techworld FROM IDG Insights

Issue: **October 2019**

THE COMPLETE GUIDE TO CLOUD MIGRATION

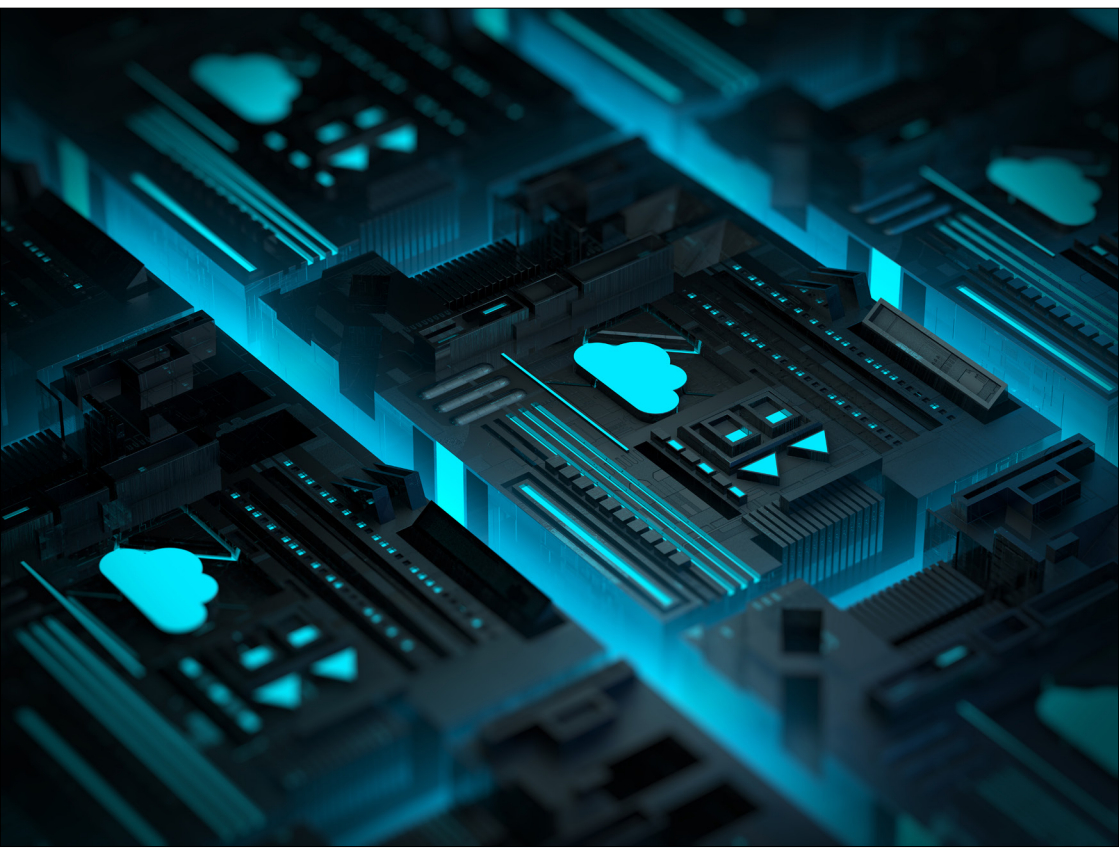

CREDIT: ISTOCK

Produced by IDG

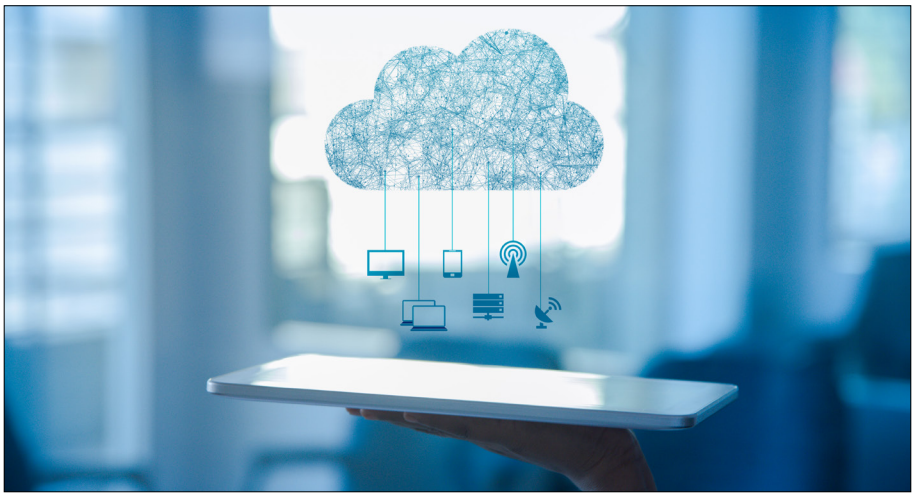

Credit: iStock

Cloud cover

Since the advent of the term ‘cloud computing’ back in 1996, when it popped up in an internal document at Compaq, the distributed computing technology has been on the lips of everyone from giant technology vendors to business and IT leaders and the founders of world-changing companies such as Uber, Netflix and Facebook.

Ever since, the major cloud vendors have been putting their marketing machines into overdrive to convince enterprises of the benefits of cloud computing – flexibility, cost reduction, near limitless scale – and entice them to move their precious applications and databases into their data centre farms. However, according to IDG’s 2018 Cloud Computing Survey, the average enterprise environment is still predominantly non-cloud (53 percent), so there is still work to be done.

That process of migration isn't as simple as it is often made out, even if Amazon will literally load all of your data onto a truck and move it for you through its Snowmobile service. Companies that have successfully navigated a cloud migration, whether that's a 'big bang' move or a more phased approach, do preach of the benefits in setting their organization up for the future.

The focus of this guide is to bring together a selection of real-world stories of enterprise cloud migrations, from all-in moves to key databases or applications here and there to test the waters. All of these case studies include key lessons learned and some of the biggest benefits and pitfalls of performing a cloud migration. Scott Carey

Contents

Credit: iStock

Eleven years on

How Netflix completed a historic cloud migration with AWS

In its 20-year history, Netflix has grown from a DVD rental website with 30 employees to a global streaming service with over 5,000 titles, 130 million subscribers and $11 billion annual revenue that has drastically transformed the entertainment industry.

At the Consumer Electronics Show in January 2016, Netflix CEO Reed Hastings launched the service in more than 130 countries.

"While you have been listening to me talk, the Netflix service has gone live in nearly every country of the world," he announced on the Las Vegas Convention

Centre stage. “Today, right now, you are witnessing the birth of a global TV network.”

This was all made possible by a radical transformation of a previously traditional IT operation, as Dave Hahn, a senior engineer in Netflix’s cloud operations and reliability team, explained at 2018’s Service Desk and IT Support Show in London.

