


THE DOPPLER®

CLOUD POWER

GETTING THE MOST OUT OF HYBRID CLOUD



Expert advice to help you succeed in the cloud.



ROBERT CHRISTIANSEN
CHIEF CLOUD STRATEGIST

One of the top three reasons we see cloud transformations stall or outright fail is poor change management. On page 26, explore how to overcome the organizational challenges that are systemic to every large cloud initiative.

THE DOPPLER EFFECT

Digital transformation, fueled by the last decade's explosion of apps and data, is driving a fast-emerging world where change is constant, and speed is everything.

We live in a world that's hyper-connected, where everyone and everything shares data. Plus, we have expect instant gratification. The digital, hybrid cloud transformation is today's reality, but it's complex. Enterprises around the world, like yours, are in various stages of end-to-end business digitization, in pursuit of accelerated time-to-value in order to achieve critical business outcomes. We believe that those who act – at speed – on a continuous stream of insights from edge to cloud, securely, will win.

HPE is focused on accelerating every cloud strategy, so you can go further, faster. Strengthened by the acquisitions of Cloud Technology Partners (CTP) and Red Pixie, we can share insights, learnings and best practices from hundreds of successful cloud engagements around the world.

In this special edition of The Doppler, our experts share their expertise across three key areas of the hybrid cloud journey:

Move to the Cloud: The journey to cloud begins with a strategy and plan — to define and achieve the right mix of public and private clouds for your apps and data; to address the people changes and skill advancement; and to implement financial consumption models that will achieve cost savings. On page 4, we dis-

cuss how to accelerate strategic planning and migration by outlining our best practices and methodologies proven to help you succeed from the outset.


Innovate in the Cloud: The cloud isn't a destination, it's a new way of doing things. A hybrid cloud strategy enables enterprises to launch new products and services faster and immediately demonstrate operational excellence through agile development and cloud-native technologies for continuous innovation. The cloud has changed the requirements of IT teams. Developers expect fast and easy access to resources for rapid innovation. It takes more than just technology to see success with cloud – it requires a culture of continuous change and innovation, across the entire organization. Starting on page 52, you'll see how this is possible.

Run the Cloud: Optimizing your cloud operations with automation, security, governance and cost controls can be a daunting task. This is especially complex when you have multiple clouds across various locations. On page 72, we detail how to ensure your cloud environment continues to operate smoothly and safely, so you can stay focused on your organization's real differentiation.

The future belongs to the fast. HPE is here to help you make your hybrid cloud vision a reality with a run and sprint model to help you go further, faster. [Learn how.](#)



Robert Christiansen
Chief Cloud Strategist



Hewlett Packard Enterprise

THE DOPPLER is published quarterly by Cloud Technology Partners, a Hewlett Packard Enterprise company.

Join thousands of IT professionals who get our weekly email reports by visiting

hpe.com/info/doppler

What's Inside

16



60

**Six Key Activities to
Sustain your Cloud
Transformation**

**Go Big on DevOps:
Transforming the
Enterprise beyond
Automation**

PREPARING FOR THE CLOUD

- 6** Transformation is Hard, Here's Why It's Worth It
- 10** Hybrid Cloud vs Multi-Cloud: What's the Difference, and Why Does it Matter?
- 16** Six Key Activities to Sustain Your Cloud Transformation
- 26** Navigating the Cloud's Most Challenging Blocker: Your People
- 34** Cloud Economics: Are You Getting the Bigger Picture?
- 44** Six Key Enablers for Large-Scale Application Migration to the Public Cloud

INNOVATING IN THE CLOUD


- 54** Why Cloud is Making Computing at the Edge Relevant Again
- 56** IoT for the Enterprise: From Prototypes to Production
- 60** Go Big on DevOps: Transforming the Enterprise Beyond Automation
- 68** Data Persistence in a Containerized Environment

RUNNING HYBRID CLOUD EFFICIENTLY

- 74** Five Essentials for Managing Hybrid Cloud
- 78** Go Big on Security: Five Steps to Overcome Security Concerns in the Cloud
- 84** Why You Need an Enterprise Cloud Managed Service Provider
- 88** Getting Compliance Under Control
- 92** HPE's Evolution to a Hybrid Cloud Environment



PREPARE TO MOVE TO THE CLOUD



The journey to cloud begins with a strategy and plan – to define and achieve the right mix of clouds for your apps and data.

- 6** Transformation is Hard, Here's Why It's Worth It
- 10** Hybrid Cloud vs Multi-Cloud: What's the Difference, and Why Does it Matter?
- 16** Six Key Activities to Sustain Your Cloud Transformation
- 26** Navigating the Cloud's Most Challenging Blocker: Your People
- 34** Cloud Economics: Are You Getting the Bigger Picture?
- 44** Six Key Enablers for Large-Scale Application Migration to the Public Cloud



TRANSFORMATION IS HARD, HERE'S WHY IT'S WORTH IT.

Complex, and often long, cloud transformation results in an organization that can execute at a higher rate of speed, learn from failures and constantly improve.

Many enterprises are struggling with the same transformation challenge: how to create an organization that can quickly adopt new technology while responding to immediate market needs and new competitive threats. The transition to becoming that kind of operation is not simple. It involves all aspects of an organization developing new muscle memory, based on new ways to operate that are accepting of learning and failure. It also requires empowering staff to make rapid decisions driven by common organizational direction and principles.

Cloud is the platform component that enables a more agile, responsive organization. But consuming cloud in a way that a business can also transform requires a coordinated effort across multiple domains. All aspects of the people, process and technology components of an organization are impacted as it matures from current to new operational models. Figure 1 outlines the key changes that must occur in order to maximize velocity. The primary change is the move from approval based processes to a focus on exception handling and constant organizational improvement.



Figure 1: New and Old Worlds

Many organizations start their cloud adoption journey with the intention of saving costs on IT operations. While the cloud can bring cost savings, its larger benefit to an organization lies in enabling faster time to market for new products and capabilities. The real value of cloud transformation is the organization's new ability to quickly consume the latest technology and rapidly adapt and respond to market needs.

Determining an organization's cloud transformation maturity requires careful measurement of both current maturity levels and the level of effort necessary to transition to the next higher state. Figure 2 shows the 10 dimensions commonly used to measure current cloud maturity and track organizational plans to achieve the next level.

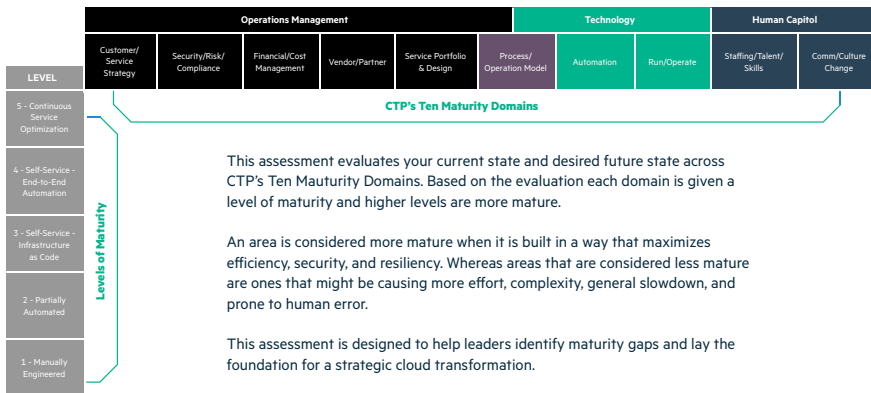


Figure 2: Cloud Transformation Office

Cloud transformation was originally heralded as a path forward for organizations to become more competitive, more agile and more able to recruit the best talent. But the reality is, transformation has become a sticking point for many organizations. There are several intersecting reasons for this:

- **Years to Build Can Take Quarters to Untangle** – Many of today's business systems took years, even decades, to build. Consequently, they possess a level of unknown complexity that takes time to understand, untangle and improve. Many organizations are not prepared for the sustained programming investment necessary to untangle these complex systems.
- **Business Cycles** – The usual business cycles associated with market changes, mergers & acquisitions and divestitures takes a toll on organizational change. These dynamics introduce time-bound projects that consume resources and attention. But cloud transformation is often not something that can be time bound; rather, it moves at a pace driven by the organization's level of acceptance and readiness.
- **Changing Regulatory Landscapes** – Many industries, both in the U.S. and abroad, continue to see regulations change. This changing landscape makes risk averse leadership reluctant to invest in programs that might not meet new regulatory standards.

With a successful cloud transformation, come several key elements that can accelerate growth for the organization:

- **Agility** – An organization's ability to respond to new competitive threats and market conditions is a key measure of its cloud transformation maturity. The ability to quickly consume new technologies and apply them to market conditions is an important characteristic of successful players across markets.
- **Innovation** – A company can live and die on the ideas its staff can come up with and quickly incubate into products. A cloud mature organization with a learning culture can encourage staff to try new things, test new products and fail fast.
- **Cost Optimization** – Cost savings are no longer the focus, as the ability to optimize cost to market conditions emerges. Cloud mature organizations can measure their investment against key value chains and determine if that investment is optimized based on the company's P&L.
- **Fail Fast** – An organization's ability to rapidly try new ideas, learn from them and iterate is a key element of a successful cloud transformation. Most important is the ability for an organization to respond positively to failed projects, take the learnings and share them with the wider organization. Encouraging risk taking is critical to rapid innovation.
- **Recruiting Top Talent** – Organizations define success by the talent they can hire and nurture internally. Recruiting in today's fast-paced world is significantly different from traditional recruiting for IT positions. Organizations must focus on finding people who are adaptive, self-learning, highly engaged and dynamic. These traits ensure individuals evolve with the company, demanding improvement in everything they do and lifting the organization during their tenure.

Upon the successful transition to a modern approach to technology, organizations will be able to quickly move from idea to pilot projects to production, as the benefits of their new capabilities are realized. The value of this complex and often long cloud transformation is in creating an organization that can execute at a higher rate of speed, learn from failures and constantly improve, while developing a pool of talent that is always looking for new ways to improve and innovate.

Written by Joey Jablonski, former CTO of Cloud Technology Partners (CTP) and currently VP of Data Engineering & Analytics at iHeartMedia.

HYBRID CLOUD VS. MULTI-CLOUD: WHAT'S THE DIFFERENCE, AND WHY DOES IT MATTER?

A baseline to understanding how these two architectural cloud approaches are being utilized by major enterprises.



Although it's been a long road, enterprise IT is finally achieving a general awareness of the benefits of cloud computing. While a picture of cloud is emerging in their thinking and strategic planning, the path between the "here and now" and the rosy cloud future tends to be murky. These companies' future cloud environments are variously described using phrases such as multi-cloud, hybrid cloud, cloud bursting, distributed cloud and even fog computing. While any of these represent topics sufficient for a whole series of articles, let's just look at the two terms that are often confused, yet are likely to be the most important over the next few years: hybrid cloud and multi-cloud.

Hybrid Cloud

No large enterprise, no matter how well prepared, can simply leap to the cloud in one fell swoop, even if the goal is to migrate completely to a public cloud provider such as AWS, Google Cloud Platform or Microsoft Azure. There is going to be a necessary transition period, during which the enterprise will have some resources, systems and workload capabilities that have been migrated to public cloud, while others remain in the enterprise data centers or colo hosting centers. This interoperability is a common example of a hybrid cloud.

Unless an organization is literally "born in the cloud" (built on the public cloud for essential infrastructure and product/service delivery, plus supporting SaaS services such as web-based email, Salesforce and Zendesk), every enterprise's cloud journey must include preparation for simultaneously supporting a cloud infrastructure and a legacy infrastructure. This requires conscious decisions about the level of integration vs. isolation that will be achieved between the data center side and the cloud side.


For many organizations, it may be tempting to simply graft a separate cloud environment alongside their traditional data centers, so as to minimize disruption of the existing internal operations and the introductions of new tools into existing environments. However, this path leads to increasing complexity, as more and more functions have to be simultaneously performed in multiple environments. So while hybrid cloud architectures vary, it is a best practice to anticipate the need to develop and deploy integrated platforms and architectures wherever practical.

Here are some characteristics that are typical of successful hybrid cloud environments:

- A centralized identity infrastructure that applies across multiple environments
- Persistent, secure high-speed connectivity between the enterprise and the cloud environment
- Integrated networking that securely extends the corporate network, creating a segmented but single overall network infrastructure
- Unified monitoring and resource management

Don't Put all your Eggs in One Basket

While it might make sense to diversify in order to minimize perceived risk when deciding to go to the public cloud, there are several challenges one must keep in mind.



Multi-Cloud

This term seems relatively self-explanatory: deploy cloud infrastructure on more than one public cloud provider, with or without an existing private cloud. However, the motivation for WHY companies might consider multi-cloud approaches and architectures is where things get interesting.

Risk Reduction, (“Don’t put all your eggs in one basket!”)

When organizations decide to go to public cloud, a typical concern is the perception of risk associated with dependency on one external firm, such as Amazon, Google or Microsoft. In response, it is common to wonder whether it makes sense to minimize that perceived risk by using more than one cloud provider, thus maintaining a complete and separate environment in each one. This provides an additional option in case the relationship with one provider becomes untenable for some reason, and, in theory, makes it possible to maintain services in the event of an outage at one provider. There is an instinctive logic to this approach; however there are also some realities that argue against it.

The first challenge is the complexity of maintaining an additional complete set of architectures and operational relationships, one for each provider. Given that most companies will already be operating in a hybrid cloud, this makes a total of three environments that must be maintained and operated. That doesn’t make multiple cloud providers impossible, but it needs to be understood. Note that there are vendors who offer valuable third-party products and services that can help provide standardized abstraction layers, theoretically minimizing the complexity of managing multiple cloud providers. A good example that comes to mind is Pivotal Cloud Foundry, especially known for enabling applications to run on multiple clouds.

