TOTAL UPTIME
The Cloud Availability Platform

4 Cloud Gotchas to Avoid

Data center, on-premise, public-cloud, private-cloud, multi-cloud, hybrid-cloud, IaaS, PaaS, SaaS… it's a nightmare even for the most seasoned IT experts.

There is no question that you understand the tremendous benefits of everything the digital transformation (DX) has to offer. Cloud adoption rates are increasing significantly, and IT budgets are morphing to follow suit. But distributing applications and infrastructure around increases risk, introduces complexity and challenges availability at every turn.

To embrace DX and to come out on top, there are four gotchas you must avoid. Understanding them and their potential impact to your organization is critical to cloud success.

1: Availability

Assumption: Moving my applications to the cloud will provide unprecedented availability.

This is a common assumption, and unfortunately, it is rarely true. The more you distribute your infrastructure and applications, the greater likelihood that you will negatively affect uptime for two primary reasons: reduced control and increased complexity.

Let's dig a little deeper. When all enterprise infrastructure is in a single location either on-prem or at a colocation facility, you have 5 basic points of availability to worry about. These are simplified concepts, of course, but they primarily consist of:

- Security
- Power
- Network
- Cooling
- Infrastructure (servers, routers, switches etc.)

To improve availability, you implement redundancies in each of the 5 areas. While this may lack simplicity, the benefit is complete control. Yes, you are responsible for more, but you're also in a position to respond to problems in a manner appropriate with your corporate policy. Furthermore, you get to determine the quality of the components used and the level of care and maintenance taken to ensure availability.

In this scenario, if something goes wrong you could walk over and deal with it. If a router crashed, you'd reboot it. If a hard drive failed, you'd replace it, and so on. You're directly in control over availability and the mean time to repair (MTTR). Yes, it takes more effort, and it is usually more expensive, but provided you aren't running a low-budget operation, availability is usually very high.

Enter the Cloud

If you move absolutely everything to a major cloud provider such as Amazon Web Services, Microsoft Azure, Google Cloud and the like, then you've replaced five points of availability with one. On the surface, this seems extremely wise and rather enticing. But when they have an outage – and they do – you are completely helpless. You have traded control for something else, like convenience, scalability or lower operating costs. One provider, in this case, is probably unwise. It's a single-point-of-failure.

Of course, the major public cloud providers should be incredibly reliable – and they are indeed very reliable – but incidents do happen. We all know that, and we've seen the stories in the news time and again. Usually, on average once per year, one of the major cloud providers has a significant event.

Moving everything doesn't happen in real life

Our example of moving everything to the cloud is quite uncommon in reality. According to RightScale's January 2018 survey, 81% of enterprises have a multi-cloud strategy and 51% have a hybrid-cloud strategy. This suggests that rather than eliminating their own data center completely (with the 5 points of availability we just considered), enterprises are actually adding on to it. If you simply add one cloud to your existing environment, you now have 6 points of availability to consider. A 20% increase in responsibility.

But again, that's still an overly simplified example. Ask yourself this: Do you use Salesforce, Office 365, Azure ADFS, a hosted ERP system? What else do you use? Again, RightScale's survey found that the average number of clouds used in an organization is 4.8. This means, that the 5 basic points of availability have almost doubled for the average enterprise.

THE TAKEAWAY:

To maintain availability, you just need to know what the tradeoffs are and have a strategy to mitigate them. Moving to the cloud is great. It decreases responsibility and cost and increases your ability to scale, allowing you to add capacity at will to meet demand, which could be a significant boost to availability! Those are worthy benefits indeed.

2: Security

Assumption: Moving my applications to the cloud will make them more secure.

If you consider the prior on-prem/colo data center scenario, you have two primary security challenges. Physical and Logical – that is, securing what can enter the building, and what can come over the wire.

You protect the physical with a secure building, access control, cameras etc. You protect the logical with firewalls that have full layer 7 application scanning and intrusion prevention capabilities. Assuming you keep the firewalls up-to-date with the latest patches, signatures and so forth, have configured alerting and review the logs on a regular basis, your applications should be relatively well protected.

Now extract your favorite critical application from your on-prem data center and put it up in the public cloud. Interconnect it with a VPN tunnel back to your data center, a common strategy to allow it to communicate with data sources or other technologies at your primary site. From a physical security perspective, you have no idea who has access, but that is probably not your biggest concern. Logically, you've just taken that app outside of the corporate umbrella you created and placed it out in a more open environment. You've introduced a secondary access point.

To make matters worse, the default firewall at any major cloud provider is just an access control list (ACL) of ports, protocols and IP addresses. Unless you specifically add a Web Application Firewall (WAF), you may have created a back-door that is considerably weaker than what you've implemented at your corporate data center.

Even with the addition of a WAF subscription in the cloud, further complexity has been introduced and it challenges security. You still have two entry points. Two WAFs to keep updated, two different vendors, two different UIs or APIs to master. It's not a great situation at the cost of security.

Security Mitigation

Many vendors today like Cisco, Fortinet and others provide virtual versions of their physical firewalls that you can run at major cloud providers. Some even provide a unified management console for distributed firewalls. If your vendor does, seriously consider it! Your cloud security will be stronger if you keep the vendor, software version and configuration the same at both sites, especially if they are both interconnected and allow network entry via either location. Yes, you still have twice as much to manage, and it may cost more, but by removing some of the complexity, you can regain some security.