“We flipped on the service for another 130 countries, and millions of new customers that we hadn’t previously been servicing,” he revealed. “I think you can imagine the amount of work and thinking and architecture design we had to do to open up to 130 countries, and millions of new customers just in that moment; the technical architecture, the research, the billing systems, the kind of people that we needed and the thinking about these kinds of problems in order to make that happen.”

The journey began when Netflix decided to move from its own data centres to the public cloud.

Migrating with microservices

In 2008, Netflix was running relational databases in its own data centres when disaster struck. A data centre failure shut the entire service down and stopped DVD shipments for three days.

The company’s owners faced a choice: turn Netflix into a world-class data centre operations company or move the service to the public cloud.

Netflix was growing fast. The thousands of videos and tens of millions of customers were already generating an enormous quantity of data. The company would struggle to rack the servers in its own data centres fast enough to handle the ever-growing volumes, but the

cloud would let it add thousands of virtual servers and petabytes of storage within minutes.

A migration to the cloud was the clear choice. Netflix soon became a poster child customer for Amazon Web Services (AWS), choosing the company for its scale and broad set of services and features.

The move would require a complete rearchitecting of the company's traditional infrastructure, though. Netflix could have forklifted all of its monolithic enterprise systems out of the data centre and dropped them into AWS, but this would only have brought all of its old data centre problems to the cloud. Instead, it chose to rebuild the Netflix technology in AWS and fundamentally change the way that the company operated.

"Software's like anything else; if you can design it for the environment that it's going to be living in, it will do more of the things you want it to do, more often and more regularly," Hahn explained. "So we chose to move to microservices."

This made the infrastructure much more agile by breaking aspects of the service up into multiple microservices, managed by Netflix's own small teams who understood how their service worked and interacted with other systems. This was pretty groundbreaking at the time.

This provides clear, specific insights that make it easier to change the service, which leads to smaller and faster deployments. It also allows them to isolate services to understand the various performance profiles, patterns and securities in each microservice, and move away from any individual piece that's causing a problem.

"I don't have to assemble all of these pieces built by other people in order to have a singular deployment,"

Hahn said. “Any Netflix service team can deploy their service at any time. It requires no coordination, no scheduling, no crucible to get to production.”

Benefits of the cloud

It took Netflix seven years to complete the migration to the cloud. In 2016, the last remaining data centres used by the streaming service were shut down. In their place was a new cloud infrastructure running all of Netflix’s computing and storage needs, from customer information to recommendation algorithms.

The migration improved Netflix’s scalability and service availability, and the velocity by which the company could release new content, features, interfaces and interactions. It also freed up the capacity of engineers, cut the costs of streaming, drastically improved availability, and added the experience and expertise of AWS.

“The other thing is that the cost model is really nice for us,” Hahn added. “You pay for what you use. That allows us to do a lot of experimentations.”

This gives it greater freedom to test new features and improve existing ones, such as the rows of content recommendations that are personalized every day.

“These large recommendation algorithms require a lot of compute work,” Hahn explained. “If I want to find out if a new one we’re playing with does better, I don’t want to turn off the old one because you still need recommendations. I can now spin up an entirely new set of machines in the tens, or hundreds or thousands in an afternoon, and chunk through my data and see if we’ve done better. And I only pay for the portions I use. It allows us an amazing amount of freedom in experimentation.”

Content delivery

The cloud is only one part of the Netflix user experience. Everything that happens before they hit play takes place in AWS, but the video content that follows comes from a separate system: Netflix OpenConnect, the company's proprietary content delivery network (CDN). The OpenConnect appliances store the video content and deliver it to client devices.

CDNs are designed to deliver Internet-based content to viewers by bringing it closer to where users are watching. Netflix originally outsourced streaming video delivery to third-party CDN suppliers, but as the company grew, these vendors struggled to support the traffic. It needed more control over the service and user experience.

The company decided to design a CDN tailored to its needs. It now installs OpenConnect appliances that store and deliver content inside local Internet service provider (ISP) data centres, which isolates the Netflix service from the wider Internet. Popularity algorithms and storage techniques help distribute the content in ways that maximize offload efficiency. The system reduces the demand on upstream network capacity and helps Netflix work more closely with the ISP networks that host its traffic.

"We designed OpenConnect caching boxes to hold our content and, wherever we can, we install them inside of your Internet service provider's network, so that when you see those video bits you aren't actually transiting off of your operator's network," Hahn revealed.

The new system cut the appearances of the loathed buffering wheel by an order of magnitudes. It also allowed Netflix to make the CDN software

more intelligent. Now, whenever a customer presses play, their device can get its content from numerous places on the Internet.

The investment paid off when a fire in an ISP data centre in Brazil burned down Netflix's entire stack of machines. Customers who had been streaming from the ISP didn't experience any change in their user experience. "Their devices already knew somewhere else to go get the data," Hahn explained. "It didn't interrupt even one frame of streaming, when we literally burned down the base."

Chaos engineering culture

Netflix developers are well known these days for their unique approach to engineering culture. A self-service chaos engineering tool called the Chaos Automation Platform was pioneered to test problems in their production environments, so they can be sure that their software will behave as they want during a failure.

"People press the play button on Netflix thousands of times a second," Hahn revealed. "If the systems cannot auto recover, if they cannot handle bad situations, if they cannot self-repair, by the time I get a human involved, in the best-case scenario, minutes have gone by.

"You can get an idea of how many of our customers we've disappointed in the three or four or five minutes it may take to get a human involved, and in the right place and working. Chaos engineering is an excellent inoculation to failures."

They use the chaos engineering method to ensure Netflix can survive a failure in one of three AWS regions it uses. Every month, they turn off one of the regions and test that they can move all the customers that it was

serving to another one within six minutes. To embrace chaos without causing destruction, Netflix had to create a corporate culture that supported such ideas.

The central principles were formalized in the 127-slide Netflix Culture Deck, which Facebook COO Sheryl Sandberg said “may well be the most important document ever to come out of the Valley”.

A central tenet of the policy is balancing freedom with responsibility. Teams are given ownership of their microservice and encouraged to act independently, but not recklessly.

“Netflix managers do not set out tasks for their employees to do or design their projects,” Hahn revealed. “Their job is to give them the appropriate context, so they can make the decisions to hire excellent, stunning colleagues for them to work with, and to stay out of their way.”

The company avoids making too many rules beyond a set of fundamental principles such as never touching customer data. Hahn described the approach as building guardrails but not gates, and claimed he can count on one hand the number of times he’s had to tell an engineer exactly what to do.

“By making sure that that context is widely and regularly shared I can have someone design a billing system, someone else working on algorithmic systems, SREs on our reliability teams, and somebody else working on customer service, and they’ll understand the same context and march towards the same goals,” he said.

