example, if device #1 has a probability of failure that's 1/100 (and the same is true for device #2), then the probability of failure of both devices at the same time is:

1/100 * 1/100 = 1/10,000

This means that if you have your primary data (on device #1) and two backups of it (on devices #2 and #3, correspondingly), and if all devices have the same characteristics and no common failure causes, then the probability of failure of all three devices at the same time will be:

1/100 * 1/100 * 1/100 = 1/1,000,000

This is why having more copies of your data means you will have less risk of losing data during a disaster. In short, if your data is important to you, be sure to make at least two backup copies.

Note: Another reason to create more than two copies of data is to avoid the situation when the primary copy and its backup are stored in the same physical location.

2. Store the copies on two different media

In the section above, we assumed that there were no common failure causes for all of the devices where you store your data copies. Obviously, this requirement cannot be fulfilled if you save your primary data and its backup in the same place.

For example:
RAID is not a backup solution and disks from the same RAID aren't statistically independent. Moreover, it is not uncommon after one disk failure, to experience failure of another disk from the same storage around the same time.

That's why the 3-2-1 rule suggests that you keep copies of your data on at least two different storage types, such as internal hard disk drives AND removable storage media (tapes, external hard drives, RDX, USB drives, or BluRay disks), or on two internal hard disk drives in different storage locations.

3. Keep one backup copy offsite

Physical separation between copies is important. It's really not a good idea to keep your external storage device in the same room as your production storage. If there was a fire (knock on wood!), you would lose all of your data.

If you have no remote or branch offices (ROBO), storing your backups to the cloud might also be an option. Tapes taken offsite are still popular among all company sizes. Amazon is offering a Virtual Tape Library in glacier.

What are the most common scenarios where you can leverage this approach?

Recovery

- Quick Rollback: If a small amount of data was accidentally deleted a Quick Rollback a incremental data restore. Instead of restoring an entire VM or VM disk from a backup file, Veeam Backup & Replication will recover only those data blocks that are necessary to revert the VM or VM disk to an earlier point in time. Quick rollback significantly reduces the recovery time and has little impact on the production environment.
- Instant VM Recovery: If one of the nodes gets corrupted an Instant VM Recovery will be performed. With instant VM recovery, you can immediately restore a VM into your production environment by running it directly from the backup file. Instant VM recovery helps improve recovery time objectives (RTO), minimize disruption and downtime of production VMs. It is like having a "temporary spare" for a VM, the TRON users continue to have access. while you can troubleshoot an issue with the failed VM.

  In many respects, instant VM recovery gives results similar to failover of a VM replica. Both features can be used for tier-1 applications with little tolerance for business interruption and downtime.

- Guest OS File Recovery: If we need just a file or a couple of files from the backup, we could use Guest OS File Recovery. It can restore files in place or to a different location for the following File and Operating Systems: Windows (Fat, NTFS, Refs), Linux, Solaris, BSD, Novell Storage Services, Unix, Mac and other file systems.

Testing

- Testing a new deployment of TRON Odyssey. Beside disaster recovery matters, instant VM recovery can also be used for testing purposes. Instead of extracting VM images to production storage to perform regular DR testing, you can run a VM directly from the backup file, boot it, change the IP to an unused one, deploy the new software and make sure the VM guest OS and applications are functioning properly.

II. Via Replication

Replication is your safety net. If a production VM with a Node goes down, you can immediately fail over to another VM replica, giving users access to the services and applications they need with minimum disruption while you resolve the issue. With replication you can avoid data loss and dramatically improve your recovery time and point objectives to less than 15 minutes for ALL applications and data, minimizing the impact on your infrastructure.