But an important note here is that as soon as you depend on an “abstraction” provider, you have now re-created the single provider failure point you were trying to avoid in the first place. In addition, there is nearly always lag time between cloud providers releasing a feature, and the abstraction providers being able to support that feature. This creates an agility penalty. Given that enterprise agility and time-to-market for new products and features are critically important motivations for organizations to move to the cloud in the first place, giving away some of that agility is counter-productive. Finally, because the goal is to support duplicate environments with consistent capabilities regardless of which cloud provider is operating underneath the abstraction layer, it requires that each cloud provider has the same underlying capabilities. This leads to the next challenge.

The second (and larger) challenge to the “distribute your eggs” approach arises from the drive to the lowest common denominator. By definition, if the goal is to operate duplicate environments, then all the capabilities that are relied upon must exist in both environments. Of course, it is obvious that while the three main cloud providers have services and features that significantly overlap, they are not even remotely close to being identical. The result? Any full-fledged implementation of the “don’t put all your

eggs in one basket” multi-cloud approach is by definition limited to using the lowest common denominator set of features shared by the two cloud providers. This again results in an agility penalty, because when new cloud provider features and services are being considered, it is necessary to wait until BOTH providers offer the feature or service before it can be used in this form of multi-cloud implementation.

Architectural Similarity (“Like for Like”)

It’s not uncommon to find different technology stacks in different divisions or departments, because of acquisitions or high levels of autonomy among groups. One division might be heavily built out on the Microsoft ecosystem with SQL Server, .NET and C#, while another has a history of Linux, Java and other open source technologies. We sometimes see a pattern where individual departments may choose to extend workloads into the public cloud based on ease of migration to a given public cloud. For example, Microsoft Azure offers ease of migration for Microsoft workloads, so it’s not uncommon for a department or group to choose Microsoft Azure as their public cloud for that reason, while another department may choose AWS.

It’s important to note that this pattern is not usually consistent with best practice. While it offers some cloud benefits (e.g., OpEx over CapEx, scalability and agility), it creates two or more separate public cloud footprints, adding operational complexity, and limiting the ability to achieve a cohesive view of costs. It essentially becomes two duplicate environments.

Feature Availability (“Best of Breed”)

While the multi-cloud approaches above are fairly common today, we feel that a different model should be considered as a more successful one moving into the future. This promising multi-cloud architecture can be thought of as “Best of Breed.” With this approach, the mindset is that the agility risk imposed by insisting on duplicating all features in both environments actually costs the enterprise more than the stability that is theoretically gained by deploying two totally interchangeable cloud provider feature sets. Here the guiding principle is that in order to reap the full benefits of cloud, being able to take advantage of the best service and feature advances is of utmost importance.

A good best of breed approach involves selecting a primary cloud vendor. This vendor is where the main center of gravity for cloud operations lives, with the primary identity and security designs centralized around the main provider. It is of course straightforward to utilize new services and features from the primary vendor, but the enterprise also explicitly leaves open the possibility of reaching across to another cloud vendor for a specific service, capability or feature that is either not available from the primary cloud vendor, or does not meet requirements as well.

A reasonable question might be, doesn't that add complexity? The answer is, yes, it does. However, under the right circumstances, the benefit is worth the complexity. In this model, there is an architecture assessment process that explicitly considers the option of using a second (or third) cloud provider, given that the value for the use case justifies the extra effort. These types of scenarios can include the following:

- Reaching out from the primary cloud provider to use an API-driven service on the second provider. Because authentication can be handled at the individual request level, an entire duplicate identity infrastructure on the second cloud provider isn't required.
- Utilizing a particular query-friendly data store in the second cloud provider, populated via messaging queues or object storage originating in the first cloud provider. (This can be effective if egress data volumes aren't too high.)
- Machine learning training can be performed on a second cloud vendor, especially if the source data is publicly available. Then the results can be brought over to the primary cloud to build and deploy real-time scoring applications.

While the above are only three examples, they illustrate how this model provides a balanced approach, avoids the lowest common denominator problem and provides access to the latest cloud innovations, all while keeping complexity in check.

Of course, every enterprise is different, and there may be compelling reasons and priorities in a specific case that indicate a different approach. There are certainly architectures and options available other than those just discussed, but this should provide a solid baseline to understanding how two major architectural approaches to cloud are being utilized by major enterprises. Maybe yours can utilize them as well!

Written by Neal Matthews, Principal Cloud Architect, HPE

6 KEY ACTIVITIES TO SUSTAIN YOUR CLOUD TRANSFORMATION

All cloud initiatives require a heavy dose of people and process investment.

"Move to the cloud or die" is a common headline in today's tech journals. Such an alarming statement is not far from the truth. Any long-term technology roadmap worth executive consideration must have a significant cloud component. Unfortunately, research finds most multi-year IT transformations dominated by the technology prospects. All too often, leaders neglect to anticipate the people and mindset changes required to sustain a shift to a cloud centric universe. If you are sitting down to map out your long-term cloud strategy, please make sure to anticipate the budgets and expertise required to address the cultural shift, skill gaps, learning capacity and leadership skill changes.


The reality is, migrating just one workload to the cloud will impact your current IT operating model. Such a change will require a well thought out change plan, and an action list that scales to the volume of the workloads you anticipate migrating. This article highlights six key activities you must do well to ensure a sustainable and scalable cloud transformation program.

Cloud Transformation is More Than Just Change Management

Many managers don't realize the difference between "change" and "transformation." Understanding that difference is more than semantics. Change management consultants at Prosci define change management as "the discipline that guides how we prepare, equip and support individuals to successfully adopt change in order to drive organizational success and outcomes." So, "change management" means implementing finite initiatives, which may or may not cut across an entire organization. The focus is on executing a well-defined shift in the way things work. It's not easy, but the end results are predictable.

Typical change management initiatives are a process improvement or the implementation of a new tool. For example, when a global publishing company needed to extend a new CI/CD automation framework into its federated business units, there were shifts in

The reality is, migrating just one workload to the cloud will impact your current IT operating model.



roles, processes, tooling and teamwork. The change affected dozens of roles in the operations group. By applying well known agile change management principles and techniques — such as building a coalition of leaders, getting early results, engaging stakeholders and managing execution — the CI/CD automation was adopted successfully, and significantly improved the company’s aggregated software development cycle times.

Transformation is a more encompassing problem domain. Unlike change management, it doesn’t focus on discrete, well-defined change outcomes, but rather on a portfolio of projects, which are interdependent or often closely coupled. More importantly, the overall goal of transformation is not just to execute a defined change, but to recast the organization’s capabilities and to discover a new or revised operating model based on a future vision. Transformation is much more unpredictable and iterative, and requires a lot of experimentation, aka “learning fast from trial and error,” as the Lean Startup business books explain. Transformation also entails much higher risk. And, even if smaller successive change initiatives lead to certain accomplishments within the transformation portfolio, the overall success of the transformation is not guaranteed.

Reinventing oneself as an ISV to a SaaS provider is a transformation example. When an incumbent global software company realized that small startups potentially threatened their leadership position, they realized their current on-prem software solution was losing its competitive advantage and evolving too slowly. To respond to the emerging threat, the leadership team launched a cloud transformation strategy, with the goal of building a new SaaS-based product offering. This included a number of major “must-do” activities: implementing public cloud platforms, refactoring the core software architectures, introducing over 23 new technologies, and rebuilding, from the ground up, a new customer support model. The transformation also required leadership to align all initiatives to a new DevOps and “Customer-First” mindset.

Understanding that you are dealing with more than change is the first step in scoping your cloud transformation. In essence, a true transformation requires a shift in people, mindset and culture, not just a new set of cloud technologies and some process changes.

Planning Your Sustainable Cloud Transformation Initiative

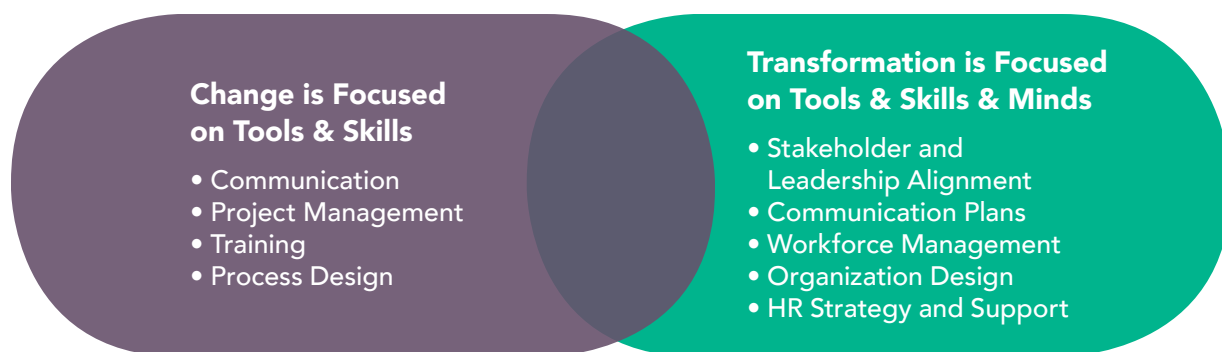


Figure 1: Transformation vs. Change

Key Activity 1: Keep your executive and sponsorship team actively involved in the cloud creation process

Your cloud initiative will fail without sustained executive and sponsorship alignment. The politics, resistance to change and focus on maintaining the status quo impact leadership and employees alike. The imperative presented with aligning leadership is that they control access to the resources and funding required to design, build and operate your cloud transformation. Funding is ephemeral: key sponsors can change their minds, especially in the rapidly changing technology space where there's a new fangled distraction invented every day.

Here, there are no easy solutions. One approach is to achieve consensus early, to continuously challenge the correctness of the cloud vision and to constantly renew consensus as to the transformation's direction. Too often, big programs try to over-mitigate risk upfront, based on the thinking that a fixed plan leads to fixed outcomes. One lesson learned from the agile world is you never know all the information upfront, so don't waste energy building a plan that pretends you do.

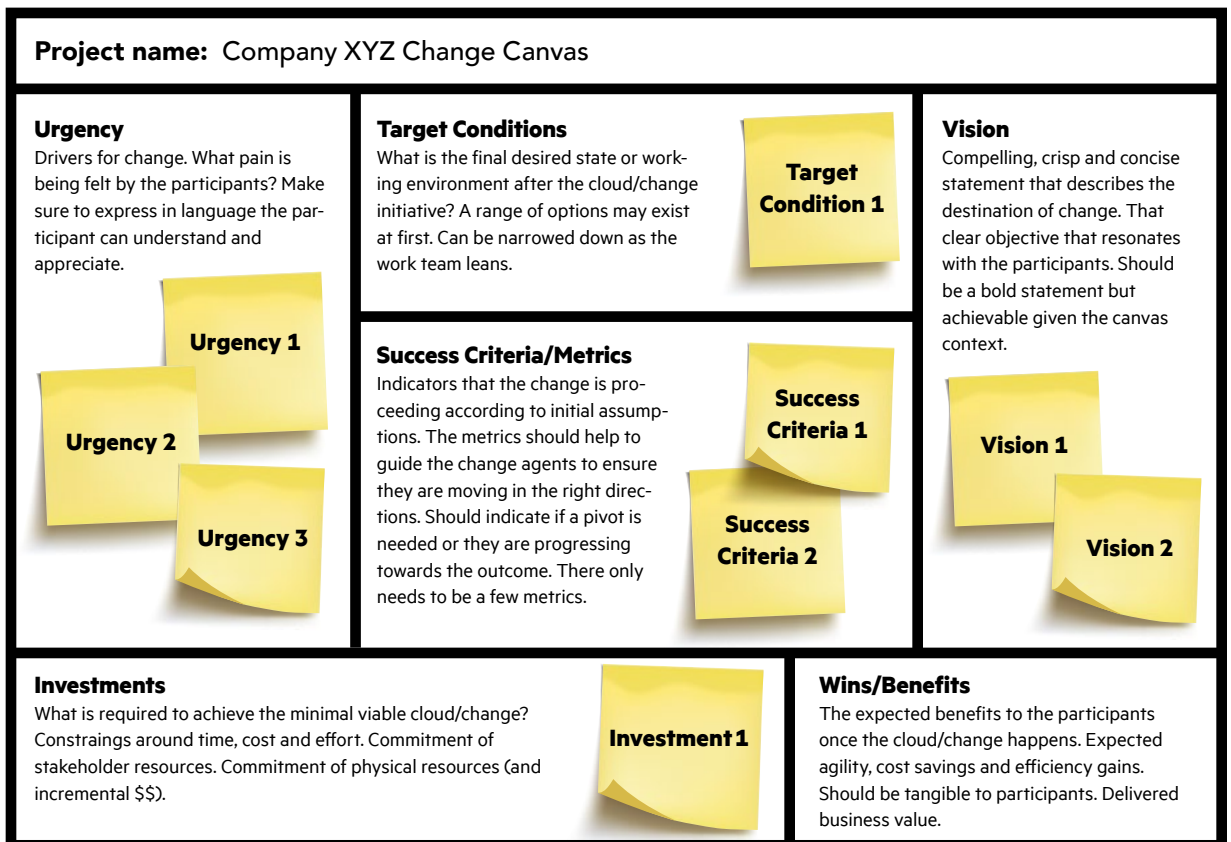


Figure 2: Using a white board and sticky notes to develop this Cloud Change Canvas, teams iterate and brainstorm to maintain and evolve the canvas for the lifespan of the transformation.

A powerful technique to achieve continuous executive alignment is to build a “Cloud Change Canvas.” A change canvas is based on the business model canvas concept, a tool created in the lean startup universe. A change canvas provides a one-stop view of the business’ urgencies, success metrics, target conditions, visions, investments and desired wins/benefits driving its cloud journey. Canvas authoring is a group activity — very visual, with a powerful style to engage participation. It’s like a business case, but it’s lightweight to create and easy to evolve as the team builds and learns. In summary, the exercise forces your leadership team to take ownership and sustain involvement in the cloud transformation process.