“That allows us to keep those teams loosely covered. We don’t have lots of structures and controls, but we keep them highly aligned.” **Thomas Macaulay**

Credit: Spotify

Google Cloud hits the right note for Spotify

How Spotify migrated from on-premise to Google Cloud Platform

Spotify announced that it was going all in on Google Cloud Platform (GCP) back in 2016, committing a reported $450 million (around £343 million) over three years. In Spotify, Google got itself an anchor customer, not just because of its brand and scale, but also its reputation as a data driven, engineering-centric company.

It has since shut down both of its US data centres and will be free of on-premise infrastructure by the end of the year following a complex migration. During 2018's Google Cloud Next conference in San Francisco we

heard from members of the Spotify and Google Cloud teams who were involved in the migration about how they went about it and some key lessons learned.

Why migrate?

Director of engineering at Spotify Ramon van Alteren started by explaining why the company had decided to go all in on cloud infrastructure in the first place.

"If you think about the amount of effort it takes to maintain compute, storage and network capacity for a global company that serves more than 170 million users, that is a sizeable amount of work," he said. "If I'm really honest, what we really want to do at Spotify is be the best music service in the world, none of that work on data centres actually contributes directly to that."

As well as freeing up developers from worrying about provisioning and maintaining infrastructure, van Alteren said the company also wanted to start taking advantage of some of the innovation coming out of Google Cloud, specifically the BigQuery cloud data warehouse, Pub/Sub for messaging, and DataFlow for batch and streaming processing.

Van Alteren also said the driving force behind the move to cloud came from the engineers tasked with maintaining those data centres. "A big part of that was asking what their job looks like once we move to cloud," he said. "The net result was that group of engineers, some of the most deeply respected at Spotify, ended up being advocates for our cloud strategy."

Services migration

The migration plan was formulated back in 2015 and was broadly split up into two parts: services and data.

The services migration focused on moving nearly 1,200 microservices from on-premise data centres to the Google Cloud Platform.

The three main goals during migration, according to van Alteren, were to minimize disruption to product development, finish as quickly as possible to avoid the cost and complexity of running in a hybrid environment, and cleanup, ensuring Spotify didn't have any lingering services running in its data centres.

One of the first things Google and Spotify did was to build a small migration team of Spotify engineers and Googlers, and build a live visualization of the entire migration state, so that engineers could self-serve to see the progress of the project.

That visualization looks like a set of red (data centre) and green (Google Cloud) bubbles. Each bubble represents a system, and the size of the bubble represents the number of machines involved.

"This had a number of interesting side effects, one that saved me a lot of time as a programme manager to save doing status updates," van Alteren said.

"Next, it created a real sense of accomplishment for teams that were migrating and they could see the impact they were making."

The services migration started with mapping dependencies, as the architecture at Spotify means every microservice relies on somewhere between 10 to 15 others to service a customer request. This means that a 'big bang' migration, where everything stops, was not an option as customers expect constant uptime from the service.

Instead, Spotify engineering teams were tasked with moving their services to the cloud in a focused

two-week sprint, where they effectively paused any product development. This also allowed those teams to start assessing their architecture and to decommission anything unnecessary.

One thing Google Cloud brought to the table specifically for Spotify during the migration is its Virtual Private Cloud (VPC) option. “This allows you to build similar to an internal network which connects multiple projects and they can cross talk,” van Alteren said. “This gives teams good control of their own destiny, they get to do what they need to for their service and if they shoot themselves in the foot they shoot themselves in the foot and not the whole company.”

The second blocker was latency caused by the Virtual Private Network (VPN). As it started the migration, Spotify found that shifting 1,200 or so microservices took up a lot of VPN bandwidth.

“To be honest, the VPN service at that time was more or less scaled for an office of 25 developers,” van Alteren revealed. “When we showed up with four data centres that didn’t work so well, we collaborated with Google and got this to work solidly pretty quickly. So we built multiple gigabytes of network pipes between our data centres and Google Cloud to get that dependency problem to disappear.”

Once these blockers had been removed Spotify was able to start moving user traffic over to the cloud. “Another key realization in the service migration was that we could decouple service migration from moving our user traffic,” van Alteren said.

“So we deliberately separated these road maps to focus on getting these applications ready to run on Google Cloud and a separate road map that allowed

us to gradually connect more and more users with GCP, which allowed us to control the reliability, user experience and our migration speed, and how much traffic was flowing over these VPN links as well."

Once that migration was in full flow, the central migration team started to secretly induce failures in those cloud systems, recording how the teams reacted and recovered on the new architecture.

Peter Mark Verwoerd, a solutions architect at Google said: "While that was fun to break things and see teams scramble, it helped ensure the monitoring systems were properly extended to the new cloud deployment, if a team didn't notice, that would be a big red flag. Finally, we had this playbook they could start going to for failure modes in the cloud they may not have had in the past."

By May 2017 each migration sprint had been completed and traffic was being routed towards Google Cloud. Then, in December 2017 Spotify hit 100 percent of users and had already closed its first of four on-premise data centres. Shortly afterwards, the second data centre was closed and the final two, both in Europe, were closed down towards the end of 2018.

That road map is "a pretty strong signal for people with long-tail applications still running in data centres that they need to get a move on," van Alteren added.

Data migration

Next up to talk through the data migration was Josh Baer, senior product manager for machine learning infrastructure at Spotify, who described the experience of moving one of Europe's biggest on-premise Hadoop clusters to the cloud. Due to a highly complex dependency graph it was a challenge to move 20,000

daily data jobs to GCP without causing downstream failures, according to Baer.

Spotify started by assessing the possibility of a 'big bang' migration. "Shut down the Hadoop cluster, copy all of the data over to GCP and start things back up again," he revealed.

Unfortunately, even with a 160 gigabit per second network link, it would have taken two months to copy the data from the Hadoop cluster to Google infrastructure. "We wouldn't be much of a business if we were down for two months," Baer added.

The strategy Spotify landed on was to copy a lot of data. "As you moved your job over to GCP you would copy your dependencies over, then you could port your job," he explained. "Then if you had downstream consumers, you may have to copy the output of your job back to our on-premise cluster so they weren't broken. As the bulk of our data migration lasted six to 12 months, we were running a lot of these jobs to fill gaps on our dependency tree."

Naturally a migration like this ate up network bandwidth, so Baer and his team learned quickly to over-provision and to avoid using VPN whenever possible.

Each migration sprint started with two options for the team involved: they could lift and shift – something they called 'forklifting' – where appropriate or time-poor; but ideally they would rewrite.