Replication has multiple benefits:

- Restore normal operations while your replica is running by failing back to the production VM, to a new location, or just making the replica VM your new production live VM. You can maintain multiple replica restore points, so if your latest replica is corrupt, rolling back to a previous restore point is available as an option.
- Plan your entire failover in advance, and start it with a single click using the Failover Plans. Add VMs from replicas, move them up or down to get a boot order and set a delay for each VM so that they don't start before a previous one starts up. You can fist start the TRON Full node and the set the TRON Solidity Node to start second as it depends on the Full node. It's that Simple!
- If needed you can also facilitate data center migrations or perform maintenance on your production hosts with the Planned Failover feature. Planned Failover shuts down the source VM, replicates any changes to the target VM and starts the VM—all with no data loss and little downtime.

To achieve minimum data loss we run or TRON infrastructure with a 15 minute replication interval. Once started the TRON Witness Node needs at most 5 minutes to sync with the network. If we would be an elected SR among the first 27 and this outage would happen, this translates to about 11 blocks missed while full synchronization with the rest of the network is performed. This is however further improved, because we run two separate TRON Witness Nodes with different priority and once the first one fails the second one takes over.

Ok, I understand that Infrastructure on premises is susceptible to failure, but we run our workloads in the Amazon Cloud as per TRON Recommendation. We do not have to worry about any of these? Right?

Actually, you are not right and this is why. Implementing cloud technology is about simplifying logistic efforts. In this white paper I will talk about Amazon Web Services (AWS) as it's the most popular provider, but no matter it's private cloud, public cloud or hybrid cloud, there are clear benefits and the cloud became the new normal for every piece of technology we interact with, where this is your mobile phone, your tablet, your PC, or a dPOS blockchain network like TRON with hundreds or thousands of nodes.

Businesses that are migrating to the cloud are looking after very clear benefits: eliminate the burden of maintaining an on-premises infrastructure, lower the costs and enable business scalability and flexibility. This is also true for Tron SR's and SR candidates who do not run their own infrastructure.

In order to understand what it is, let's start with the definition of cloud computing:

Cloud computing is the on-demand delivery of compute power, database storage, applications, and other IT resources through a cloud services platform via the internet with pay-as-you-go pricing.

So you use the providers compute infrastructure and you pay as much as you use. It's like using like an on-premises infrastructure and run VMs on it, except the part that you don't have to build it, maintain it and secure it, you get all these in one package.

But you would backup your on-premises infrastructure, wouldn't you?

The same applies to cloud computing. Let's take a look at AWS Shared Responsibility model on the next page. It explains very clear who is responsible and for what.

How responsibility is split between the user and cloud provider.

