An IT Checklist for Cloud Based Global Architecture Solutions

- RELIABLE
- FLEXIBLE
- SECURE

Moving on premises solutions to the cloud is compelling for a number of reasons but there are still some areas of potential risk around picking the right vendor. Similar to on premises solutions, you still need to assess product features, capabilities, professional services, support and cost. In addition, you need to investigate and validate vendor claims related to reliability, flexibility and security.

This white paper is designed to help you assess and evaluate cloud solutions as viable replacements for on premises infrastructure. It was written to provide an understanding of how software development and company processes can impact vendor aptitude. Business buyers will likely focus on whether a particular cloud solution will be a good fit for the business problem. That leaves IT and other tech savvy decision makers to help the company determine if the solution will provide the right capabilities in a secure and reliable manner for years to come.

Reliability

Uptime and Availability

In the early days of cloud, led by Salesforce and other CRM vendors, concerns of reliability surged. Reliability was top of mind for those companies looking to make the switch to cloud. It wasn't that their existing on premises systems had great uptime. They had control over fixing any issues that came up, and that was viewed as more reliable.

The ultimate goal is 99.999% uptime.

Now, reliability of the cloud is more understood and the discussion of reliability for cloud venders has coalesced around the idea of uptime comparable or greater than that of on premises software. The ultimate goal is 99.999% uptime. This equates to about 5 minutes of downtime over the course of a year. This is orders of magnitude harder and more expensive than achieving 99.9% reliability which equates to about 8 hours of down time over the course of a year.

Let's pause here and look at these numbers: What is 8 hours of downtime within a year that contains 24 hours a day, 7 days a week, and 365 days? How many corporate IT run services can claim to have this kind of uptime? For most organizations, the answer is not many.

But I hear you saying, "uptime like that isn't required for most applications." Of course! You are absolutely right, and that is the point. Compared to on premises solutions most cloud companies cleared the uptime requirement years ago.

Availability brings the discussion of uptime back to the user. What most businesses need is for solutions to be available during operating hours. A multi-tenanted cloud environment providing 99.999% uptime and effectively achieves this for all businesses. If your business is UK based and only operates between the hours of 9 to 5 then a system that is always available between those hours would meet your business needs. A minimum uptime of 30% is all that would be required assuming that the system was always available from 9 to 5.

The reliability requirements for your solution will of course depend on the nature of the usage and its importance to the business. Look for companies that post their historic reliability in a public environment and who's past performance over the last 12 months meet your availability requirements.

Redundancy

Design the architecture so that any individual failure has a fallback within the architecture.

In order to create a highly reliable system you need redundancy. Since some failure at all levels of a system is inevitable, the task is to design the architecture so that any individual failure has a fallback within the architecture and can be picked up without human intervention.

Areas of consideration for redundancy include:

- Hardware level redundancy and or VM level redundancy

  In this age of cloud, hardware is still required somewhere to run the software. The decision to co-locate hardware, or pay someone else to design and maintain the hardware (AWS, Azura, Google, IBM etc.) for your proposed cloud provider, comes mostly down to cost these days. Most cloud providers end up doing a mix of the two based on the age of the business, amount of computing power necessary, and the desire for more or less capital or operating expense.

  Managing hardware requires specific skills and cycles but can yield cost savings and uptime advantages if done correctly.

  Whether your service provider manages the hardware themselves or is using a Platform as a Service (PaaS) provider, you need to think about process redundancy and how processes share resources and fail over.

- Process redundancy

  Processes within the larger service need to be available in order for the service to run. Obviously some processes are more critical to the operation of the service than others. Your service provider should have a map of those processes that need active-active redundancy or N+1 redundancy vs those that can operate with passive redundancy. Again, for the purposes of using a service from a vendor, you will need to assess the importance of the service to your business and then feel comfortable with the steps taken by the service provider to ensure service reliability.

- Network redundancy
  If your service provider co-locates their service, it's important to understand how the location is setup vis-a-vis network redundancy. Are there multiple routes to the internet? If one carrier becomes unavailable for some reason, can another carrier pick up the load and handle the traffic? What SLAs are available to your service provider around bandwidth, redundancy, and availability? Your service provider should have answers for these questions and express knowledge of which elements need redundancy and which elements can live without it.