"This was useful for teams that didn't feel comfortable just porting over their jobs using the forklifting path because they may have inherited these data jobs and hadn't really looked at them, and if they were going to dig into them they might as well rewrite them," Baer said. "The biggest thing with rewriting was it required a much

larger time investment from teams and as engineers what usually happened is as they started writing it they want to rearchitect it too. Towards the end and middle of our migration we had to tell people to stop the rewrite path, just migrate their stuff, and if they really wanted to rewrite it, then it was already on GCP, so we could still hit our migration targets."

Spotify is now running its data stack entirely on BigQuery, running 10 million queries and scheduled jobs per month, all-in-all processing 500 petabytes of data.

Lessons learned

Max Charas, a strategic cloud engineer at Google warned: "This migration strategy is very tailored to Spotify both technically and organizationally, so if you wanted to do something like this it might look very different." That being said, some key lessons were learned from the migration.

The first was to prepare. Charas said: "We prepared for probably two years before the migration and each migration took around a year. We tried to build a minimal use case to show the benefits of moving to GCP, but that couldn't be a small thing to show the true value."

Second was to focus. Van Alteren said: "It is truly amazing what you can do with team of engineers focused on a single thing, we had sprints of a week moving 50 to 70 services. It will also help your business stakeholders, who will be happier with a short period of time with no product development instead of a long time. If you try to do other things at the same time, you will slow your migration down to a crawl."

Third was to build a dedicated migration team to "act as guardrails to help them know what they need to

know, pass on past experience and learnings and just be resources they need", Charas said.

Last was to "get out of hybrid as fast as you can – all these copy jobs are expensive and complex", Baer said.

Results

The results for Spotify have been more freedom for developers and greater scale, without sacrificing its quality of service.

"Quality of service is something we measured diligently and there has been no degradation there," van Alteren argued. "The benefits we derived includes our event delivery pipeline, that carries our royalty payments for rights holders, but is also a core part of our product development. When we moved to cloud that pipeline was carrying at peak 800,000 events per second, look now and we carry 3,000,000 a second, having that much more information available for product development is insane."

And cost savings? "This is a key thing to keep an eye on as we move from a centralized buying position to a distributed buying position where everyone is capable of spending money for your company," van Alteren admitted. "So it depends. Currently, we have grown in size so it is really hard to compare and I can't give you figures." Scott Carey

Credit: iStock

Why GoCompare is going all-in on Microsoft Azure

GoCompare's owner is in the process of moving its price comparison site to Azure in an all-in migration, with help from Rackspace

Popular British price comparison site GoCompare is in the process of moving its entire technology stack to the Microsoft Azure cloud in a bid to better scale for unforeseen peaks in demand, and start to leverage data science tools to make its comparison and switching products more effective for customers.

GoCompare has evolved into more than its insurance comparison brand fronted by infamous operatic mascot

Gio Compario, played by Welsh tenor Wynne Evans. The GoCo group now includes the energy auto-switching service weflip and Look After My Bills, which it acquired for £12.5 million in July this year, building out its portfolio of financial services comparison services.

The group decided to take a phased approach to the migration, with the help of partner Rackspace, starting in September last year. As a company wholly reliant on its website, the key to any migration was maintaining uptime and minimizing risk.

"We have been fortunate to have a stable platform over the years and don't suffer from reliability or uptime problems, so it was important to maintain that and not introduce risk," Kieron Nolan, group director of technology at GoCo Group told us. "That incremental approach was about maintaining that stability and lack of disruption."

The initial delivery of a full disaster recovery replica in Azure and the migration of staging and development environments to the cloud platform was completed in March this year, and the full production platform is due to move over by the end of the year.

Rackspace provided a consulting role throughout the process. "We wanted an experience from a partner that had done this before and help us make the right decision and understand the different options out there," Nolan explained.

Major benefits

Once that work is done the organization can start to take advantage of Azure services such as Databricks and other data science products to make its services more predictive and intelligent. For example, this could

provide customers with products that find the best insurance policy for their needs, or the best energy deal depending on their previous usage.

Shorter term, the business has seen an immediate impact on its website response times, which have improved by as much as a second, as well as boosting its ability to scale flexibly in response to unforeseen demand for the site.

For example, if the 'Money Saving Expert' Martin Lewis were to mention a GoCompare service on his ITV show, the business could see a Monday evening spike in traffic. "Those historically we struggled with, but with Azure we have the elastic model to be better placed to take advantage of that spike," Nolan said.

What vendor lock-in?

In the current climate of hybrid/multicloud hype, an all-in cloud migration could be seen as unusual, but it wasn't a concern for GoCo.

The organization was heavily a Microsoft shop already, with a .NET tech stack and SQL back end, and plenty of related skills in the building.

"When we looked at migrating to the cloud, from an architecture and migration point of view, Azure made the most sense," Nolan argued, "to evolve that stack and not re-platforming anything. We wanted to get to the cloud, within reason, with the architecture intact and tweak elements at our own pace."

The wider group does have different cloud providers and skill sets, however, for example My Voucher Codes is built on an AWS stack, as is weflip.

"The strategy we have is from a disaster recover and continuity point of view we build in Azure," Nolan

said. Then, group-wide the “strategy for resiliency is to have brands with different providers, that is the way we approached this, so if there was a major outage it wouldn’t take out the whole group. At a brand level that mitigation is within the stack itself.”

Nolan admitted that GoCompare is “heavily invested in the Microsoft stack, but from our point of view that isn’t a problem”.

Lessons learned

By all accounts the migration was a pretty seamless experience, but that doesn’t mean the organization didn’t learn any important lessons.

Primarily, Nolan was glad that GoCompare and Rackspace decided to focus on the migration, and not reinventing its stack. “Deciding specifically not to re-architect everything as part of the migration was a key decision and allowed us to focus on the migration and not re-engineering,” he explained. **Scott Carey**

Credit: iStock

Travel giant migrates its search app to the cloud

Why Amadeus moved its core Master Pricer app to Google Cloud

The Spanish travel tech giant Amadeus has moved its core Master Pricer travel search application – which powers metasearch engines likes Expedia and Kayak and processes up to hundreds of thousands of transactions a second – to the Google Cloud Platform, in the first step of an ambitious cloud migration strategy.

Speaking at the 2019, Google Cloud Next conference in San Francisco, Denis Lacroix, SVP of core shared services R&D, explained that the decision to move such an important application to the cloud came through pure necessity.

"Back in 2017 we started to see a possibility that we might be running out of space in our data centre," he said. "Master Pricer is compute-intensive and uses a lot of machines, a lot of floor space, and is growing in traffic 30 percent every year, so pretty much exponential growth. We could see that within three to five years we would be out of space. So we had two choices: either we pour more concrete and extend the data centre, or we actually move the workload to the cloud."

Migration

His colleague Dietmar Fauser, SVP of technology platforms and engineering, added that of all the Amadeus applications, Master Pricer comprises "modern distributed Linux applications with in-memory data management technology, so there is no Oracle database underneath, which is why they were relatively easy to be containerized and put on [Google Cloud]".