Key Activity 2: Assess your current cloud maturity and set direction

Most likely you are not starting from scratch. Even if you are, the fact remains that you can leverage the learnings of others who have cloud transitioned before you. One common approach is to conduct a cloud maturity assessment. If you are familiar with CMMI Institute's Service Capability Maturity Model, then you are familiar with the tactic.

	Service Operations				
	Customer / Service Strategy	Service Portfolio & Design	Vendors / Partners	Financial / Cost Mgmt	Security / Compliance
5 — Continuous Service Optimization	Enable Optimal Cloud Operations	Full Lifecycle Management of Cloud Service	Cost, Quality & Compliance SLAs / SLOs	Continuous Cost Optimization	Secure, Auditab
4 — Self Serve - End to End App/Infra Automation	End to End Visibility	Application Landscape Templates	Provider Roadmaps Integrated	Total Cost of Operations	Application Control
3 — Self-Serve - Infrastructure as Code	Application Aligned Services	Workload Profile Cloud Templates	Real Time Service Reporting	Total Cost of Service	Infrastructure Control
2 — Partially Automated	Application Oriented Cloud Services	Initial Standard Cloud Services	Vendor Management Framework	Real Time Cost Integration	Cloud Control Oper
1 — Manually Engineered	Heavy Lifting per Application	One off Service Design	Disparate Contracts	Manage to Budget	Compliance Docu

Figure 3: Cloud T

In the cloud context, a Cloud Transformation Office yields an effective cloud adoption strategy. It defines the characteristics that determine the stage of maturity, transformation activities within each stage that must be completed to move to the next stage, and outcomes that are achieved across multiple stages of organizational maturity. We developed the Cloud Transformation Office to assess companies across three focus areas: service operations, technology and human capital. The process is a simple gap analysis. The initial step is to assess the current state. Then, based on the organization's existing cloud capacity and readiness, a future state is identified and recommendations developed to achieve the "To-Be" state. It is a power tool to set your cloud journey in the right direction.

Key Activity 3: Conduct a stakeholder analysis

Assessing and profiling your key stakeholders is imperative. By definition, transformation means challenging shared beliefs, and potentially breaking organizational infrastructures designed not to change. A systematic and structured approach to managing your stakeholders' interests helps to ensure understanding and commitment to the

		Technology		Human Capital	
Compliance / Risk /	Process / Operation Model	Automation	Run / Operate	Staffing / Talent / Skills	Communication / Culture Change
Compliant le Services	Proactively Managed Services	Applications Auto-Scale / Auto-Heal	Automated Service Optimization	Industry Leadership in Cloud / DevOps Skills	Quality of Service Focused Delivery
ns deploys w/ Framework	Continuous Application Changes	Application & Infrastructure Provisioning	End to End Cloud Service Supportable	Optimized Cloud / DevOps Skills	Outcome Based Customer Oriented Delivery
re deploys w/ Framework	Self Service Infrastructure Changes	Infrastructure as Code	End to End Cloud Service Supportable	Federated Cloud / DevOps Skills Crosstraining	Application to Cloud Teams Partnering
Control Framework ational	Automated Service Provisioning	Unified Automation Framework	Cloud Service Patterns Operational	Cloud / DevOps Skills Maturing	Initial Team Demonstrate New Operational Model
ce Controls mented	ITIL Controlled Infrastructure	Automation Framework	Unique Operations per Application	Cloud / DevOps Skills Developing	Cloud Vision and Objectives are Outlined

	Level of Cloud Awareness	Commitment to Cloud	Influence over the Cloud	Impact by the Cloud
Stakeholder 1	Low	Medium	High	High
Stakeholder 2	High	Low	Medium	Low
Stakeholder 3	High	Low	Medium	Medium

Key	■ High	■ Medium	■ Low
-----	---	--	---

Figure 4: Cloud Stakeholder Assessment

cloud project. Also, grouping individuals and groups into discrete buckets facilitates building support networks to anticipate and manage people resistance. The analysis can be broken into two steps:

Step 1: A common stakeholder assessment approach is to identify and prioritize groups of individuals who:

- Are directly impacted by the new cloud initiative — including sponsors and change agents
- Are indirectly impacted by the new cloud initiative — including enablers
- Influence or power the new cloud initiative — including those who are with you and those who are not

Step 2: Once you categorize qualifying constituents into manageable segments, plot an analysis matrix into columns labeled “Level of Awareness of the Cloud,” “Commitment to the Cloud” and “Influenced and Impacted by the Cloud.”

The resulting single page data graphic provides a high-level road map of who you need to engage frequently, at mid-milestone and at major-milestone. For example, “I need your commitment and active roles in creating my cloud goals; this is the target you need to meet on a weekly, if not daily, basis.”

Key Activity 4: Design for effective communications and engagement

With a stakeholder map in hand, you are ready to establish an effective communications and engagement strategy to move leaders and teams to build sustained commitment to the new cloud vision. A cloud transformation plan needs to define and track engagement to address key content and timings. The tactical matters, such as engagement owners, mediums and measuring success metrics, all need definition.

Establish guiding principles to drive the change engagement. For example, the transformation adopters must own the cloud design, build and run phases. If they don't, then you own the risk of building the wrong people, process and technology solution. To further help drive commitments, make certain that:

- Success is measured in business outcomes, not just deliverables — business value must result from the cloud implementation, the sooner the better
- Information flows two ways and is open and honest — a dialogue to provide feedback is established and acted on
- Roles and responsibilities are clear — when people know what is expected of them, they can deliver results and respect accountability

Key Activity 5: Understand the re-skilling, education and training implications

One of the cloud's dirty little secrets, often not discussed, is that the newfound automation in these hyperscale platforms can render a significant portion of the operational staff unnecessary. While certain job functions may be automated out of existence, the people in these functions still have great value to the enterprise. This is where the importance of re-skilling, education and training come into play.

Such efforts entail a range of formal and informal activities aimed at building cloud knowledge, skill and experience. You want your training program to construct a bridge from the traditional operating world to the new cloud centric world. Your people training dollars demonstrates your company's commitment to its employees. It also enables the evolution of valuable competencies and services.

A unified cloud learning plan augments existing in-house expertise and development capabilities with vendor provided materials. In addition, it creates learning communities and networks comprised of early adopters.

A typical training program is a four phase life cycle:

- **Assess** — Determine existing skill gaps, learning objectives and certification tracks
- **Communicate** — Broadcast curriculum tracks, knowledge sharing opportunities and expert networks
- **Celebrate** — Reward and recognize those that participate
- **Measure** — Establish online portals and dashboards to report training participation and to track cloud specialists

Key Activity 6: Start building cloud and demonstrating value fast

This is the execution portion of the transformation, where you start to iterate and increment your demonstrable cloud solution. As opposed to a design-then-build “Big-Bang” technology project, the more successful path is to break the problem into small changes.

This is where you introduce the Minimal Viable Cloud (MVC) methodology. An MVC is derived from the lean startup community and their Minimal Viable Product (MVP) concept. The approach is to build a cloud implementation with just enough features to sat-

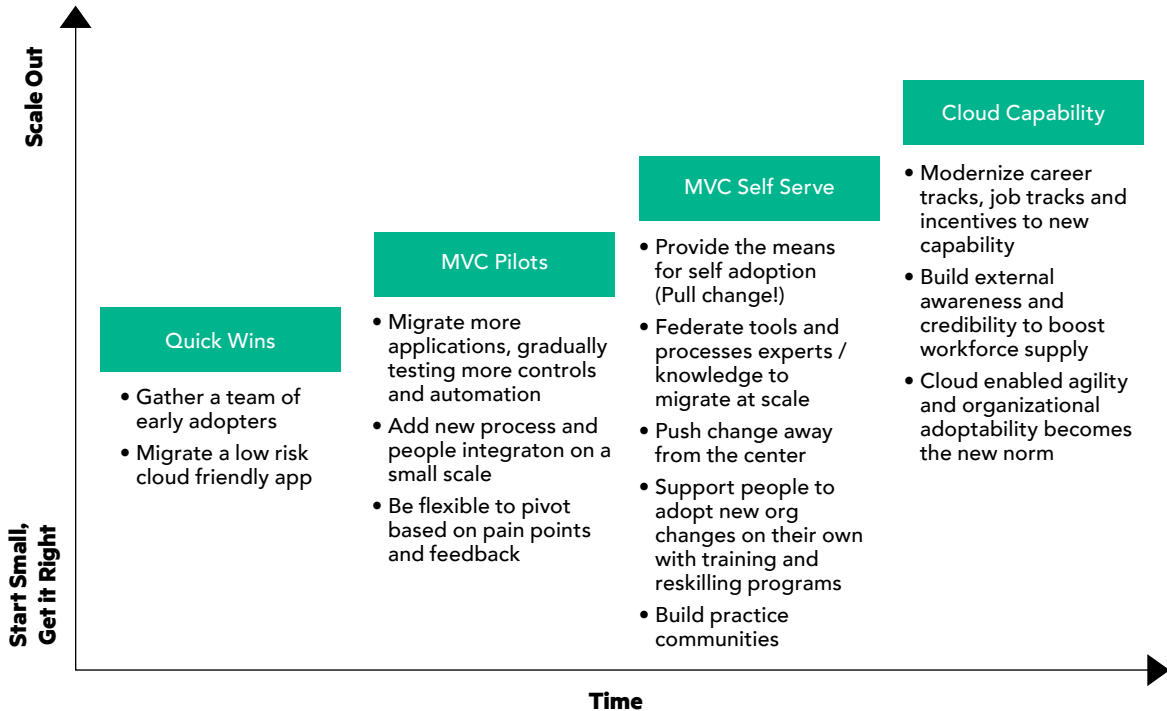


Figure 5: The power of MVC's — as you learn, you scale!

isfy early customers, and to provide feedback for future development. Time box yourself to a six- to ten week implementation window. This ensures you keep the requirements minimal, while turbo charging the learning and buy-in. Make sure to migrate a real-life workload, an application that tests the boundaries of acceptable business risk criteria.

Making it count motivates the organization to do it correctly. No pain, no gain, as the adage goes. Once you complete your initial MVC, continue to add workloads and capabilities to it to test additional controls and services. You will not always get it right the first time, so be ready to pivot if needed. Learn to tolerate failure. Trust in yourself to always respond quickly.

Summary

The six key activities I've outlined above are intended to help you deliver a scalable cloud transformation. At the end of the day, they propose to manage your three most common program risks:

- Managing the people resistance to change
- Building the right solution
- Institutionalizing the change beyond the program timeline

Weaving in modern day lean and agile practices provides the change management method to break up a very large problem domain into manageable pieces. This article provides a practical reference guide for you to now plan your long range cloud transformation.

Written by Paul Barnhill, VP, Principal Architect and DevOps Strategist, HPE

Learn more about the
HPE Transformation Program for Cloud

NAVIGATING THE CLOUD'S MOST CHALLENGING BLOCKER: YOUR PEOPLE

A simple model to help frame the organizational obstacles you'll encounter and the solutions you'll need to succeed in your cloud adoption program.



The Hard Reality

We see the same story play out over and over again within the enterprise -- people and process slowing down cloud adoption. Part of the problem is that we cloud technologists typically subscribe to a pure technology solution. We invest a heavy trust in the power of our designs. Who dare question the obvious benefits of on-demand and API driven cloud infrastructures?

Unfortunately, we all know too well the many dead bodies laid to rest in the “You build it and they will come” project graveyard. So why do most cloud adoption programs flounder to achieve their promise for agility and speed? We find the majority of cloud initiatives inadequately invest resources to address the most common obstacle: People’s natural response to resist, and even obstruct, technology and process change.

Why Do People Resist the Cloud?

The bottom line is the cloud is a transformative technology. It threatens the status quo of the technologist’s traditional roles and responsibilities in the large enterprise. The cloud forces companies to adopt new paradigms for how they consume and provision their infrastructure. The cloud compels them to build out an entirely new array of service capabilities, which in turn forces the adoption of new processes and people competencies. Infrastructure is no longer a physical asset and an operational and service management problem. Infrastructure in the cloud era is a software design and a development management problem.

A major cloud initiative introduces an entirely new set of technologies. This means workers need to retool and acquire new skills. For example, network engineers no longer manage firewall and router hardware configurations, but instead must learn how to manage virtual private clouds and build their network architectures engineered as code.

The cloud displaces the demand for certain roles and increases demand for others. The enterprise compute and network consumers can expect to provision their infrastructure by a self-service portal. No longer do they provision their hardware via service tickets and by centralized edict. Automation eliminates the need for long meetings with cadres of system engineers and project managers. In the cloud, any provisioning task can be reduced to the click of a button or a software API call.

An Introduction to the Organizational Change Management Performance Model

Is a cloud adoption initiative any different than any other organizational change management project? What we’ve learned over the years is the resistance to the cloud is really the same people opposition you see in any major enterprise transformation project. Consequently, we can turn to the same tools provided by organizational change management and process re-engineering fields. Figure 1 illustrates an introduction to what is commonly called the “Organizational Change Management Performance Curve.”

Change Management Performance Curve How Employees React to Change

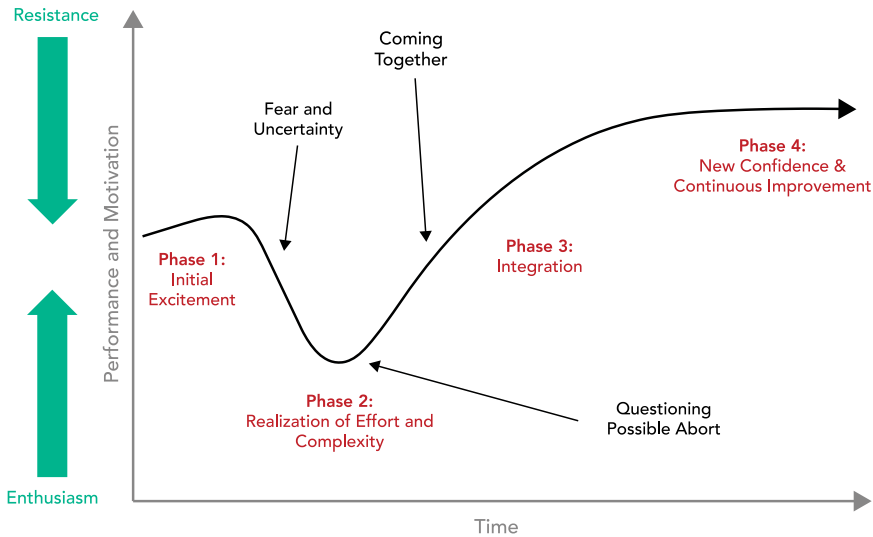


Figure 1: How Employees React to Change

Let's walk through the typical transition phases from the employee's perspective.

