RKON | Whitepaper

Why 50% of the companies that go to the cloud regret it. (And why the causes may not be what you think.)

By Marc Malizia
Chief Technology Officer, RKON Technologies

A recent EMA survey estimates that as many as 50–60% of companies that made the move to cloud computing failed to achieve any significant portion of the benefits they intended to gain.

To avoid this happening to you—and keep you from losing valuable ROI in the early stages of the transition—we've outlined some false assumptions companies make about cloud computing:

1. **Thinking It's All or Nothing –** Clients figure out quickly that only in very rare cases can they put all of their systems into the cloud—and that there are many reasons why certain systems aren't a good fit either technically or financially.

2. **Thinking It's Cheaper –** The promise of cost savings erodes pretty quickly when customers find they can't go to the generic public cloud offering with the low price structure that attracted them in the first place. This is the dominant reason the vast majority of cloud initiatives fail.

3. **Thinking It's Turnkey –** Don't think that once you figure out your solution, the provider is necessarily going to help you migrate to the cloud—and help you come up with a design for a hybrid approach.

4. **Thinking It's the Cure All –** It's clearly not. As we stated in false assumption #1, certain systems aren't a good fit for the cloud—technically or financially.

Public Cloud Use of Templates

The first surprise comes when companies are told that systems they're running have a big list of requirements that need to be fulfilled to qualify for the cloud. For example:

1. **OS Versions:** Often cloud providers will only allow one version. Public cloud providers primarily run on Linux, and although they try to accommodate Microsoft OS, it's oil and water by nature. How many applications or systems are you running on old versions and why? Do the software vendors tell you your system only runs on an older version and they won't support it otherwise? Maybe the client has other systems that are integrated and in order to communicate properly, you'll need older versions as well. Anyone who's gone through an OS upgrade understands the interdependency problem this creates and the need to work around it. Going to the cloud often means that systems need to be

upgraded in advance, so now there is an upgrade project in front of the migration to the cloud initiative. Customers either realize they need to standardize their environment first or take the huge risk of migrating what they have, as is—and then finding out what doesn't work after the fact.

2. **Virtualization:** This is a typical requirement of public and private clouds, so if the manufacturer won't support this, it's a problem.

3. **Security Limitations:** Public cloud is a shared environment and the security infrastructure will change how those systems can be accessed from the outside. This becomes especially complicated when the targeted system needs to integrate and talk to other systems outside of the cloud since, in many cases, the security infrastructure the cloud company uses will change the routing. In these instances, when you have these systems in-house under the same roof, the security zone approach can simplify matters.

Once you take one of these systems out, however, a detailed routing interdependence map needs to be developed to ensure it still works the same way. If those systems integrate or share data with other systems outside of the cloud, it may require security policies that are so specific and discrete that they render developing a functional security policy and future change control inefficient.

4. **Performance Issues:** Clients are often fooled into comparing their current specifications with the cloud providers' and thinking they're equal—only to find out after a migration that for some reason the throughput isn't the same and there are performance problems. From network, to storage, to the efficiency of their internal network, the challenge can be hard to identify. For example, cloud providers often can't guarantee things like I/Os or storage performance overall. As a result, many companies have gotten big surprises thinking that storage is all the same when it comes to performance-oriented systems. Additionally, many providers don't have things like high-performance solid state integrated into their solutions.

5. **Specialty Solutions Are Not Supported:** When it comes to WAN optimization, Citrix presentation, VDI, custom load balancing configurations, custom security policy and security solutions, IDS, DLP and SEIM are often either not offered or lack the full feature set. And, often after clients have signed on the dotted line, they figure out that the little things they chose in the past are going to get thrown out. The ultimate decision comes down to what's more important—the value the specialty solution brought to the table or the value of the cloud initiative, because you often cannot have both.

6. **Compliance, Sensitive Data Requirements:** These limitations often drive clients to eliminate some systems from the scope of their cloud initiative or to consider alternatives that public cloud providers present as a solution—such as their private cloud, for example. Keeping some systems in-house or evaluating private cloud solutions are often the reaction to the limitations of the public cloud. Let's walk through the scenarios.

Private Cloud

This is recommended by cloud companies for systems that won't completely fit into the public cloud for reasons such as the ones listed above—which is the vast majority of systems. There are technical or functional limitations for most private cloud solutions, but the ROI and cost are what prevent most private cloud initiatives from succeeding.

The fact is that most "private clouds" are often a specific build for a client. The provider scopes out the client requirements, prices out an environment end to end, purchases that equipment and turns it into typically a three-year lease for

the client (under the auspices of a three-year service contract). They mark all this up 20–30% and you have basically the same solution you could have built at a co-location, but with a hefty markup. By its definition, what is built for the client is not shared, so some subtle things happen with this approach.