During the migration, Amadeus had engineers from both Google and Red Hat on site to help them get to grips with OpenShift and the container orchestration technology Kubernetes. "We adapted Kubernetes quite massively to our needs," Fauser added.

So within 18 months Amadeus had lifted and shifted Master Pricer to the Google Cloud Platform, but why did it opt for Google and not, say, Amazon Web Services (AWS)?

"First of all it's a non-exclusive deal; we intend to also partner with other cloud providers," Lacroix revealed. The reason they opted for Google Cloud in this instance is twofold: a historic collaboration with Google due to its feed for Google Flights and the vendor's expertise around Kubernetes.

"Kubernetes runs simply very well on Google, it comes from Google. So on AWS and Azure there is a little bit more attention needed to ensure that it's really smooth," Fauser added.

And why not go all-in on Google Cloud in that case? "[There is] fear of lock-in and also just commercial negotiation 101 really, to keep everyone awake," Lacroix explained.

Cloud benefits

Amadeus said the key benefit of this shift is, of course, cost savings, which are primarily achieved through the greater flexibility of cloud compared to running on premise.

"You really have to spend a bit of time to understand the economics of the platform you're working with, the levers you can use, how quickly you can scale back your footprint after a campaign run by an airline or Expedia, for example, and the faster you scale back the cheaper it gets and you can see that the day after," Lacroix said.

"Currently, we actually size things up for peaks. These peaks happen around Black Friday, January and Easter, but in between you have hardware just sitting there, just doing nothing," he added.

Fauser added that cloud also forces his developers to be more "dollar aware" as they now get instant feedback as to the cost of any change.

That doesn't necessarily mean going to the cloud is cost neutral, Fauser revealed. "You have to go with relatively high volume commitments to get a competitive price... on Google the bill only from log activity is reaching close to a million [dollars] a year already," he said.

Another benefit is latency. “Some of our large customers are based in the US or Australia, and Singapore. So when we serve them from our data centre, we have latency up to half a second,” Lacroix explained. By running Master Pricer in Google regions closer to these customers, Amadeus can quickly improve that search latency.

The last benefit is more indirect and helps with recruiting and retaining talent because “working on these type of things is so much more sexy to advanced engineers than working on mainframe”, Fauser added.

Cloud strategy

In terms of the broader strategy, Amadeus is in the process of retiring all of its mainframes in a bid to go 100 percent cloud.

“The move is a major step forward in Amadeus’s vision for cloud-based architecture that is globally distributed across both the private and public cloud, driving more speed, responsiveness, and resilience in global operations. In addition, Amadeus is now progressively migrating its shopping applications to run on the public cloud across multiple regions. This enables Amadeus to scale its system capacity faster and on-demand, which is particularly key in moments of peak travel,” the company said in a press release.

Next will be a case of moving other applications to the cloud, whichever vendor that ends up being with.

“The data that Master Pricer works with, airline fares data, is read-only data ... So that’s like a fairly easy use case. The other extreme is if you have a super large Oracle cluster. You can move that to the cloud, but it requires a lot more work,” Lacroix said.

"What is important now is to take on the engineering efforts to bring more complex applications, with more messy data management [to the cloud]," Fauser added. "The real problem is the data management – it's not so much the compute tier. You get VMs everywhere, but running complex, synchronized Oracle operations is a different ball game." **Scott Carey**

Credit: Spirit

How Spirit overcame cloud migration challenges

Lessons learned from an S/4HANA cloud migration by Spirit Airlines

Budget US airline Spirit is in the process of moving from SAP's legacy ECC enterprise resource planning (ERP) software to S/4HANA public cloud, but it has had to overcome some challenges on the way to what it says is a very different system for its users to get to grips with.

SAP doesn't break out its public cloud vs on-premise numbers for S/4, which now counts more than 10,000 customers, but only a small minority so far have publicly opted for the SaaS version of the vendor's next

generation ERP system, instead of staying on-premise or in a single tenant private cloud. This makes it all the more useful to hear the perspective of a customer that is migrating to the public cloud version of S/4, especially as the vendor continues its efforts to push customers towards the public cloud where possible.

Speaking during SAP Sapphire in Orlando earlier this year, Spirit was upfront about some of the challenges the public cloud version of S/4HANA brings with it, primarily hinging on change management and integrations to third-party software, as well as the benefits of modernizing processes and moving away from a customized world.

Why S/4HANA public cloud?

The airline has been an SAP customer for a long time, and has been running its original instance of ECC without a single upgrade since 2013. Recognizing that ECC support is being dropped by SAP in 2025, the airline decided to modernize ahead of time.

The move to S/4HANA public cloud hinged on several motivations, including a reduced IT maintenance burden, streamlining its financial reporting cycle, and moving away from highly customized business processes to more best practice.

“The more you customize, you are moving away from best practices. Just because it is good for the user, doesn’t mean that its best practices,” SAP BI architect at Spirit Airlines, Pani Pothur, explained during the breakout session. An example of how Spirit has already improved here is that it has already reduced the number of cost centres from 600+ to 150+, and has consolidated from 20 company codes for invoicing, to just one.

Lessons learned

On the technical side one of the challenges Spirit faced was the lack of integrations for some third-party software it was reliant on, such as OpenText for invoicing and an Excel plug-in called Z Options.

After seeing that the invoicing capabilities in S/4 probably didn't scale for the airline, it also decided to procure SAP Ariba in tandem, opting for the SaaS solution because of its natural integration with S/4.

"One of the key pieces of advice I would give is to take your time," Pothur said, "because the S/4HANA public cloud is a completely different architecture. So you will have to make sure that you do your due diligence before starting your journey because all of your third-party applications, your integrations, you have to focus on them all working with the S/4HANA public cloud."

He also urged customers to get their master data in order ahead of time. "You have got to make sure you clean up the master data," he added, "you don't bring the garbage into the new system."

That being said the main challenges, as always, related to change management.

"Since it is going to be different to ECC you have to make sure your key stakeholders are involved in each and every engagement process to have a successful implementation," Pothur explained. "The screens are different, the processes are going to be different, you need to make sure that your key stakeholders are involved in all the different phases, from explore to testing and documentation."

In terms of training Pothur, admitted Spirit could have taken more time to get this right. "You have to spend more time," he said. "It's completely different, you cannot

really employ the same process from your older system, so you have to use best practices, so ensure you have enough time to train everyone."

These lessons were hard learned for Spirit, which delayed its S/4 implementation date after the business realized it wouldn't be ready. "We started our journey last year, but because of all the things that we have learned so far we had to pause for a few months and re-strategize," Pothur admitted. The airline is now due to go live with S/4 public cloud in October of this year.