Phase 1: Initial Excitement

In the early phase, people form an initial excitement and have an open mind for what benefits a cloud adoption program will have for the company. In this context the status quo seems intact. Teams involved directly or indirectly generally know what is expected of them and how to react. The new initiative looks like a typical project. All the participants easily navigate the program's milestones using existing process guardrails.

Phase 2: Realization of Effort

Once they learn more details, they start to ask questions as to how the initiative will impact their roles and those of their team's. As they become more informed, they become more inquisitive. They start to realize the program's potential consequences on their day to day duties. People become aware of the magnitude of the depth and breadth of potential change. Their behavior shifts from cooperation to resistance. They seek refuge in multiple forms of denial such as rejection and diversion. They start to make comments such as "This too shall pass" and "We're special, the cloud will not work for us given our unique business requirements."

The Realization phase is the most critical time in the cloud adoption lifecycle. If questions remain unanswered, employees will start to build their own, almost always, ill-informed narratives. These false narratives, and any associated emotions of fear or panic, start to dominate the employees' motivation. They may start to handicap themselves or their teams with unrealistic goals and poorly thought-through execution plans. They may seek

“Recognizing and overcoming employee resistance is a critical aspect to a successful cloud adoption program. Just announcing you are going to the cloud is an insufficient launch point for transformational success.”



delay tactics such as extensive analysis, over engineering, or the addition of unneeded complexity. Their personal goals no longer align with the company's goals. They seek to preserve the status quo at any cost, regardless of their employer's business objectives.

Phase 3: Integration

Eventually, the chaos of realization passes. Integration is the phase where employees and stakeholders start to discover how the cloud adoption program benefits them specifically. The group learns that cloud skills are in high demand and increase their value in the marketplace. People start to have a vested interest in a positive project outcome. They set expectations and norms on others and try to align to the company's new way of thinking.

Group members in this stage may need more support than might be expected. They can become easily frustrated when things fail to work perfectly the first time. Although the team may feel good, they are also concerned that the initiative may fail, forcing them back to the uncertainties of the Realization phase. The employees need reassurance and new methods for forging through the unplanned difficulties of this phase.

Phase 4: New Confidence

The final of the four Performance Curve phases is a transition point where cloud becomes the new norm. If the change is well conceived and assimilated, the group is in agreement and performance is in full alignment to the new cloud established practices. They feel a sense of accomplishment and are open and honest about what's at stake should they be unsuccessful. Those directly involved in the project start to actively recruit new believers. Their implementation crawl is becoming more of a fast walk or slow jog.

Your People Engagement Plan — *Defining When and How to Engage Employees and Stakeholders*

A thoughtful plan to engage employees and stakeholders is a critical component to a successful cloud adoption program. Leadership needs to commission a governing body of cloud deciders, or what we call a Cloud Business Office (CBO), to draft a plan to ensure the right people are receiving the right message at the right time so they are prepared for what lies ahead.

A thoughtful change strategy needs to address three dimensions of the people engagement problem. First and foremost, management needs to identify who is impacted by the cloud. For example, the company must categorize their employees into distinct stakeholder groups such as program sponsors, change agents, influencers, resisters and those directly transitioned.

Secondly, management needs to outline the many activities needed to positively influence people adoption. These activities vary in complexity and implementation effort. Each needs to break down its messaging and purpose by key people and stakeholder groups. We find the planning process more powerful if management organizes and

analyzes their plan into the following six organizational change management disciplines.

1. Leadership and Executive Sponsorship
2. Stakeholder Management
3. Communications Planning
4. Reskilling, Education and Training
5. Performance and Incentive Programs
6. Organizational Alignment

Finally, management needs to identify when to engage which employee or stakeholder and with which engagement activity. Figure 2 helps illustrate how you can time people-focused activities into four distinct stages of engagement.

- Stage 1** Assess and Focus
- Stage 2** Design and Set the Direction
- Stage 3** Test, Learn and Demonstrate - Minimum Viable Cloud (MVC)
- Stage 4** Scale Migration and Full Operation



Figure 2: The Four Stages of People Engagement

Employee and Stakeholder Stress Point	What Challenges Leadership and a Cloud Business Office (CBO) Need to Address
<p>People are partially aware or not aware at all of the initiative and its impacts</p>	<p>“Why do we need to migrate to the cloud?” Management needs to share the economics and the business value the cloud brings to the company. There is no role in the organization unaffected by the cloud, from finance, to legal, to product development.</p>
<p>People do not understand how the cloud will impact them and their jobs</p>	<p>“Will I have a job after we migrate?” Management needs to be open and honest in all their communications. At first, they may need to get comfortable saying “we don’t know the answer to that problem now but we are working on it.” They need to demonstrate their undivided support to help their employees make the transition.</p>
<p>People are unaware of the cloud’s impact on their teams</p>	<p>“Does automation eliminate the need for my team?” Automation and infrastructure as code potentially eliminates tasks and activities currently done manually. Management needs to reinforce how automation is needed to increase agility. Automation does not always mean the need for fewer people, just a change in care and feeding skills. Remember, all this new automation needs sustenance and investment as well.</p>
<p>People worry they will not be able to master the new skills and knowledge to perform in the new cloud world</p>	<p>“I’ve always been a hardware expert, I don’t know how to manage software as code infrastructures.” Management needs to provide the skills, knowledge, and social training opportunities to allow employees to make the choice of transition.</p>
<p>People realize organizations and reporting structures change with how the cloud transforms software design, build and operations workstreams</p>	<p>“I feel the old world technology silos prevent the company from taking advantage of the cloud.” For the cloud to truly bring agility to the enterprise, the end to end software value stream must change to exploit newfound capabilities and automation. Management must make it clear that the cloud introduces new opportunities to shorten the software development cycle. This means the enterprise needs new reporting and team structures to promote cross functional collaboration.</p>

What to Expect From a Successful People Engagement Plan

A successful cloud adoption program depends on understanding the critical employee engagement activities and transition phases as outlined in a well-thought-out People Engagement Plan. Employees need help to understand how change impacts them on the journey to the cloud. Ultimately, an organizational change management program for cloud adoption needs to engage impacted employees and stakeholders to understand:

- Why the cloud program is needed and how it impacts the business's value delivery and effectiveness
- Cloud adoption readiness and steps required to mature cloud process and technology capabilities
- How, when and why the cloud-related changes are implemented
- Social, process and technical skills education and training requirements to sustain design, delivery and run software build cycles in the cloud
- Metrics and measures to align employee performance and behavior to the cloud program and the ultimate desired business outcomes

Recognizing and overcoming employee resistance is a critical aspect to a successful cloud adoption program. Just announcing you are going to the cloud is an insufficient launch point for transformational success. Walking employees and stakeholders step by step through the change process will give them an opportunity to address their individual fears and concerns. This in turn improves their performance and enthusiasm for a successful cloud program outcome. Once they see the light and achieve high levels of confidence, your cloud resisters quickly become your cloud evangelists.

Written by Paul Barnhill

**Automated out of Existence?
Keep Your Career Relevant**

[Watch the Video](#)



CLOUD ECONOMICS: ARE YOU GETTING THE BIGGER PICTURE?

Every company has its own ideas on how best to determine cloud ROI. For many, it's capital expenses (CapEx) versus operational expenses (OpEx). Cloud computing shifts IT spending to a pay-as-you-go model, like utility billing; you only pay for what you use, when you use it. For startups or new applications, this well-established argument is sound. Why purchase hardware, data center space, power and supporting software when they can be rented by the hour at a higher utilization rate? For enterprises, however, the decision is not so simple.

Enterprises have invested in data centers and equipment that is already bolted into racks and partially depreciated. For these organiza-

Identifying and making the right economic decisions requires additional investigation into the detailed costs and benefits of enterprise cloud.



tions, identifying and making the right economic decisions requires additional investigation into the detailed costs and benefits of enterprise cloud.

Thinking Beyond CapEx vs. OpEx

Most enterprises have hardware utilization rates significantly below 20% because of the excess capacity required to handle peak demand. As such, many companies carry up to five times the required hardware, networking, and data center space during steady state business cycles. If their computing demand is spiky, utilization rates outside of peak cycles are commonly below 10 percent. As a result, enterprises are spending much more on compute and storage than is required. Figure 1 depicts the traditional model where cloud shifts fixed CapEx expenses to variable OpEx expenses. To understand the full value of cloud for your enterprise, you must look beyond the CapEx vs. OpEx benefits and assess the other value drivers at play.

The Traditional View – CapEx vs. OpEx




Figure 1: shows the challenges of the traditional model. The blue line depicts hardware purchases which are completed periodically to increase capacity ahead of projected utilization. The red arrows indicate the cost an enterprise bears in excess capacity while the red shaded area represents the danger of underestimating utilization.

Closely associated with utilization is IT infrastructure's ability to scale up and scale down to support business agility. Traditionally, high utilization reduces IT spend, but limits agility and negatively impacts innovation and business growth. Conversely, the cloud can provide significant savings (near 100% utilization) and nearly infinite agility. The value of this agility is challenging to calculate so we have a tendency to ignore it. That is a big mistake.

Are you quantifying agility?

The business value of cloud is more about agility and utilization than any other cost consideration. Consider that the cloud provides us with the ability to provision and de-provision nearly unlimited resources as needed with complete control. Moving from 20% to near 100% utilization provides significant cost advantages and even greater value in the ability to quickly solve business problems without waiting for software and hardware procurement and installation. With the cloud, businesses can enter into new markets, accommodate new customers, avoid compliance penalties, or just move fast when they need to move fast, all while concurrently maintaining fully-utilized hardware and networking resources.

The value of the cloud scales with the value that its agility brings to your business. The faster you need to change, the more valuable it will be.

What's more difficult to quantify is the value of countless ideas and projects that are started, only to be stalled or placed lower on the priority list because by the time the resources are provisioned three to six months later, other initiatives have taken priority. With the agility of the cloud, projects can be conceived, provisioned and deployed within 24 hours - allowing for real-time planning and execution.

Uncovering the Real Costs & Benefits of Cloud

Most cloud ROI calculations don't factor in agility or fully capture the costs of poorly utilized hardware. While charts illustrating the differences between CapEx vs. OpEx are relevant to the potential value of cloud, a more complex assessment is necessary to reveal the full picture.

In the following pages, we go beyond simple CapEx vs. OpEx decisions and explore how leading enterprises are calculating their true value of cloud.

Key Themes to Remember

1. Cloud delivers both tactical and strategic value - so analyze hard and soft costs separately.
2. Quantify the value of agility as it may be your greatest overlooked driver for cloud adoption.
3. The value of cloud varies greatly depending on the enterprise and the industry. While some companies achieve near 50% cost savings, others only achieve 10%. This is typically a good indication that value is there and must be explored.
4. Focus on economics early in the cloud adoption journey and involve members of the financial team throughout the analysis. Organizations that are proactive in identifying their cloud ROI streamline their cloud initiatives and achieve greater success.
5. Identify what the key drivers are for your organization to adopt the cloud:
 - Is there a mandate to reduce costs?
 - Are you losing market share to competitors because of your inability to rapidly develop and deploy applications?
 - Have you experienced outages that are damaging your reputation in the industry?
 - Are you in need of more advanced, and cheaper, business continuity and disaster recovery capabilities?
 - Are you interested in expanding your global footprint and nervous about the latency in your current environment?

Defining ROI and TCO in the Cloud

Return on Investment (ROI): The financial gain from an investment in cloud divided by the cost of that investment.

Let's first look at cloud ROI. While businesses often discuss their "cloud ROI," more often than not they're missing the bigger picture. Most cloud ROI calculations are focused around IT cost savings and how they affect the bottom line. Instead, cloud ROI should be more about the value that is returned to the organization.

Value drivers that are often overlooked in typical ROI calculations include accelerated time to market, improved developer productivity, decreased provisioning time and many more intangible benefits of cloud.

Total Cost of Ownership (TCO): The sum of all direct and indirect costs of the IT estate including all application development, maintenance and support, operations, data center, network and BC/DR.

Cloud TCO defines what will be spent on the technology after adoption - or what it costs to 'run the engine.' Typically, a TCO analysis looks at the costs of the "as is" on-premise infrastructure and compares these costs with the costs of the "to be" infrastructure state in the cloud. TCO analyses are much simpler to calculate than ROI analyses; however, they only give the stakeholders a narrow view of the total financial impact of moving to the cloud.

So, the difference between a TCO and an ROI analysis is that a TCO defines the spending and savings, whereas the ROI determines what value is generated, while taking spending and savings into account. It's critical that you understand both, and their differences, in order to effectively define the full value of cloud for your business.

Identifying the Types of Savings

Start by defining the hard benefits of cloud in terms of direct and visible cost reductions and efficiency improvements. These costs are typically easier to identify and easier to assign a clear value to.

Soft savings are those value points that are more challenging to accurately quantify, but can be as valuable, if not more valuable, than cloud's hard savings. Any holistic cloud economics evaluation should have a thorough analysis of both hard and soft savings.

Hard Savings	Soft Savings
<ul style="list-style-type: none">• Reduced spending on compute, storage, networking, security• Avoidance of hardware and software purchases (CapEx)• Reductions in operational costs, backup and DR• Reduction in operations-oriented personnel	<ul style="list-style-type: none">• Reuse of services and applications that allow you to define and redefine solutions using the same cloud service• Increased developer productivity• Improved customer satisfaction• Ability to change business processes quickly around new and emerging opportunities• Increased global reach

How Much Does Human Error Cost You?