1. **OS Templates:** Private builds generally use templates that are hardened versions, with their monitoring and support tools built in and locked down. There are often excessive integration issues with the private cloud provider's standards that create compatibility or functional issues. Again, many clients find themselves having to perform an internal standardization project in preparation for moving to a strict environment—or they just risk performance and uptime by combining the equivalent of a data center migration and an upgrade at the same time.

2. **Loss of Incremental Scalability:** This is a HUGE consideration. Since the provider only buys what the client needs at that moment, when growth occurs, they end up having to buy another "unit" and sign another three-year contract. Eventually you will require forklift upgrades for systems out of capacity or, worse yet, you will be burdened with a whole system when in reality you may only need a small portfolio of that capacity. You end up having to pay for the excess unused capacity and pay for the provider's markup—leading many clients to wonder why going to a big provider with economies of scale is more expensive than doing it themselves. In a true multi-tenant or public cloud you need more memory or CPU—and it's simply a matter of allocating it. In a private cloud, you often have to finance a whole new system—or you pay for and deal with the logistics of upgrading a server and the associated down time.

3. **Stuck With Legacy Solutions:** Requirements change quickly now and one of the promises of the cloud is not locking into solutions that you don't know you will need one or two years down the road. In the private cloud model, when those requirements change you'll find you can't get rid of those costs. That's because the financial structure of many of the providers is such that there's no way to repurpose those technologies. And, since they aren't multi-tenant, they can't offset those costs via reallocation. In this model, clients risk the technology becoming obsolete long before the three-year financing or the five-year depreciation cycle.

4. **Loss of Economies of Scale:** Remember, private clouds are specific builds based on the client's current requirements. Clients might be in a world-class data center but are still often limited in their ability to tap into the provider's scalable solutions.

5. **Loss of Access to Mature Turnkey Solutions:** DR, for example, can be a very inexpensive solution if clients are able to tap into a multi-tenant environment. In a private cloud, the client is often responsible for purchasing every expensive infrastructure and bears the entire financial burden when the solution becomes too expensive.

6. **Loss of Specialty Skill Sets:** Cloud companies rarely offer skill sets to properly manage specialty infrastructure and, worse yet, don't allow the client to manage elements of the system at the cloud location despite the limitations in their offering.

In summary, private clouds have their place, but you need to understand why you're moving to private before you do.

Microsoft Azure Cloud

In some cases, Azure can be a useful tool. It has many features and benefits that make sense, especially when it's offered like Office 365.

Multi-Tenant Private Cloud

This option solves many of the problems associated with the flexibility of public clouds and the cost of private clouds. Very few companies have opted for this type of environment because of the complex engineering required to build a secure high-performance flexible environment that is multi-tenant and cost effective.

The General Challenges of Moving to the Cloud

Thinking that choosing one cloud partner is the right strategy

Vendor consolidation is a good tactic to solve inefficiencies with unnecessary redundant providers. Unfortunately, that same approach in IT isn't realistic. Companies have got to start getting used to a federated model. We have seen incredible money lost when the decision is made to choose one cloud provider for cost savings. Almost everyone ends up at the same spot—whether they choose Azure for Office 365 or SaaS providers for key applications that don't require customization. Unfortunately, the only strategy that works—assuming money is no object—is one where you look at each system and determine what's possible first and then draw a strategy from the findings. The result will surprise you.

Tech people who blame the technology for poor performance

Poor performance is often due to other factors that are difficult to verify. For instance, those same people may have created an excessively noisy design that does not follow best practices. Or, maybe the team lacks discipline to properly maintain those technologies.

CFO and CIO getting on the same page (hidden costs)

The big challenge is avoiding duplicate costs that occur when you combine interdependence of technology and depreciation schedules. When you put part of your operation in the cloud, a portion of the infrastructure is still "on the books," resulting in duplication of costs.

Interoperability and interdependence

Clients rarely understand internally which systems are connected or dependent on one another. In some cases clients with poor designs may have little or no flexibility and are faced with all-or-nothing decisions. This is a tricky topic where the answer isn't always obvious. Application interaction is often very subtle, and clients tend not to discover interoperability until the systems have been separated and suddenly there's a performance problem. The most common response to that is clients saying that the new provider is the problem—not realizing that separating a system from the rest of the systems and infrastructure will almost always have consequences that were missed in the feasibility phase. Think about moving a system that uses current storage, backup infrastructure, load balancers, or web firewalls, telco connections and internet access. Unless those systems are dedicated to the application or at the end of their useful life, there will be duplicate costs.

Another infrastructure element that is often overlooked is Active Directory, which is duplicated in many cases but not seen as part of the system when first considering it. Going to Office 365, for instance, will require you to choose whether you maintain an on-premise Active Directory and whether you decide to use identity federation.