But remember, the average enterprise has 4.8 clouds. That's more than the two in the example we just gave. Are you prepared for that level of management? And how will you protect sensitive data? The kind of Personally Identifiable Information (PII) that falls under HIPAA or PCI? This impacts security with a new element of "confidentiality" that has much stricter standards for safeguarding. For example, did you choose a PCI Certified cloud?

THE TAKEAWAY:

Great security is attainable with a multi-cloud or hybrid-cloud strategy. Just understand the risks so you can mitigate them. Strongly consider managing security for all entry points with a vendor that provides a single pane of glass.

3: Performance

Assumption: Moving my applications to the cloud will improve their performance.

Cloud performance is directly affected by location and implementation. Without careful planning, performance will decline when applications are moved to the cloud. Understanding the many differences between traditional infrastructure and cloud infrastructure is paramount to maintaining or improving application performance.

There is no one-to-one comparison matrix between traditional and cloud because every provider offers something slightly different. You must specifically evaluate and test in order to determine what the impact will be before moving any applications to the cloud.

Points to consider:

- **Are you moving your applications from physical servers to cloud servers or containers?**

  If you answer yes, you must evaluate the impact of moving from dedicated hardware to shared hardware. Whether virtualized or containerized, you must also consider the impact your new and unknown neighbors will have. Plan for the unknown. Most clouds offer auto-scaling, and this might just be your lifesaver.

- **When moving an application, will it receive similar resources?**

  Again, not a difficult question, but something to consider. 1 CPU core on a dedicated machine is not equal to 1 CPU core on a virtual or cloud machine. RAM is different, disk speed is different. And if you're switching from a VM to a container, it's a whole different animal. The only way to know if you're comparing apples to apples is to either have the hardware specifications from your cloud provider (which are almost impossible to get) or to load test your application in both environments to determine what cloud resources will provide equal or better performance.

- **In what geographic location will you place your application?**

  The cloud makes it easy to place an application in a different geographic area to mitigate risk, but it can also introduce undesired effects. If your users are internal, and the app is internal, but you want to move it to an external public cloud, you just moved the goal posts. Consider geographic proximity and network latency between your users and the application. How will you connect the WAN? A little packet loss on a VPN tunnel can make applications quickly unusable and increase end-user complaints exponentially.

Lastly, many organizations assume that the right SLA is what guarantees performance in the cloud, but this is simply not the case. SLAs provide an expected range of availability and performance, but there is no room to negotiate a better SLA unless you are an enormous enterprise with buying power.

THE TAKEAWAY:

Great performance is completely possible in the cloud. Just consider every angle. Cloud can help too! For one thing, it makes it easier to implement CDN or other edge caching. The right cloud provider can implement TCP optimization across the WAN too.

4: Integration

Assumption: Interconnecting the network will be easy.

All public and private cloud providers offer VPN tunnels or direct connections. But unfortunately connecting all the dots is easier said than done. It's best to create a detailed logical diagram of how you envision everything interconnecting to avoid any surprises.

Points to consider:

What is the primary point of external access to the application?

Will access be through a public IP on your corporate firewall, tunneled back to the cloud? Will it be a public IP at the cloud provider NATed to the VM or container? Consider how public access will be achieved and note the entry points to the network. If the application is only internal, this reduces risk and complexity.

How will you integrate the clouds?

Will you use VPN tunnels? Direct connections? A combination of the two? Or no interconnection at all? If you create connections, will you build them between separate endpoints and use dynamic route distribution such as BGP for automated routing in the event of failure? The more automation you can implement at the outset, the less frustration you will have during a network incident.

Will you use separate subnets for each of your clouds?

Always implement separate subnets for your public, private and on-prem networks. Choose subnets that are less likely to be existing in customer environments (like 192.168.0.0/24), in the event you connect to them in the future. Implement granular firewall rules between all sites as if they were all untrusted. With multiple points of entry, never assume traffic on the LAN got there the right way.

How will you balance the traffic or redirect it in the event of failure?

Everything fails eventually. How will you ensure that infrastructure is not overwhelmed with traffic while other components sit idle? How will you redirect users from one site to another when disaster strikes to ensure you meet service level agreements?

THE TAKEAWAY:

Integrating your cloud environments securely and reliably is completely possible. Just consider every point of interconnection and failure and plan for outages. If you implement monitoring, automation and intelligent traffic distribution there is no question you can master the cloud.

Conclusion

You can and should look to the cloud. It is the future, and it is bright!

Mitigate the availability, security, performance and security challenges and you'll be in a great position to come out on top. If you consider every component and how they interconnect and impact your organization, there is no question you will become your organization's Digital Transformation hero.

Total Uptime can help!

We built our global platform to solve these four challenges (and more). We've helped Fortune 500 and small business alike take advantage of the cloud while still avoiding these the gotchas. Here are a few success stories for further reading:

- **Informatica**

  Enterprise cloud data management company improves control of application traffic and increases security and availability - https://totaluptime.com/case-studies/informatica/

- **Pittsburg State University**

  Public university meets Disaster Recovery requirements for their distributed authentication infrastructure https://totaluptime.com/case-studies/pittstate/

- **Morris, Manning and Martin LLP**

  International law firm achieves automatic ISP failover to maintain constant connectivity to their IT resources https://totaluptime.com/case-studies/mmmlaw/

- **XML Travelgate**

  Travel integrator increases availability and control with cloud load balancing https://totaluptime.com/case-studies/xmltravelgate/