A key savings that is often overlooked is cost avoidance. You may notice that cost avoidance is not listed in the grid above. This is because cost avoidance can land in either hard or soft savings. Some mistakes are easier to quantify (such as having your e-commerce platform go down for an hour on Black Friday), but others, such as an impact to your reputation, are much harder.

The cloud's ability to help enterprises avoid costs is extremely powerful. For example, many of our clients are using various levels of cloud and DevOps to improve uptime and make operations, development and deployment much more efficient.

If you do not take the time to understand how much a lack of automation (and ensuing human error) is costing your company, you are missing a major opportunity to reduce future costs.

How Cloud Starts

Cloud initiatives often start as “shadow IT,” with business units requiring better agility than central IT can provide given their physical hardware inventory. As usage grows, Compliance and IT staff become concerned over governance and security. They use the cost savings (“TCO”) over internal or outsourced infrastructure and associated costs to build a business case to define a larger, broader cloud strategy. Once companies migrate, they reaffirm that although hard savings justified the hybrid approach to compute, the real benefits are found in the soft savings, coming full-circle to why businesses adopted cloud in the first place.

Agility: The Cloud’s Most Undervalued Benefit

Cloud agility is the ability to rapidly change an IT infrastructure in order to adapt to the evolving needs of the business. This is becoming increasingly important in today’s disruptive markets and the reality is that many enterprises are plagued with IT infrastructures that are so poorly planned and fragile that they are limiting business growth. Cloud agility provides a huge strategic advantage and significantly increases a business’ chance of long-term survival.

To help quantify the value of agility for your organization, start by breaking down its three components:

1) Your degree of change over time

The degree of change over time is the number of times that the business reinvents itself to adapt to market demands. While a pulp and paper production company may only have a degree of change of five percent over a five year period, a technology company may have an 80% change, and a company in a downturn market may have a 40% percent change over this same period.

2) Your ability to adapt to change

The ability to adapt to change is a number that states the company’s ability to react to required change. For instance, a large manufacturer may need to rapidly pivot in order to take advantage of a new market opportunity, but may not have an IT culture that can change at the rate required. Thus, they don’t have the ability to adapt to change, so, no matter what you do to promote the use of disruptive technology like cloud computing, they won’t be able to take advantage of it.

3) Your relative value of change

The relative value of change is the amount of money made as a direct result of changing the business. For instance, a retail organization’s ability to rapidly establish a frequent buyer program to react to changing market expectations, and the resulting increases in revenue from making that change.

Cloud Financial Maturity Spectrum

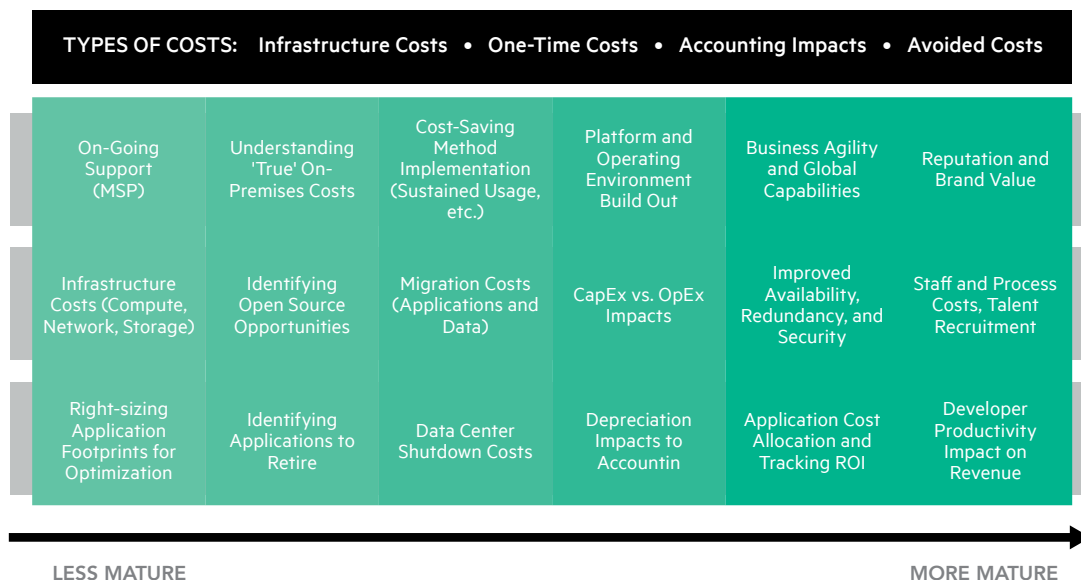


Figure 2: Achieving Cloud Financial Maturity

4) Your individual perspective on agility.

Finally, it is also important to take into consideration how different individuals at the same company view the benefits of agility. A CIO's perspective will vary widely compared to how a head of Infrastructure or head of engineering values their organization's ability to change. This is a key nuance that greatly affects investment decisions. The most successful companies focus on the overarching business agility the cloud provides, a strategic perspective that is generally in line with how CIOs view their organization's ability to change.

What We're Seeing

The value of the cloud scales with the value that its agility brings to your business. In other words, the faster an industry needs to change, the more value cloud brings. No matter the industry, achieving financial maturity during cloud adoption is challenging and requires exploring all the economic considerations we've already discussed.

We see that a company's cloud financial maturity on their cloud investment follows a consistent flow as the company matures. Companies early on in their cloud journey are focused on identifying hard costs and other clear cloud value drivers, which are easier to identify but provide a limited perspective to total ROI. Figure 2 depicts these considerations in light green at the beginning of the spectrum. As companies mature through their economic analysis and cloud adoption journey, their financial maturity increases as they start to understand the value of agility and other soft costs, which are depicted in dark green.

CASE STUDY

25% Savings in 18 Months

A multinational software and manufacturing firm faced rapidly growing IT operating costs tied to its legacy IT outsourcing provider. With a monthly IT spend of \$14.5M, the company identified cloud adoption as a key driver for cost reduction. To determine the potential savings, CTP conducted a TCO study comparing the current environment to target cloud endpoints. In the graph below, we show the company's 18-month TCO. In this case, the TCO compares using an IT outsource model vs. a hybrid cloud (public cloud services + an IT outsource provider). This TCO model calculates a payback period of 6 months for a hybrid environment, with savings growing to \$4M per month (25 percent reduction) at month 18.

Monthly TCO: IT Outsource vs Hybrid (Cloud + Outsource)

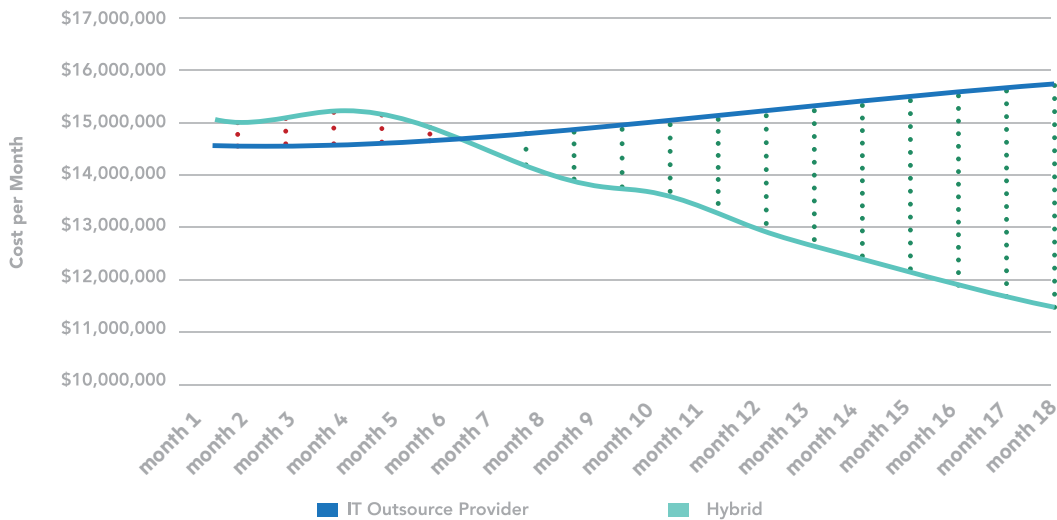


Figure 3: 18-Month TCO for an enterprise manufacturing company highlighting growth to a 25% savings.

The next graph shows the company's total run rate comparison between lift-and-shift vs refactoring vs colocation for a single enterprise application. Note that the run rate for colocation is lower, at first, in this model. However, it goes up significantly over time, relative to the cloud run rate which does not grow at the same rate, and in many cases shifts downward as the cost of compute and storage declines over time.

For this application, payback is at six months if the application is lift-and-shifted, or at 18 months if the application requires a major refactoring. Of course the cost of migration, whether it be lift-and-shift or refactoring, impacts the savings and when they are realized. Keep in mind that you need to consider the existing and future run rates of each approach prior to making your final decision.

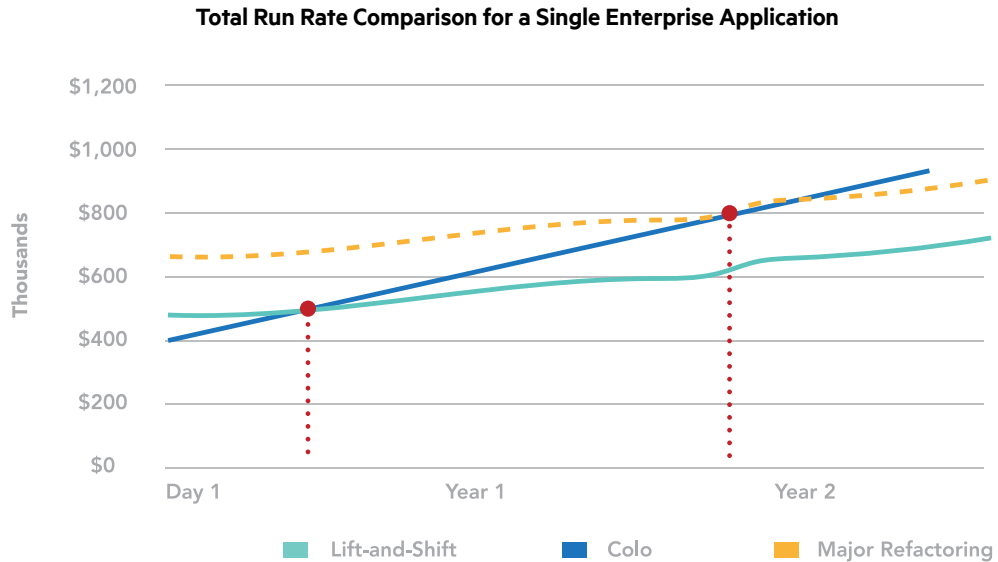


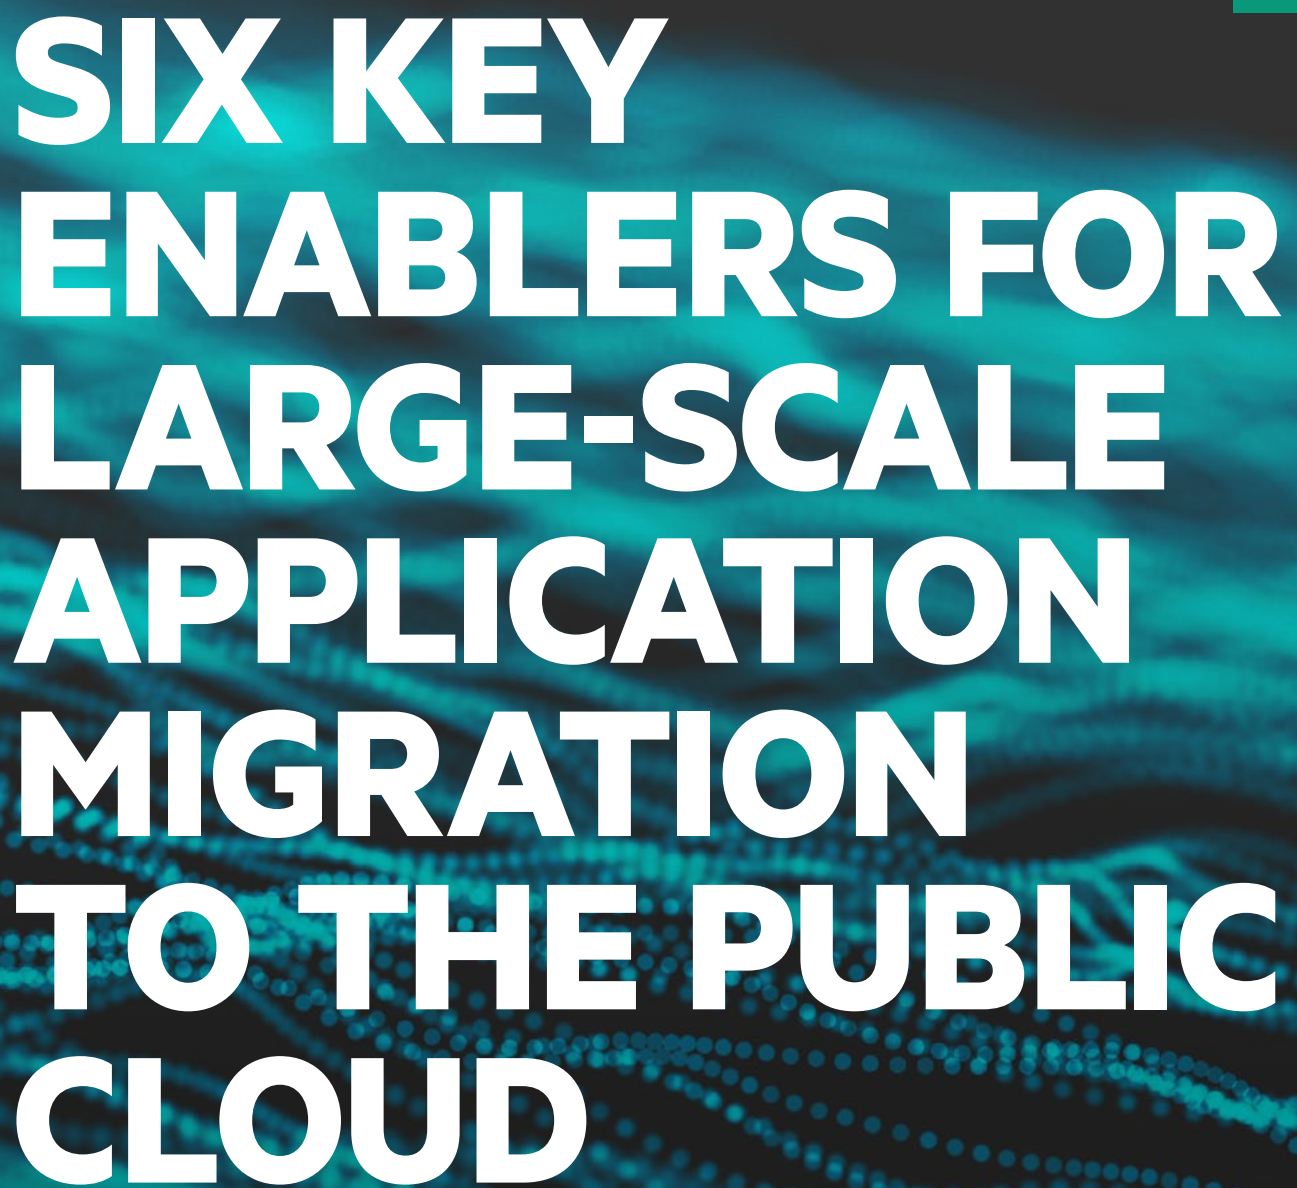
Figure 4: Application-level analysis for a single enterprise application.
 We see that this type of application-level analysis enables granular and data-driven justification for additional investments in cloud.