If you want to continue hosting Active Directory, but don't want the hassle of implementing identity federation, then you can have Microsoft Office 365 handle domain names through a process called partial re-delegation. Your organization retains ownership of the domain name, but certain functions such as e-mail and web hosting are redirected to Microsoft Office 365 servers.

Infrastructure with a wide variety of depreciation cycles

These unfortunately don't align with the way those systems are interdependent technically, so every system being considered for the move typically is connected to and relies on many others that create all-or-nothing scenarios. This leaves you with a limited option: to move one system, duplicate those supporting systems (and costs) at the new location or provider, and retain those supporting systems at the current site. Moves to decrease complexity at first glance often end up with excessive waste in duplication or managing multiple locations.

Misunderstanding of IT costs—and not considering other variables.

Clients tend to just look at the hardware/software infrastructure costs (that is, everything except the applications) when determining their costs. Rarely do clients have a true unit cost for everything they do, including power/facilities and the human resources required for implementation.

Typical costs for running IT tend to be around 1/3 hardware/software, 1/3 people and 1/3 power/facilities. RKON has evidence based on our own cloud that the people cost may be closer to 40–45%. Clients tend to evaluate cloud decisions on factors that are easy to measure, not realizing they often account for only a small portion of the overall cost.

Migration to the cloud is complicated, to say the least.

The journey can be costly and problematic—and needs to be given careful consideration before you get started.

Clients tend to underestimate the complexity of the work and time it takes to migrate to the cloud. There are several factors that need to be planned out by an experienced data center migration professional.

Preparing for the move: Cost of standardization

Unless the cloud company will allow old versions of technology, there is usually a significant upgrade and standardization project that needs to be performed before moving to the cloud. As mentioned before, rarely will the cloud company take your infrastructure as-is. Clients often find that custom application providers won't allow new versions of operating systems without it having a major impact on cost. When they do allow it, rarely are clients at the current versions. Which means they are taking a tremendous risk moving to the cloud while migrating at the same time. The other option is to go through an upgrade/remediation before moving, which is not something cloud companies tend to advertise.

Lack of turnkey solutions

Service providers won't take responsibility for migration or the end result. That is often left up to the customer to figure out. Most providers don't have migration teams, but even when they do, there's no design team that is brought in to do an analysis to see if the new system will work. It's rare when a client can afford to hire a competent team to run the IT organization AND the right technical team to oversee the decisions to see if it will all work. This requires a very mature architecture team with significant experience in all areas of infrastructure—and one that has worked with diverse multi data centers and multi cloud providers. Companies' default thinking is that it's the provider's job to get it up and running since they will make money off the system over the next three-year contract—and that they should set it up for free (which is never the case). Even when the provider can help with the transition, they rarely can help on the client side.

"We trust our people"

We do too, but as Reagan said, "Trust but verify." Let's face it, loyalty has gone by the wayside in corporate America. The typical CIO's lifespan is two years. and demand for IT people is still so high that anyone can pack up and find a new

job at the drop of a hat or the first sign of trouble. The challenge that everyone brings their own skill sets and knowledge to the table, based what he/she perceives is the "right way" of doing things. Teams turn over, don't understand or agree to the last team's plans and end up changing direction before those solutions are finished—wasting that up-front investment in the unfinished project.

In short, when unchecked, companies end up with the long-term impact of short-term thinking. The layers build on these bad decisions until the team is left with no incremental budget for new initiatives while still having to manage a complex network of half solutions that have little hope of delivering ROI but remain integral and immovable elements of the network. Knowing so many data center initiatives take much longer than people think to complete (often two to three years before ROI is achieved), it's not uncommon for clients to end up with partial, uncompleted solutions across two or three data centers or service providers. That, in and of itself, becomes the complexity of future decisions: how to untangle the mess and the immovable stream of the underperforming expenses that have been created.

Solution

The end vision is easy to understand, but it's the migration in getting there that's fraught with problems as described above. It's important to come up with a long-term strategy that has some flexibility built in. Most of the focus will be on those systems or elements that are connected to the target system but still have useful life left—in which case, moving early would incur duplicate costs, and what do you do in the interim?

The difficulty is figuring out a way to gracefully move to cloud when the opportunity makes sense from a technology interdependence and financial depreciation cycle standpoint—and to find a provider that can take ownership and accommodate each phase of the process.

The ideal solution is to find a cloud company that can:

- **Help** you with legacy infrastructure not quite ready to move to the cloud—offering assistance in remediating existing infrastructure and managing it at the client site until it can be moved

- **Offer** the economies of scale combined with flexibility—taking your existing standards without moving them into a private cloud-dedicated customer build

- **Assist** in preparing and moving systems—and take ownership for the migration