Benefits

In terms of the benefits of moving to S/4HANA public cloud, Pothur identified a whole host of efficiencies he hopes to achieve later this year.

First, the old ECC system required a lot of manual work to get financial reporting done. As S/4 is built on the in-memory data store HANA, data is processed in closer to real time, the finance team will have the ability to do a 'soft close' ahead of time.

Moving to S/4 public cloud also reduces the maintenance effort required from IT. "In terms of the amount of effort, the amount of maintenance activities, the upgrades and everything to do, it is like a really time-consuming," Pothur said. This has total cost of ownership implications for the business also.

"Moving to the cloud will definitely reduce the time, effort, maintenance and cost for your company, because everything will be taken care of by SAP," he added.

Lastly, Pothur wanted employees to have better mobile access to the ERP system when they are on the move, with Ariba's mobile capabilities also being heralded as a major improvement for the business. **Scott Carey**

Credit: Asics

Asics to bring runners better workout data

The sportswear company turned to the cloud data warehouse after its Amazon Redshift instance started to creak

Asics Digital, the data-centric subsidiary of Japanese sportswear giant Asics, recently shifted its data warehouse from a creaking Amazon Redshift instance to Snowflake's cloud-native solution in a bid to better support its growing portfolio of digital services, such as the Runkeeper and Asics Studio apps.

The 70-year-old Japanese brand acquired the GPS fitness-tracking company Runkeeper in 2016,

forming part of a digital strategy that mirrors that of its competitors like Adidas, UnderArmour and Nike, all of which offer a run tracking application of their own.

Combine this with the homegrown Asics Studio app – a subscription service which provides users with a range of curated workouts – and the company clearly sees value added digital products as part of its future direction in the lucrative health and fitness space, which is estimated to be worth £5 billion in the UK alone.

How it got there

Back in 2015 Asics was running much of its business on Amazon's Redshift data warehouse, but as its data volumes continued to rise the company was quickly running out of storage.

"We had to keep taking clusters down for maintenance to increase the size or clean out the data we thought we weren't going to use," Chris Drouin, manager of advanced analytics at Asics Digital said.

Add to that the sort of concurrent user issues that other Redshift customers complain of and the IT team started to look around for alternatives.

Snowflake was built to offer customers almost limitless scale and concurrency by effectively spinning up new cloud instances (S3 on AWS, for example) for each workload to effectively run as a standalone data warehouse, but all under the same roof, so data-science queries never tread on the toes of BI, for example.

Drouin said that the "migration wasn't perfectly simple, but as Redshift and Snowflake primarily source data from Amazon S3, for some parts of our infrastructure it was as simple as getting Snowflake to start picking that data up.

"In other cases we changed some load processes, because Snowflake gave simple and flexible options there, such as handling semi-structured data. We have a fair bit of information that comes in JSON blobs and Snowflake can read those natively as a data type and access them with near-native performance levels."

Asics moved its first set of production data to Snowflake in 2016 and ran both systems in parallel until the organization was confident enough to commit to a full migration.

What are the benefits?

Before Snowflake, Drouin said employees could be waiting in a virtual queue for queries to end due to the lack of ability to run concurrently.

"Snowflake solved a lot of those problems for us by decoupling storage and compute," he said.

The company started by spinning up a Snowflake instance for its extract, transform, load (ETL) warehouse, which is relatively small but always on. Next was the business intelligence (BI) warehouse for reporting, which powers a set of internal Tableau dashboards.

The biggest benefit so far from a pure operations point of view is the reduced maintenance obligations on the IT team at Asics.

"We spent a lot of time on maintenance with our previous solution," Drouin explained, "so not having to schedule periods of downtime and subsequently backfilling, or making good on missing data that would have come in, has been a huge benefit."

Naturally this has cost and productivity implications, such as saving a member of the team from overseeing a maintenance period scheduled for a Saturday morning.

What next?

By running workloads in their own 'standalone' data warehouses, data scientists at Asics are now able to run ad-hoc queries without stepping on those operational workloads. The Asics Digital business unit now stores 17 terabytes of data with Snowflake, including data points from 50 million users across 180 countries, from GPS run information to buying history.

"We have been trying to build what we call a single view of the customer and that was more challenging in the past when we didn't have all of our data in one place," Drouin said. "Now that we do, that enables us to quickly spin up an instance and see the history of our relationship with you as a customer."

It also allows Asics to start delivering the accumulated knowledge it has on runners back to them in the form of useful insights. For example, incorporating third-party weather data into the Runkeeper app to nudge users to run at a different time, or consider an alternative workout if a storm is on the horizon.

Asics isn't the first fitness company to turn to Snowflake for this sort of thing either. We spoke to Strava earlier this year about how it was running queries overnight while its Redshift instance creaked under the weight of its growing data volumes. The San Francsico-based company stores more than 120TB in Snowflake today.

Strava is also using Snowflake to allow its data scientists to produce things such as its global Heatmap or to quantify effort through heart rate, or optimize its Grade Adjusted Pace metric. This is the sort of thing that makes Drouin most excited. Now Snowflake can do much of the heavy lifting on the back end, the advanced

analytics team at Asics can start to “find new sets of data to bring to customers and provide insights into their workout and gear performance”, he said.

“Those are the things on our radar. There will be new things in the fitness apps such as Runkeeper and Asics Studio to power new features and functionality in those apps,” he teased. **Scott Carey**

Credit: Centrica

Energy giant turns to Microsoft and the cloud

Centrica shuts down UK data centres in move to the cloud

The UK and Ireland arm of energy giant Centrica has finished migrating its enterprise systems and data from an on-premise data centre to Microsoft Azure, shifting to Office 365 and Dynamics 365 in the process.

The move to cloud services was announced by Microsoft back in June 2017, and the company predicted completion by the end of the year, or Q1 2018 at the latest. This included moving employee documents and data to OneDrive and Sharepoint, migrating all

mailboxes, and running Dynamics CRM on a hybrid model for the next few years.

Speaking to us at Microsoft's Future Decoded event in London in 2017, Centrica's UK and Ireland CIO Ash Jokhoo said: "We realized that we have lots of on-premise activity in terms of files for our employees and so that was the real first movement, how to improve the collaboration between our employee base and move files to an ecosystem that makes sense for our organization."

Centrica, which is the parent company of British Gas and Bord Gáis in Ireland, had been a Microsoft shop for some time, using Office and Dynamics CRM on premise, so it "just made sense to move forward with an integrated suite" when it came to moving back office systems to the cloud, according to Jokhoo.