What's Next?

Cloud provides the perfect opportunity to change the way your organization runs IT. With cloud, IT has a much more positive influence on the business and is better aligned with strategic goals. IT will no longer be a drain on company resources, and many enterprises will find that the newfound efficiency and agility of cloud adds huge value to their bottom line. After all, IT is there to serve the business, not the other way around.

While few will dispute the hard cost savings of cloud, there is always a need to create a compelling business case that clearly defines the value of agility and other soft benefits. While it is important to remember that there are rare scenarios where cloud does not make economic sense, generally, you'll find a strong ROI and reduced TCO. Ultimately, you won't know the full costs and benefits until you define your business case and run the models.

Written by Greg Janecek, Director of Enterprise Sales North America, HPE, David Linthicum, former SVP, CTP and currently Chief Cloud Strategy Officer, Deloitte, and Stuart Robertson, former Director, Global Alliances, CTP and currently Director, Business Strategy, Immuta



SIX KEY ENABLERS FOR LARGE-SCALE APPLICATION MIGRATION TO THE PUBLIC CLOUD

A prescriptive approach and framework for a successful cloud migration journey.

In working with many clients across hundreds of projects, we continually hear similar themes around large scale migration efforts. Here are the most common areas our clients struggle with:

Scale and Velocity Requirements for Migrating Large Numbers of Applications: Most organizations take months to create future state application deployment architectures, and then migrate a single application. When an enterprise has hundreds of applications to migrate, this challenges the viability of migration.

Lack of Visibility into Complexities and Dependencies: If you don't have visibility into an application stack, you can't migrate it. Whether this is due to poor documentation, legacy or acquired assets, or decentralized deployment practices, it is common for organizations to lack visibility into application components, configurations and dependencies. Even though an enterprise may have a list of all assets and their details within a CMDB, how do you validate those assets, their architecture patterns and dependencies to determine which assets must move, and whether they are compatible with the target cloud platform?

Suitability of a Migration Approach: How do you determine the specific migration approach for an application (Rehost, Replatform, Refactor), and which applications can leverage cloud native components?

Migration Sequencing and Execution: How do you determine and account for the right sequencing and configuration changes for the various dependencies, and how do you account for the configuration changes required during the execution phase? Making a configuration change can take only a few minutes, but making the change out of sequence, or missing a change, can result in many hours or even days of troubleshooting.

Migrating Mission Critical Applications with Zero Downtime: Most enterprises have mission critical applications that must be up 24/7, with large data sets and databases constantly updated with multiple upstream and downstream integration points. Databases are especially difficult to move, and if moved incorrectly, these assets can have a profound impact on the organization.

Let's look at the six key enablers that must be in place for any migration project to succeed.

Six Key Enablers

Enabler #1 - High Value Data Collection

It is essential that businesses understand their application estate. We strongly believe in “High Value Data Collection,” which focuses on defining upfront: the context, the data to be gathered, how it will be consumed and the final outcome.

Taking on a migration initiative requires a deep dive into the enterprise’s portfolio. Gear your discovery and analysis effort toward addressing the main business drivers and specific goals. Based on those goals, target an analysis exercise at the estate, application and infrastructure levels, as well as specific to an individual component. We recommend that every organization perform an analysis at each of these levels based on various business drivers.

The following are the typical use cases:

- Defining general strategy for transformation across the landscape, understanding common patterns and identifying first movers (Estate Level Analysis)
- Addressing challenges for a specific application portfolio for a line of business (Application Level Analysis)
- Addressing specific pain points, such as middleware or database transformation (Component Level Analysis)
- Addressing more specific business drivers, such as “exit a data center,” which may require an infrastructure-centric analysis

The type of data required for each of these use cases varies from general application information to asset details, detailed architecture and dependency information.

Even though creating a data model to define exactly what is needed for each of these analyses should be a no-brainer, many organizations struggle with this. Issues range from not finding required information to being overwhelmed with the amount of data, as well as the time and effort needed to gather it.

We recommend that organizations spend up-front time and effort creating a data model that defines the use cases with the following requirements:

- Asset information
- Additional analysis attributes
- Data gathering mechanisms

The data gathering mechanism can range from self-service questionnaires to discovery/monitoring tools, to CMDB sources. Many discovery tools have additional capabilities for analysis, including cost analysis, architecture recommendations and platform recommendations.

In summary, organizations need to enable high value data collection through the proper definition of use cases, asset details and other functional data that is key to the analysis. They also need a robust discovery mechanism that can gather all the required information and maintain it in a repository to use in further stages of analysis and eventual migration, if required.

Enabler #2 - Objective Assessment and Analysis with Rules

Once we have the required visibility into the landscape and the various assets, an assessment and analysis determines the fate of an application. As described above, an analysis effort can be either at the estate level, the specific business and applications level or the infrastructure level.

The outcome from an estate level analysis typically provides visibility and direction for the organization, including first movers. It also shows where the organization needs to focus its efforts in the short, medium and long term to meet business objectives.

In the context of an application level analysis, a typical ask is: which applications are best suited for migration and what platform and architecture patterns are suitable? The key application metadata used in such an analysis falls into four categories: business, technical, operational, security and governance. Examples of technical metadata categories include: architecture, technology stack, automation, performance, scalability, dependencies, data size and data velocity.

A suitability analysis requires defining cloud-ready characteristics and a scoring mechanism that can be applied to all applications to determine whether they should be moved to the cloud. For example, an organization can define the following characteristics that deem an application not suitable to be moved to the cloud platform:

- Application is large, single instance and/or monolithic, and cannot be broken into services
- Application has an external dependency that cannot be reached from the cloud platform
- Application is not compatible with list of approved, compatible cloud libraries
- Application has contractual, legal or licensing issues due to a third party packaged application or technology needed for running in a cloud environment

Once you deem an application is suitable, you can choose the appropriate cloud platform and migration pattern. We recommend a quantitative approach in which each of the application's characteristics are scored against an endpoint and summed up to determine the suitable cloud platform. In many instances, the target cloud platform decision is based on other factors, including licensing, contractual availability of specific services or affinity.

Various factors determine the migration pattern, including business function, objectives, business cycle, criticality and priority, application architecture and effort required. For example, a custom, off-the-shelf application may be more suitable for Rehost (rather than Replatform or Refactor). You can apply a set of rules to the characteristics of an application to determine the migration pattern.

Determining cloud suitability and migration pattern requires creating and defining what suitability means and then carefully analyzing the characteristics of the application using objective and qualitative means.

Enabler #3 - A Prescriptive Migration Plan with Attention to Details

Once an application analysis determines the application is suitable for migration and its migration patterns are known, you need to develop a detailed migration plan. Besides the overall goal of the migration, a migration plan should include execution level details, including all migration tasks and owners.

The execution level details will vary based on the chosen migration pattern. For example, a Rehost pattern will mainly have infrastructure tasks with few application configuration changes, as well as a test and validation plan. A Refactor pattern migration effort will need to include details on each of the components to be changed, including its current and future state, functionality, accompanying code, deployment details, and test and validation plans.

At a minimum, the migration plan should include the following:

- Business drivers and expectations
- Functionality
- Current and future state architecture
- Patterns and approach
- Asset list, dependencies
- Desired configuration or changes
- Tools (migration, deployment, monitoring, logging)
- Detailed migration tasks, deployment plans with RACI
- Readiness, test, validation, cutover and operations plans

For migration efforts to achieve the desired velocity and scale, the migration plan for execution needs to be very prescriptive. In addition, reusable approaches and artifacts and self-service approaches for different stakeholders will greatly help reduce latency.

In the process of creating a detailed migration plan, one should account for approaches that reduce the amount of downtime required for cutover during a migration effort. You can use various approaches to achieve close to zero downtime for an application migration, including: blue/green deployments, incremental replication, isolated environment for testing and validation before cutover, or temporary DNS records for testing.

For example, you can develop standard migration checklists for each of the patterns, which you can reuse across all applications. As many organizations have a common set of tools for testing and validation, a common set of tests across patterns is useful to the operations team for onboarding applications at the desired velocity.





Enabler #4 - Automation with a Factory Approach

The key to address the scale, velocity and safety challenges in a migration project is to incorporate automation and reusability. As more and more tasks are automated, the migration process becomes easier and can be scaled, increasing the safety and velocity of migrations.

Migrations at scale require a factory approach. This uses a mix of tools and processes to improve quality, accuracy and precision. In terms of tooling, selecting the right mix of tools and investing time and effort to developing the required automation will achieve the required scale and velocity. Besides the tools, you need a process and an orchestration layer that addresses all facets of a migration, including sequencing, configuration, dependencies, testing and validation. You must also have an enablement program in which additional staff are trained and deployed based on the demands of a large scale migration.

In addition, an automation and factory approach will help reduce downtime during a migration, user errors introduced by manual tasks and traceability requirements.

In many instances, an existing CI/CD and governance system may suffice for applications where the deployment is highly automated. However, if automated deployment is not available (e.g., for COTS applications and the Rehost pattern), you need to develop it yourself or use out of the box automation provided by third-party tools.

Enabler #5 - Strong Governance for Visibility, Continuous Process and Quality

Executing on the migration plan requires addressing all moving parts during the migration project. Besides detailed migration plans, checklists and automation, you need effective play calling, essential for the smooth execution of migration tasks.

Migration requires a factory and network operations center like approach in which runbooks exist for the majority of the tasks and scenarios and all relevant staff know exactly what to do in which scenarios. You will also need a solid governance process to oversee the migration execution.

In a factory approach, you can dedicate a workbench (a self-contained group with all the tools and skills) for a specific pattern. The workbench concept is also useful in scaling, where additional workbenches are deployed to increase the velocity of the migration effort when required. But multiple workbenches do pose additional challenges in terms of governance and tracking, to make sure all parts of the factory are working to their optimal levels.

Enabler #6 - Key Cloud Performance Indicators (KPIs)

Once assets are migrated to the cloud, there needs to be a comprehensive governance and monitoring framework to continuously track various aspects of the target platform and applications. Even though many enterprises have mastered this aspect in their data centers, the cloud platform poses unique challenges where the same set of tools and processes may not work due to various factors, such as multi-tenancy, scale and lack of visibility/granularity to the infrastructure layer.

Besides deploying the typical infrastructure and application monitoring and logging solutions, you need to monitor additional elements when applications are in the cloud. A few examples we consistently hear from clients include cost management, posture and compliance. To keep tabs on the environment, the cloud platform requires a different set of tools and processes. We recommend defining cloud specific KPIs to effectively monitor the target platform.

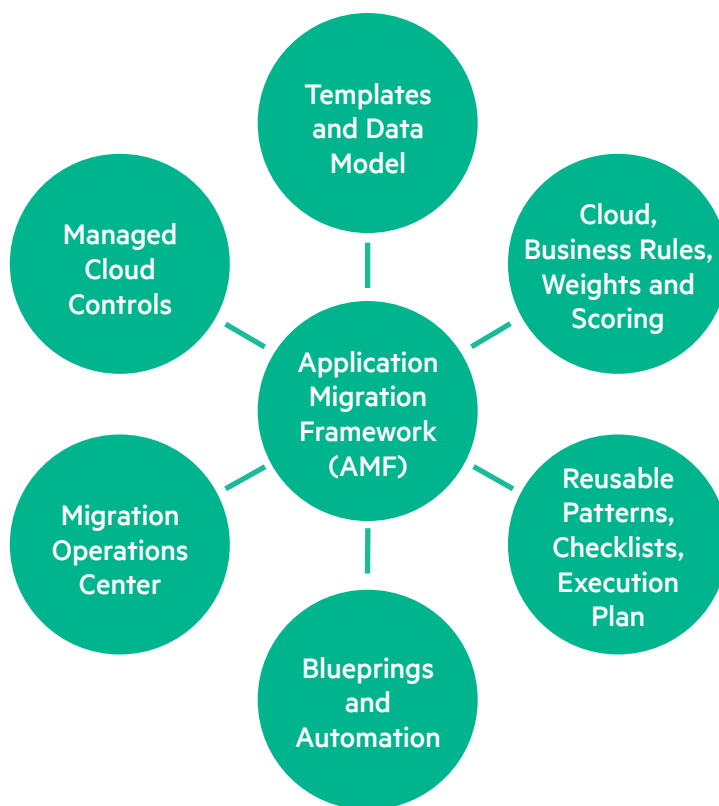


Figure 1: CTP's Application Migration Framework (AMF)



Conclusion

A large-scale migration project is a transformation effort that requires careful analysis, planning, development of the required automation, and flawless execution. CTP's Application Migration Framework (AMF) is an extension of the HPE Application Migration for Cloud service developed to address the multiple challenges and use cases encountered when migrating large portfolios of enterprise applications in a short period of time. It can take years to develop and refine the knowledge, skills, tooling and processes to quickly and efficiently migrate large application portfolios. AMF lets organizations reduce that time to migrate an application.