  If your service provider is using a Platform as a Service, then the carrier redundancy is up to that provider. Your service provider should have a good understanding of those elements as well even if they aren't managing the carrier redundancy as part of their business.

- Geographic redundancy
  Data centers should be active-active and designed to handle less than half of the traffic each. You should talk to your service providers about what is required to continue service in the case of the failure of a data center either due to power, network, or some kind of physical event. Does it require logout and login or something more involved? Is action from someone on your team required? If so, do they need specific skills? Is action from your service provider required? If so what are the steps required; do they have people ready to complete those steps? Is it automatic? The more automation the better but fully automatic systems have a higher cost. It shouldn't surprise you to learn that there are manual steps for some of the less frequent failure events. The important thing is to make sure they have plan in place and you and they are ready to implement it in case of a failure.

Flexibility

As a consumer of a service you might wonder, "Why do I really care how that service is built?" If it works, and does what I need it to do, it could be built on Cobol and the value would be the same to me.

Flexibility is what really matters. Over time solutions can become complex. If they are built as a wad of highly modified code, they can become fragile. Not fragile in the sense of reliability, although that is certainly possible, but fragile in the sense that the system can't be easily updated.

Make sure you are comfortable with their ability to serve your optimistic growth requirements of their service.

What does that mean for your business? Down the road, or even tomorrow, it can mean your provider will be slow to add features, integration points, security fixes, or really changes of any sort.

This will force you to move off of your existing platform and suffer all the associated disruption of a move: retraining employees, system setup, testing, etc. If the system you are on can't keep up with innovations, you will eventually have to jump to a new service.

How much excess volume does the service provider system have?

What typical fluctuations do they see and how much lead time do they need to add capacity?

How can you tell if the company you are considering is going to keep up with the times? Ask questions about how the service is architected, and with what technologies it's built. If the answer doesn't articulate the use of modern technologies (or too modern: read new, still in development, and untested by the market) that could mean the organization has a fragile architecture and may be unable to keep up with changes driven by the rapid innovation that transforms consumer behavior.

Ask how the system was built, what the underlying third party software components are. Microservices are increasingly being used to ensure that a system is flexible and extensible. If your service provider is using them, it's a great indication they will be able to keep up with the latest market demands. Ask about RESTful interfaces. Much of this will come out in a good architecture discussion and that includes other elements covered in this white paper.

Security

Security may be very to critically important depending on the industry in which your company operates. Definitely, for cloud businesses that have access to, or may store, sensitive data, keeping that information encrypted and out of the hands of nefarious people or groups should be a top priority. But how can you tell the difference between those companies that claim high security and those that have put the systems and practices in place to be highly secure?

If security isn't covered on a company's corporate web site, it could be red flag that they haven't taken proper steps to secure data within their environment. Companies may talk about two forms of security verification. One is assurance which depends on third party attestations. Attestations provide assurance that security controls

are in place and operating effectively. The second is compliance. Generally speaking, attestations are stronger. However, consider the source of the attestation. An attestation from a reputable security consultant is stronger than one from cousin Bob.

Balance the needs of your business based on your industry against the other business requirements that you have for the solution.

Conclusion

The move from on premises to cloud-based architecture has clear business benefits. But the technical aspects deserve equal consideration. Most cloud vendors offer solutions at or above the reliability, flexibility and security of on premises. Hence, cloud systems are rapidly subverting on premises business models. However, no architectural decision should be made with haste. Not all vendors are created equally. You and your team will want to invest in a solution that will both meet your immediate needs and innovate for future breadth of service and competitive longevity.

YOUR SOLUTION EVALUATION CHEAT SHEET:

1) What is the vendor's uptime? _____
2) What areas have redundancy?
   - ☐ Hardware level
   - ☐ VM level
   - ☐ Process
   - ☐ Network
   - ☐ Carrier
   - ☐ Geographic
3) What contingency plans exist for possible failures? _____
4) Is the system flexible? _____
5) Does it incorporate microservices? _____
6) Does it have a RESTful interface? _____
7) How much volume does the system have? _____
8) How fast can it scale? _____
9) Does the company have security assurance and compliance? _____
10) Is the vendor transparent about security policies? _____

4000 Executive Parkway, Suite 400
San Ramon, CA 94583
925.201.2000
www.five9.com

2006_G_ITCHECKLISTCLOUD_2017