Under global CIO Mike Young, Centrica embarked upon a 'Digital Workplace initiative'. Speaking at the time of the partnership announcement, he said it would equip "our global workforce with Office 365 and modernize field services within DynoRod by using Dynamics 365". In terms of migration Centrica planned on moving 90 percent of its enterprise data into the cloud, according to Jokhoo. However, "customer data is a completely different proposition", he said, with that sitting in an SAP ecosystem. The CIO aimed "to move forward together with both technologies in a way that serves our customers".

The speed of migration was rapid after the announcement in June 2017. "I think we are running at two speeds in terms of migration: fast and faster," Jokhoo said at the time. "We've done 20 million files across our organization and 20-22,000 Outlook mailboxes already."

The company also moved in terms of adopting Skype for Business across the organization, racking up more than 20,000 hours of video calls in 2017.

Centrica's infrastructure and operations team were also running adoption sessions alongside their migration efforts to help employees adapt to working in a more cloud-based, collaborative way.

"We have a culture that is evolving to meet the needs of our changing customers, which means we have to have agility and the way we go to market has to speed up," Jokhoo explained.

While added collaboration is always the stated goal for organizations making a shift like this, there are, of course, cost savings. "I think there is always an upside in terms of savings when moving from on-premise to the cloud," Jokhoo said. "But there is a vast amount of investment and cost-to-achieve that needs to go through first. Although savings and cost profile does absolutely play into the strategy, the collaboration element is far more important."

Centrica was able to shut down some data centres and legacy systems, though Jokhoo insisted they had to change "the mindset also. If we are going to attract the best people to join our organization to transform and disrupt the industry we need the best tools and the Microsoft suite is absolutely a part of that."

Jokhoo admitted that none of this work would have the sort of direct impact on customers as some of the data mining being done within British Gas, or the IoT innovation of Connected Home, as both have given customers better insight into their energy use, but there are non-direct impacts this work can have on customers. "We are in an extremely competitive environment and if

we can get the cost-to-serve where it needs to be, we can pass those savings on in terms of prices for services and energy," he said. "That's really the key here and the long-term play as we go through this transformation."

Additionally, Jokhoo predicted Centrica's ability to serve customers would improve as employees have access to more dynamic tools. "Clearly there will be a more agile customer service organization on Office 365 and a more agile field force organization on the Dynamics platform, so service will be better," he argued.

What next?

Centrica is already eyeing some of the advanced capability a shift to Azure could bring the business.

Jokhoo mentioned how "intelligent edge and being able to compute at the edge is going to be a fascinating journey for everyone in this industry. I see that as a beacon and will be really interested to see how that evolves and how we may be able to integrate that [into Connected Home]." Then there is the possibility of building out augmented and mixed reality capability that could reduce the workload of its 12,500 engineers. "This is not something we are actively building out," Jokhoo said, "but we do have an R&D department and helping customers over digital channels to show them how to reset a boiler or pressurise their radiator system or even, within the realms of safety, to help them through the process of connecting up their devices. The help and support is really where it will come into play." **Scott Carey**

Credit: iStock

Five pitfalls to avoid when migrating to the cloud

Not analysing your apps, business and costs can be major mistakes

Mistakes can be costly. They also can be so painful they keep you from venturing any further ahead. Of course, that's true with almost anything tech-related, but IT managers will tell you that there are some common, and potentially damaging, pitfalls that anyone looking at a cloud migration should work to avoid.

Migration mistakes can cost the enterprise money and time, and eliminate or reduce any expected increases in agility as well as speed and cost savings. Those stumbles and losses could cause business

executives to back off from a bigger cloud migration. It also could cause executives to lose faith in their IT leaders. “This is part of the learning curve,” said Deepak Mohan, an analyst with IDC.

“The negatives are attributed to the cloud and not to these mistakes that need to be corrected... If a company does not realize the cost savings and they fail to see the results they thought they’d get, the result is that there is a drop in faith and a lowering of confidence in your cloud strategy. And that will cause a slowdown in adoption.”

Part of the issue is that the cloud is really a different beast for a lot of IT shops. Everyday IT jobs, such as data storage, app development and resource allocation, need to be thought about and handled differently.

“What we learned is that while it’s easy to get started, cloud is completely different from IT,” revealed Temujin Baker, senior manager of engineering and architecture for King County, Washington. “How you run your business in the cloud is different than how you run it” on premises. “There are changes in how you do your work, the skills that are needed, the process.”

King County, with 220 connected sites and more than 13,000 employees, began migrating to the cloud in 2015, using AWS and Microsoft Azure. It started with 1,600 on-premises applications and now has 30 of the 1,600 running in the cloud, with plans to move 120 more to the cloud this year. The county will consolidate and modernize some of its on-premises apps, so not all will move to the cloud.

Aaron Barak, King County’s director of operations, said 2015 was all about seeing where the pitfalls lie and figuring out the best ways to avoid them. “The first year

or two was really about learning and getting our staff ready to support the move," he said.

And while Daniel Morreale, CIO of Hunterdon Healthcare System in Flemington, New Jersey, said there definitely are pitfalls in the path to the cloud, there are ways to avoid them.

"I expect that each instance is going to be a little bit different, so there will always be an opportunity to make a mistake and then learn what you did wrong," he said. "Making my data accessible from anywhere in the world and the capacity to eliminate a lot of my hardware and not cooling my data centre and dealing with power, it makes moving to the cloud a good value proposition."

Here is some guidance from IT people who have started their own cloud migrations, as well as from industry analysts, on some big mistakes to avoid in your own move to the cloud.

1. Not analysing your apps before the migration

One of the mistakes that Dave Bartoletti, an analyst with Forrester Research, warned against is failing to analyse a company's applications before considering any kind of cloud move.

"The first mistake is trying to do too much," argued Bartoletti, who added that IT managers need to start the migration with a small set of applications. "You need to run a portfolio analysis – either with a consultant or on your own – to determine what apps are good for the initial migration."

For instance, organizations need to split their applications into two initial groups – apps that will be migrated and apps that will be replaced.

It generally wouldn't make a lot of sense to migrate an on-premises email application when an enterprise simply could start using a cloud-based app, like one in Microsoft's Office 365.

Bartoletti also recommends IT managers check how compliant-sensitive an application is. "The cloud is a safe place to run compliance-sensitive workloads, but don't make your most sensitive data or applications the place to start your migration," he warned.

Companies also should be aware of applications that have an elastic load pattern. If an app tends to have workloads that drop and spike, then it's a good candidate for the cloud where it easily can be scaled up and down, taking advantage of cloud pricing so a company isn't paying for hardware when it's not using it.

Hunterdon Healthcare's Morreale said they made a mistake during their initial cloud migration because they didn't do enough application analysis. The single hospital – with 186 beds, 60 ambulatory locations and about 3,000 users – is in the process of moving to the cloud and is completely doing away with its data centre.