In summary, embracing each of these six enablers addresses all challenges quickly and more easily, encouraging enterprises to realize the benefits of adopting cloud. The six enablers promote fundamental change and deliver results by addressing a broad range of key stakeholder needs.

Written by Prakash Patil, VP, Principal Cloud Architect, HPE

INNOVATE IN THE CI



ING LOUD

- 54** Why Cloud is Making Computing at the Edge Relevant Again
- 56** IoT for the Enterprise: From Prototypes to Production
- 60** Go Big on DevOps: Transforming the Enterprise Beyond Automation
- 68** Data Persistence in a Containerized Environment

**The cloud isn't a destination,
it's a new way of doing things
and developing faster.**

WHY CLOUD IS MAKING COMPUTING AT THE EDGE RELEVANT AGAIN

It's been a common problem for years. If you gather large amounts of data from a device or other source, and you need to process that data instantly, then moving it to a centralized database each and every time introduces latency.

IoT must deal with this issue time and again. For example, say there is a machine on a factory floor that analyzes the quality of an auto part that it makes. If the part is not up to specification, as determined by an optical scanner, then it's automatically rejected. While this keeps a human from looking at the part, and thus slowing down the process, it also takes a great deal of time. The system must transmit the data and image back to the centralized database and compute engine where a determination is made as to the success of the manufacturing process. Then the results are communicated back to the machine.

The cloud complicates this process even more. Instead of sending the data back to the data center, it is sent to a

remote server that can be thousands of miles away. To make matters worse, we send it over the open Internet. However, considering the amount of processing that needs to occur, the cloud may offer the best bang for the buck.

Overcoming the Latency Challenge

To address the latency problem, many suggest "computing at the edge." It's not a new concept, but it's something that was recently modernized. Computing at the edge pushes most of the data processes out to the edge of the network, close to the source. Then it's a matter of dividing the work between data and processing at the edge, versus data and processing in the centralized system.

The concept is to process the data that needs to quickly return to the device. In this case, it's the pass/fail data that indicates the success or failure of the physical manufacturing of the auto part. However, the data should also be cen-

Edge computing isn't new, but its cloud element makes all of the difference.

trally stored, and, ultimately, all of the data sent back to the centralized system, cloud or not, for permanent storage and future processing.

Edge processing means that we replicate processing and data storage that's close to the source. But it's more of a master/slave type of architecture, where the centralized system ultimately becomes the point of storage for all of the data, and the edge processing is merely a node of the centralized system.

To accommodate edge processing, we need to think a bit harder about how to build our IoT systems. That means more money and time must go into the design and development stages. However, the performance that well-designed IoT systems will provide to meet the real-time needs of IoT will more than justify the added complexity.

I suspect that computing at the edge architecture will become more popular as IoT becomes more popular. We'll get better at it, and purpose-built technologies will start to appear. Computing at the edge of an IoT architecture is something that should be on your radar, if IoT is in your future.

A Few Key Points to Remember

Edge computing is about putting processing and data near the end points. This saves the information from being transmitted from the point of consumption, such as a robot on a factory floor, back to centralized computing platforms, such as a public cloud.

The core benefit of edge computing is to reduce latency, and, as a result, increase performance of the complete system, end to end. Moreover, it lets you respond to critical data points more quickly, such as shutting down a jet engine

that's overheating, without having to check in with a central process.

Although this latency reduction can aid all types of systems, it's mostly applicable to remote data processing, such as IoT devices.

Edge computing is not about snapping off parts of systems and placing them at the edge, but rather about the ability to look at data processing as a set of tiered components that interact, one to another, each playing a specific role.

The data that's processed and stored at the edge typically only resides there temporarily. It's ultimately moved to centralized processing, such as a public cloud, at certain intervals. That central location's copy becomes the data of record, or the single source of truth.

Don't do edge computing unless you have a specific need for it. Edge computing is a specialized approach to solving specialized problems. Enterprises are often guilty of adopting technology just because it's mentioned more than once in the tech press. But doing so will cost you more money and add risk — and edge computing falls into this category.


So, What does this All Mean?

Edge computing is a tactical way to solve the latency issues, built upon many tried-and-true architectures of the past. However, what's new is the element of the cloud, and the ability to leverage edge systems as if they were centralized. The new cloud element is bringing new relevance to edge computing.

Written by David Linthicum, former SVP, Cloud Technology Partners (CTP) and currently Chief Cloud Strategy Officer, Deloitte



IOT FOR THE ENTERPRISE: FROM PROTOTYPES TO PRODUCTION DEPLOYMENTS



Business and technology leaders must provide a clear set of processes for IoT solutions to move from concept all the way to a deployed and supported solution.



“Innovation leaders want to see what IoT can do for their business, while ensuring their investment dollars aren’t simply delivering them expensive, one-of-a-kind prototypes that don’t lend themselves to real world deployment.”

The Internet of Things (IoT) is one of those technological waves that is both exciting to experiment with and potentially frightening to miss out on. This often leads to what a colleague of mine refers to as ‘random acts of IoT’, accomplished while companies try to figure out what it means to their business.

While this experimental approach is traditionally good for validating concepts or ideas of how IoT can solve a business problem or be used to monetize opportunities outside the immediate customer base, many companies fail to move from the prototype or Proof of Concept (PoC) stage to a full-blown deployment. Why? Most companies fail to design an enterprise-level enablement and deployment framework for these experimental operations. This framework should be jointly created by business and technology leaders and provide a clear set of processes for IoT solutions to move from concept all the way to a deployed and supported solution. Many times these IoT experiments only have support from one side of the enterprise, either technology (majority of cases) or business, but rarely both. Because of this, IoT is often seen to be wholly owned by technology and lacks a direct connection to a business problem until further in its lifecycle when technology needs to justify the solution before continuing.

The IoT Sandbox

At the most successful IoT-enabled companies, you see a clear joint ownership and strategy model between technology and business. The IoT ‘sandbox’ that is created is based on the business’ needs and compliance requirements, as well as products and services that have been initially validated by the technology side of the house. Once the sandbox is up and running, the experiments deployed are prioritized by the continuous and transparent involvement of the business, and the overall enterprise strategy. From

a technology perspective, this involvement is ideal because the technology side has a line of sight into the business problems they are looking to solve with IoT. Furthermore, they know that if their solution is successful, then it has a clear path to production deployment. This is big for cultures looking to promote agility, innovation, risk taking and making sure team members can see the benefits of their efforts.

Managing Your IoT Transformation

Innovation leaders want to see what IoT can do for their business, while ensuring their investment dollars aren't simply delivering them expensive, one-of-a-kind prototypes that don't lend themselves to real world deployment. To do this effectively, you need to build an enterprise IoT enablement and deployment framework. This means engaging with leadership from business development, legal, sales, finance, operations, IT, etc. to get alignment and ensure that each stakeholder can ask questions and have their feedback incorporated in the acceptance criteria and deployment processes for these IoT projects.

For example, IoT deployments can cause huge headaches for legal and finance, based on their need for both capital and operational expenditures in various countries, data collection and distribution regulation, global data privacy compliance, deployment labor costs and many more complex aspects. This can quickly derail or delay a promising IoT solution when tangential teams are surprised when they get a laundry list of requests and requirements that force last-minute, heroic efforts to tackle or alter practices that are counter to their preferred process. By involving these departments earlier in the framework creation process, you are able to discuss concerns preemptively and work on mitigating them before they become a huge barrier to your IoT project's company-wide adoption.

This may seem obvious, but often technologists forget that legal and finance teams are some of the most influential gatekeepers and can make or break your project's success. Once they have approved your framework for building and migrating IoT solutions there won't be any surprises. Although IoT initiatives may run into many of the same issues that other technology projects do, the novelty of IoT sets it apart.

From PoC to Deployment

Aside from ensuring that the business has early and frequent contributions to your IoT enablement and deployment framework, there are a number of things that many technology organizations can do to further reduce the effort required to propel a PoC into an actual deployment.

First, create an IoT sandbox with proper guardrails and controls so that anything built within the sandbox is known to comply with your company's data governance, security and privacy policies, and can be migrated and scaled without having to do significant refactoring and compliance auditing. Second, tackle the non-functional requirements from the get-go so you can greatly increase the speed of your IoT efforts. It will require

a significant investment up front to build the proper IoT sandbox and the enterprise framework, but you will reap the dividends quickly as your lines of business are able to explore IoT solutions within the safety of the sandbox. Once a viable solution has been identified, the team will have the proper processes to get it approved for a wide-scale production deployment.

Victory Through a CBO

We encourage our clients to create a Cloud Business Office (CBO) to serve as the central point of decision-making and communication for your initiatives. It is critical to approach IoT solutions with a holistic view by engaging the business and technology groups to create an IoT enablement and deployment framework. This specific IoT framework can live within the CBO, or be its own entity, but it is important to lay the proper foundation and guardrails for your organization to safely and effectively explore IoT solutions. Without a CBO, having various IoT environments and states of compliance and testing can bring IoT adoption to a crawl while you try to unify your various efforts and mitigate internal fallout. In today's fast-paced technology environment the Latin proverb still holds true: *Amat victoria curam* (victory loves diligence).

Written by Jeff Maynard, former VP, Digital Innovation, Cloud Technology Partners (CTP) and currently Principal Product Manager, AWS IoT



Bringing IoT to Your Organization

With the emergence of smart sensors and connected devices, the Internet of Things has surpassed becoming “the next big thing” and claimed its place as a key disruptive technology affecting nearly all enterprises today.

HPE delivers a broad range of IoT consulting and implementation services that use the power of the cloud to solve today's most challenging business problems.

Learn how HPE can help design, build and implement your next IoT project:

www.hpe.com/loT



GO BIG ON DEVOPS: TRANSFORMING THE ENTERPRISE BEYOND AUTOMATION

Going big on DevOps means more than just IT automation and CI/CD. It's about transforming your people, processes and technologies.



One of the biggest value propositions of the cloud is agility. Enterprises are leveraging the public cloud to accelerate delivery of new services and features at speeds not seen before. But in order to achieve agility, enterprises need to do more than just learn how to leverage cloud services. They also need to assess their legacy baggage and transform the way they think about software and delivery.

Legacy baggage comes in many forms. Many of the vendor solutions that enterprises have embraced in the era of the data center do not have a viable equivalent in the public cloud. Many legacy processes were derived in the days of biannual releases, which get in the way of progress in the era of continuous delivery. To be agile, legacy organizational structures, built to control costs and dictate policy must give way to new organizational structures that promote high levels of collaboration, transparency and shared goals.

As I have written in the past — and many of my colleagues agree — agility is the new currency. Cloud is one part of the equation. Getting out of the data center business and embracing cloud APIs is a great step forward. However, embracing the cloud without embracing the need for transformational change is a losing proposition.

To be agile, legacy organizational structures, built to control costs and dictate policy must give way to new organizational structures that promote high levels of collaboration, transparency and shared goals.

The Solution: Go Big on DevOps

Many people equate DevOps to IT Automation or even CI/CD. Automation and CI/CD are components of DevOps that you will typically incorporate into your DevOps journey but they are not DevOps in and of themselves. Our definition of DevOps is “a culture shift or a movement that encourages great communication and collaboration to foster building better quality software more quickly with more reliability.” DevOps is the progression of the software development lifecycle (SDLC) from Waterfall to Agile to Lean. When we perform DevOps maturity assessments with our clients, we assess their maturity across three spectrums: People, Process and Technology.

DevOps and People

Within the people domain, we evaluate enterprises in four categories: innovations, skills, culture and collaboration. Innovation is a much overlooked area when it comes to transformation. In the days of physical hardware, it was often unfeasible to invest the time

DevOps

"A culture shift or a movement that encourages great communication and collaboration to foster building better quality software more quickly with more reliability."

and money it took to procure or create new environments to try out experiments or a hypothesis. Many enterprises outlawed experimenting with new technologies without going through a proper vetting process. This mindset often stifled innovation. I have interviewed clients who have told me they stopped trying out new ideas because it was so painful to go through the processes, or because it was discouraged due to perceived risks. Employees would experiment at home but could not bring these learnings to work. Stifling innovation often leads to employee turnover because the best and the brightest love to innovate.

In the pre-cloud days, embracing new technologies was a lot harder than it is today. For example, let's say an engineer determines that for a particular workload, a non-relational database like Hadoop is a better solution than the existing Oracle relational database that he or she is forced to use. The engineer would have to go through a committee to get approval to bring in a new vendor. Hardware and software licenses would have to be procured. People would have to be trained or hired to manage and maintain this new technology, and so on. All of this takes weeks or months to discuss, plan and execute, which almost always exceeds the time that the engineer has to complete the task. After going through this process once or twice to no avail, the engineer will not even bother to bring up new solutions anymore.