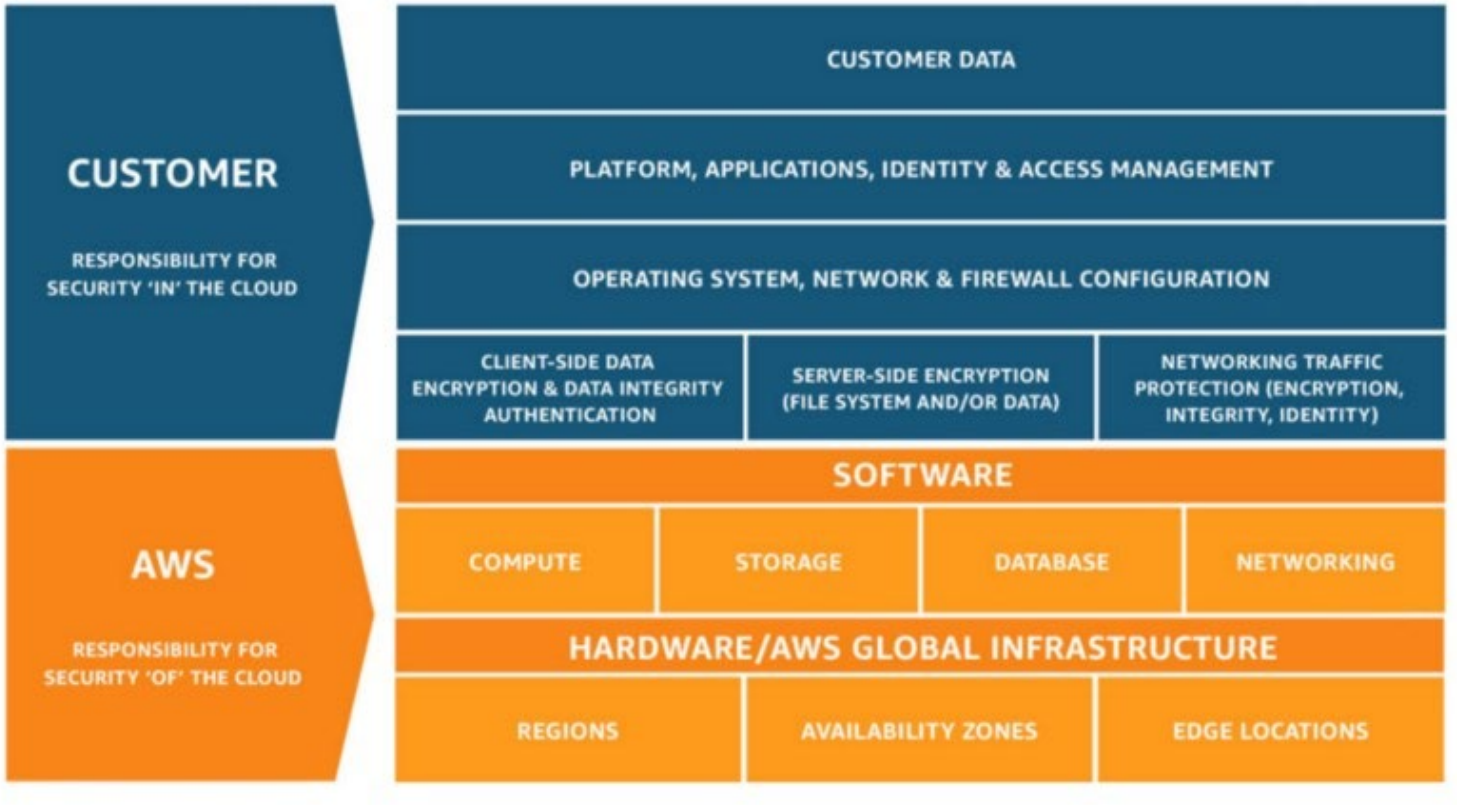

Source: https://aws.amazon.com/compliance/shared-responsibility-model/

Let's read it starting from the lower part and then move to the upper part.

Part 1 - in orange, the Public Cloud provider.
AWS is responsible to ensure the Availability of their services, to provide the selected resources (compute, storage, database, networking) and the physical security of the facilities in which the service operates. This is great, because you don't need to worry about the operational side of the services your run.

but..

Part 2 – in blue (your responsibility).
You do need to make sure your data is safe. There are many things that could go wrong with it: accidental deletion, ransomware attacks, compromised AWS account or AWS outages (unlikely, but it happens). AWS doesn't have access to your data and it's not their job to manage it. If you run EC2 instances, it's you who must perform security configuration and management tasks, including backup of your data. The same applies to updates and security patches for the OS and applications you are running.

Ok, got it ! It's clear that I need to backup my AWS workloads. How do I do it?

Everything around us is powered by data—needless to say, it must be safe. The default way to backup EC2 instances is by taking point-in-time snapshots of your Amazon EBS volumes. Recently, Amazon introduced Data Lifecycle Manager for EBS Snapshots to enhance their backup capabilities. This cool feature allows to automate the creation, retention, and deletion of snapshots taken to back up your Amazon EBS volumes.

However, these backup tools provided by AWS are limited and not strong enough for a complete backup and DR strategy. For example, EBS snapshots are not application-consistent, meaning that you can't backup applications while they are running. Moreover, it doesn't provide any options for monitoring and reporting and it lacks disaster recovery (DR) capacity, which can cause troubles for big enterprises companies, or organizations running their production workloads in the cloud.