The company started the move in June by converting its email system to Google's cloud-based G Suite service. A mistake Morreale said they made when moving to G Suite from Novell's GroupWise was not analysing what their users needed from their old email system and not figuring out how differently G Suite was going to work from what end users were accustomed to.

"G Suite and GroupWise do calendaring pretty differently," he explained. "We did not provide our executive assistants who manage multiple calendars at the same time with the tools they'd need to be efficient. We recognized that by day two and saw that we had

messed up... The needs were not adequately explained or understood. We ...admit that we blew it." Morreale turned to IT consultants SADA Systems, which sent someone in for a week to sit down with the executive assistants in small groups and one-on-one sessions to help them learn how to handle their calendars.

When King County migrated its enterprise backup, they too erred in fully analysing their app. Baker explained they did a lot of testing when it came to moving the data to the cloud. What they didn't do was test how it would work getting the data back for an actual restoration.

"The purpose of the backup system isn't to back up, but to restore," he said. "The performance of the restoration wasn't there because we hadn't done the full testing of the application. With an on-site restore, it starts as soon as you press a button. With the cloud, there's a three-hour-plus window where the service has to go back and get that data before we get a restore. That's something that we hadn't fully anticipated."

The lesson was to test the whole process the way it will be used. Barak added that factoring in the time it takes to retrieve the data and send it back, with a large outage, the delay could have extended to 24 hours.

2. Forgetting a business analysis before starting

Forrester's Bartoletti said another major mistake enterprises make is to not do a business analysis before they launch a cloud migration.

"Starting without a clear business case means you don't know what the benefit is going to be," he said. "You need some analysis to see if this will actually save you

money... Just picking up an application and moving it to the cloud sounds easy – but why are you doing that?"

During the business analysis there are key questions to ask. For instance, do you need your apps to run cheaper or faster in the cloud? And how can you optimize your apps so they will run cheaper and faster? Also, which migration tools will best fit your move?

"Selecting those tools takes time," said Bartoletti. "You don't just Google search tools for migration and use the first one that pops up."

He also suggests that IT managers profile an app's behaviour and figure out what resources it needs, compared to what it's currently using. "You don't want to go to the cloud and request a terabyte of memory when you don't need it because you'll pay too much," said Bartoletti. "You need to be 100 percent clear on what that app really needs, not what it has been using."

John Trujillo, an assistant vice president, technology in the retirement solutions division at Pacific Life Insurance, said they began migrating to the AWS cloud in 2015. "There were definitely mistakes that I could see that, to some degree, we made," he said. "One is that cloud is not a strategy... It's a tactic toward achieving business strategy and business objectives, like agility and innovation. What are the business drivers to getting to the cloud? It's not just about getting there."

For Trujillo and his team a mistake was not putting the benefits of moving to the cloud in clear enough business terms to senior executives.

"The speed of business is increasing," he said. "The work is growing. The cloud addresses those realities, but [business executives] didn't understand that... We were assuming they understood and were not really

helping them understand what the cloud is and how it leverages into business value."

3. Underestimating costs

Of course, one of the biggest draws to migrating to the cloud is the cost savings. With apps running in the cloud, companies don't have to pay IT workers to do things such as manage email, deal with costly data centres, buy as much hardware and keep applications operational.

The issue is, though, that there are costs to cloud migration, too. For instance, Bartoletti said IT managers often overlook the cost of moving storage.

"The most expensive part of migration is often moving storage or the data," he added. "It might require a 100GB data base. The cloud providers don't generally charge you for moving data in, but it could take several weeks or months if it's huge. You want to estimate in your migration project how long it will take to move your data."

Pacific Life's Trujillo agreed. "There's definitely an upfront cost in getting to the cloud. There's no doubt about that," he said. "Data storage in the cloud, properly configured and architected, can be significantly cheaper" than what we do on-premises today. "Moving it into the cloud is more of a people cost."

IDC's Mohan noted that it also can be costly to forget to reallocate people and resources.

"You need to spin away the ongoing costs and the manpower costs related to an operation" that you used to run on premises, he said. "If you're reducing your data centre needs, can you cancel the lease or sublet it and figure how you'll use the workers who used to run that for you."

Trujillo also warned that IT managers might not consider how much they're going to be distracting their IT team during a migration. IT workers, who might otherwise be focused on business-specific jobs, will be working on the migration. That will slow down or postpone other projects. Of course, the promise of the migration is that once it's in place, those same workers will have more time to be free from many of the old mundane tasks, so they can focus on business differentiation and advantage.

4. Not getting the training you need

Many IT staffs probably don't know a lot about the cloud, especially if this is their first migration. And finding cloud-savvy people can be challenging and expensive.

Without that cloud knowledge, a migration can take longer than expected and run into more problems. And, possibly even worse, once the migration is over, IT might find they've simply moved their old on-premises mess to the cloud. IT workers, and managers, need cloud training, and consultants may be needed, so IT not only understands how using the cloud works but they also understand about modernizing applications, so they take advantage of cloud services.

"One of the mistakes we made early on is not fully appreciating how scary this can be to people who have been in IT for a long time," Pacific Life's Trujillo said. "One of the critical things we didn't do well enough up front was help the rank-and-file IT staff understand that this might be scary at first, but it's really an opportunity for them to stay current and fresh in the industry."

In the absence of proper training, there also is more risk of legacy thinking being brought into the cloud."

5. Not moving beyond 'lift and shift'

Moving beyond legacy thinking will help IT managers avoid a major pitfall of cloud migrations.

"IT needs to consciously change their thinking and their processes," IDC's Mohan said. "The process change is the piece that enterprises often fail to consider."

The issue is that the cloud is much more than just picking up an application or data and dropping it in the cloud. The real benefit to the enterprise is taking advantage of cloud services, like mobile back-end services and performance management services. Using the cloud shouldn't be just about cheap storage.

"The cloud is requiring us to rethink long-cherished theories in enterprise IT," Pacific Life's Trujillo explained, adding that they caught themselves before making this mistake. "Taking your legacy thinking into the cloud is a big potential pitfall. There are very significantly new and better ways to do things in the cloud than we've all grown up with."

For instance, he said they were used to application developers or even business people coming to them and saying they needed a server. IT would tell them how long a wait it would be. "In the cloud, they don't need to come to me," said Trujillo. "They can spin up a server in a couple of minutes. A danger is continuing to be guardians of the gates in the cloud, rather than letting app dev and even business being able to spin things up on their own... We need to switch our minds from being service providers to service brokers."

In the end, analysts and IT managers say it's unlikely that companies will avoid every pitfall when it comes to cloud migration. Some problems you can anticipate, some you can't.

Build padding into your migration timeline so you can deal with issues when they come up.

"Our people don't care about the technology," said King County's Barak. "They just care about the work that needs to be delivered." **Sharon Gaudin**