Ok, so how do I do it then?

Luckily, Amazon Web Services is very popular and there are third-party solutions for comprehensive backup & DR on the AWS Marketplace. With such solutions, organizations are able to complement the native capabilities of AWS and achieve a high-level of data protection in the cloud: disaster recovery (cross-account and cross-region), file-level recovery and compliance. Businesses, companies and TRON SR's can also significantly lower their AWS bill with the ability of moving EBS snapshots to Amazon S3 buckets for long-term retention, turning their cloud investment into an even better deal.

The most effective & efficient backup for AWS.

The software you can use to protect your Amazon workloads is the Cloud Protection Manager from N2WS.

Cloud Protection Manager leverages native AWS technologies to utilize block-level and incremental snapshots.

Simply put: this is the most effective & efficient method of backup available for AWS. Plus, you get complete control of scheduling and can **easily set up automated policies so that you can have 24/7 access to your data**.

Ensure 100% uptime and availability of your data.

*With Cloud Protection Manager, you can rapidly recover entire EC2 instances across accounts or regions within 30-seconds. This **near-instant recovery time** is made possible with incremental snapshot technology and ability to create recovery sites around the world. Protect yourself from account compromises, malicious attacks, and more.*

Conclusions

It's clear that legacy approaches to DR planning and strategies are not equipped to deal with digital transformation initiatives and the evolving data center, hindering adoption and growth of new technologies and the business. It's also clear that some organizations have not prepared as they should and they need to reevaluate this as it's mandatory to be prepared in case something bad happens.

Solutions that enable orchestration and automation of DR and compliance are leading the way to greater business continuity, helping to ensure the continuous delivery of production IT services that businesses, companies and organizations rely on and compliance with industry regulations.

If you are using the cloud, it's essential for your organization to understand that while you get many benefits from "migrating to" and "running your workloads in the public cloud", you are still responsible for your data. The ROI for a solid AWS backup & DR solution is very high—just imagine what would happen to your organization in case of downtime and loss of essential data. Money is not everything, the worst thing you can loose, is your reputation.

References

1. *ESG Availability Report, 2017*
2. *Forrester, The State of Business Technology Resiliency, Q2 2017*
3. *Gartner Predicts 2017: Business Continuity Management and IT Service Continuity Management*

Software downloads

Veeam Backup & Replication

Download a Trial or the Free version of the software from here.

FREE NFR Key for NEW Veeam Availability Suite 9.5 ➔ Enrolment link

This license allows for non-production use of Veeam Availabilty Suite™ 9.5 in your home lab, without any feature limitations for one year, 2 sockets. The license works for both VMware and Hyper-V environments.

Cloud Protection Manager

Download a Trial or the Free version of the software.

About the author:

Dorian Tang - I am an IT consultant and Azure Architect. I've been using, designing and deploying both VMware and Microsoft based solutions since 2008. I'm specialized in designing and implementing private and hybrid cloud solution based on the Microsoft software stack, Datacenter migrations and transformation, disaster avoidance.

I've been in the IT industry since 2002 and I'm in consulting services since 2008 when I've started working for Microsoft in Bucharest providing Level 3 and above support for enterprise customers and did so till 2016, when I've joined Veeam Software as an engineer, helped build the local tech support organization from scratch and currently lead the local tech support team for the DACH area.

I was always curios and wanted to learn new stuff, this is why I've dedicated my time to learn and grow. I hold multiple certifications as MCSE (Server infrastructure 2012, Private Cloud, SQL Server 2012/2014, Sharepoint 2013, Cloud Platform and Infrastructure 2016, Data Management and Analytics 2016, Productivity 2016) and MCSD (Windows Azure). I am also a Veeam Certified Engineer (VMCE) and Microsoft Certified Trainer (MCT).

From July 2018 I run a Tron SR candidate under the TronLabs Romania name.