In the cloud, databases like Hadoop are available as a managed service with a simple API call. No hardware or licenses to procure, no army of people to manage and maintain, and very little time and money needed to experiment. If allowed, the engineer could test his or her hypothesis the same day without involving any other people. If the experiment proved that Hadoop was indeed the best solution for the specific workload, then the engineer could work with others to make it an official solution for the enterprise. In order

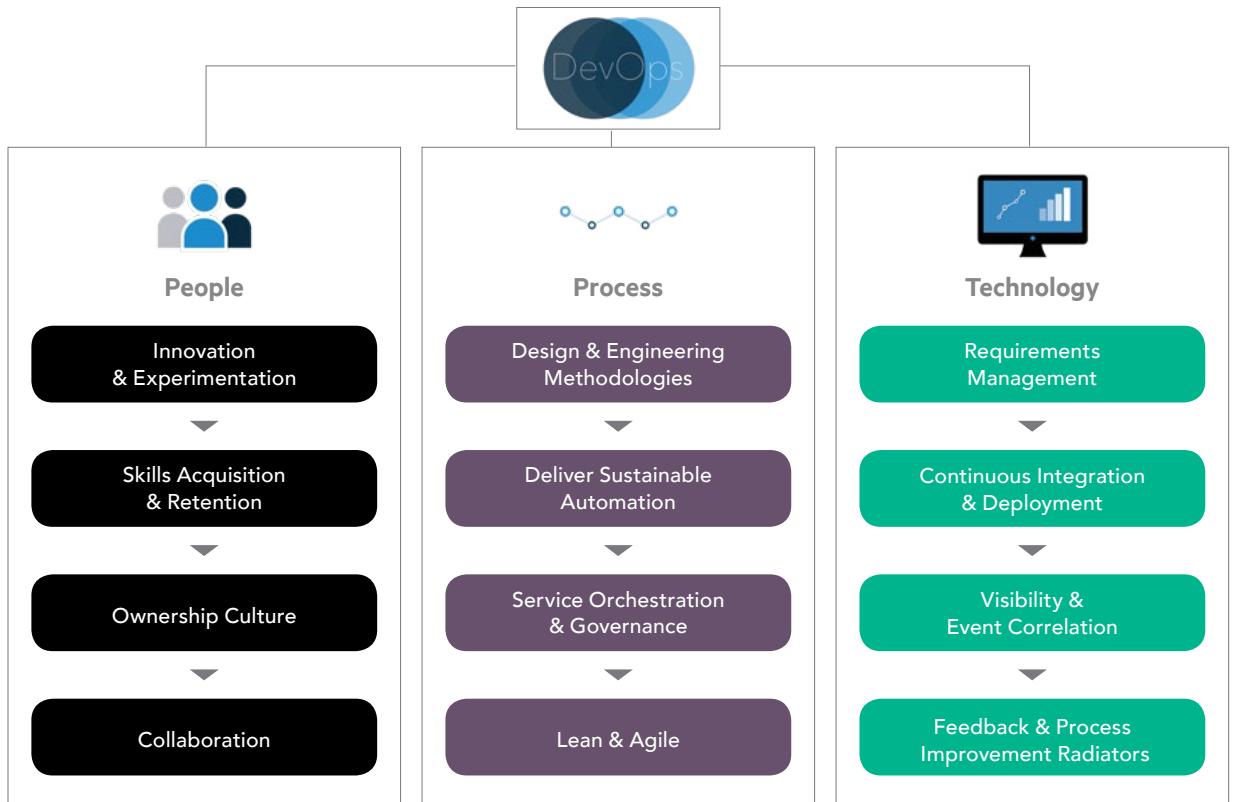


Figure 1: HPE's DevOps Roadmap for Cloud

to take advantage of the innovation capabilities that the cloud offers, people need to embrace innovation as an enabler, as opposed to it being a risk item to be carefully controlled.

People skills is another area we evaluate. Not only do people need to be trained to use the cloud provider's services, but they also must learn new methods and approaches required to take advantage of the cloud. Too often enterprises treat the cloud as just another data center instead of what it really is: a game changing agility platform. In addition to skills, we look to ensure that employees' incentives are aligned with the new way of doing things in the cloud. If incentives don't change, how can we expect people to transform to the new way of building services in the cloud?

Silos are a big blocker for any transformation. We look at how different silos collaborate and whether they work together in unison or work against each other. Traditionally, silos owned singular parts of the software development lifecycle (SDLC). One team owned development, another QA, a third owned Ops, a fourth security, etc. In the DevOps model, these are all shared responsibilities. Everyone owns security. Everyone owns quality. And not just IT people. The product owner and the business sponsor also share the ownership. When ownership is shared, everyone works towards a common goal.


DevOps and Processes


One common mistake I see in almost every engagement that involves legacy processes is a lack of focus and analysis of those legacy processes before proceeding with automation. As a result, enterprises are automating waste and not realizing the benefits of agility they were expecting. This is often the result of silos and mismatched incentives across those silos. Developers will work in their own silo and implement CI/CD. They see great improvements in both time and quality of their build process but often see little to no change to their time to market metrics. Why is that? Because there are 20 to 30 years of legacy change control processes created in the era of biannual releases that have not been addressed.

I have seen numerous instances where red tape can mire the process for weeks or months prior to being able to perform an automated build, followed by more weeks and months of red tape in order to promote the code to production after the automated build. Yet teams still focus solely on perfecting CI/CD.

If you have ever read Eliyahu Goldratt's *The Goal*, you will have learned that working on the wrong bottlenecks does not improve the overall flow of a system. Instead it moves the bottleneck to another part of the system. If enterprises only implement CI/CD without performing a value stream assessment of the complete system, they will only move bottlenecks from the build process to another part of the system, thus never achieving the desired agility. Engineers must think about the system as a whole instead of just focusing on automating a component of the system. System thinking can be a foreign topic to a silo-based organization.

Another area we look at is the SDLC practices. Enterprises planning to move to frequent deployments need to embrace lean principles and move away from traditional waterfall or immature scrum approaches. Governance is another important area. The old method





Source: Damon Edwards

of governing with an iron fist needs to give way to baking controls, policy and governance into the code. The days of holding multiple weekly review boards for architecture, security, and governance must be put to bed. These processes and mindsets simply don't work in the era of continuous deployment. In this new age, we must trust in our automation and institute proactive and continuous monitoring to check for ongoing security and compliance. Manual review by humans just doesn't scale when multiple teams are able to perform push-button deployments. We must audit ourselves in real-time in the new world.

DevOps and Technology

It is here where we finally start focusing on IT automation and the famous CI/CD processes. What many call DevOps is just one small piece of the DevOps puzzle. Running systems in the cloud requires new tooling and methods. Many of the legacy tools we have used in the past require state and physical infrastructure. We recommend born-in-the-cloud solutions in the areas of security, monitoring, logging, code repositories, etc. Providing visibility into system health and application state is crucial in providing high SLAs in the new world where deployments happen frequently. Much thought needs to go into building a robust security and monitoring framework that feeds into a central logging solution and can be accessed through a single pane of glass.

The build process should perform security and coding standards scans. Testing should be automated and part of the build process. The build process should produce a score for security, programming standards and quality. The build should fail if any one of those scores are not at an acceptable level. The goal of this approach is to not let issues progress downstream, because it is much more expensive and time consuming to fix defects later in the lifecycle.

DevOps Maturity

After evaluating our client's capabilities in the area of people, process and technology, we provide a maturity score.

The score is a snapshot of the client's current state of maturity. Next, we provide a list of gaps in each area that shows the delta between the current state and the desired future state. Clients often want to start with a maturity of level 3 so that they can get to a consistent, secure and reliable state for deployments, while achieving a higher level of agility.

ACME Maturity Model Score: 1.74




Maturity Level	People	Process	Technology
Level 1 Ad-Hoc	<ul style="list-style-type: none"> • Silo based • Blame, finger pointing • Dependent experts • Lack of accountability 	<ul style="list-style-type: none"> • Manual process • Tribal knowledge is the norm • Unpredictably reactive 	<ul style="list-style-type: none"> • Manual builds and deployments • Manual testing • Environment inconsistencies 
Level 2 Repeatable	<ul style="list-style-type: none"> • Managed communications • Limited knowledge sharing 	<ul style="list-style-type: none"> • Processes established within silos • Can repeat what is known, but can't react to unknowns 	<ul style="list-style-type: none"> • Automated builds • Automated tests written as part of story development • Painful but repeatable releases 
Level 3 Defined	<ul style="list-style-type: none"> • Collaboration exists • Shared decision making • Shared accountability 	<ul style="list-style-type: none"> • Processes are automated • Standards across organization 	<ul style="list-style-type: none"> • Automated build & test cycle for every commit • Push button deployments • Automated user & acceptance testing
Level 4 Measured	<ul style="list-style-type: none"> • Collaboration backed on shared metrics with a focus on identifying bottlenecks 	<ul style="list-style-type: none"> • Proactive monitoring • Metrics collected and analyzed against business goals • Visibility & predictability 	<ul style="list-style-type: none"> • Build metrics visible and acted on • Orchestrated deployments with auto rollbacks • Non-functional requirements defined and measured
Level 5 Optimized	<ul style="list-style-type: none"> • A culture of continuous improvement permeates through the organization 	<ul style="list-style-type: none"> • Self service automation • Risk & cost optimization • High degree of experimentation 	<ul style="list-style-type: none"> • Zero downtime deployments • Immutable infrastructure • Actively enforce resiliency by forcing failures

Figure 2: Maturity Model Matrix

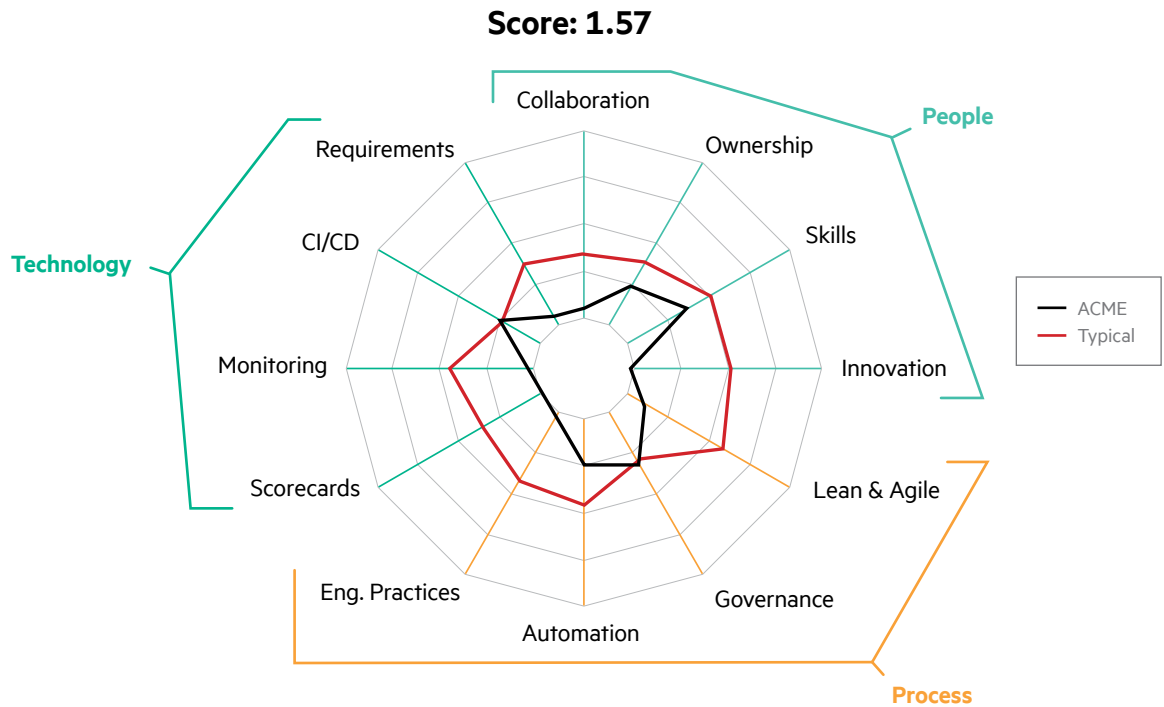


Figure 3: DevOps Maturity Model Score

The score provides us a place to begin the conversation, but where the rubber meets the road is in the details behind the score. We provide a list of gaps, in a roadmap format, and the recommendations for each gap. The number of gaps can be quite overwhelming, especially if a customer is early in their journey. But it is this roadmap that brings clients the most value.

What bottlenecks should they work on first and in what order? How can they work on people process, and technology changes concurrently? Too often we see clients only implement the technology recommendations and make little to no progress on the people and process recommendations. The end result is a suboptimal experience in the cloud and a missed opportunity to achieve the ROI that could be achieved by treating the cloud journey as a transformational project, as opposed to just a technology project.

Summary

Cloud computing is providing enterprises with the capabilities to greatly improve their overall agility. To take advantage of this opportunity we recommend that enterprises embrace the DevOps mindset. We also reiterate that DevOps is more than just IT automation and CI/CD. Embracing the cloud calls for a complete transformation of an enterprise's people, processes and technologies. Looking at the cloud any other way will diminish the returns on your investment.

Written by Mike Kavis, former VP, Strategic Accounts, Cloud Technology Partners (CTP) and currently Managing Director at Deloitte



DATA PERSISTENCE IN A CONTAINERIZED ENVIRONMENT

Four data handling lessons to consider when moving containers to cloud.



Containers have become an effective tool to facilitate rapid migration of applications, often to streamline the movement from on-premises to cloud infrastructure. Containers offer many benefits to increase the speed of migration by decoupling application dependencies from underlying hardware, operating systems, network connectivity and systems management tools. However, while containers have become a common tool for facilitating migration, they demand architectural changes to ensure proper data integrity, application scalability and recoverability.

Many organizations use containers as a lifeboat to accelerate the migration of applications to the cloud, creating a base architecture that can then evolve over time to become more and more cloud-native. The challenge is that for applications to function properly, they need changes from the architectural principles used 10, 15 and 20 years ago when many enterprise applications were first developed.

Let's take a look at some lessons learned about data handling patterns when moving containers to cloud:

Lesson 1: One of the design choices that many organizations use is to “just add a volume” to containers and trick them into acting like persistent data stores. This is an anti-pattern and treats containers almost as if they were virtual machines. This runs counter to the container model of being lightweight, easy to move around and replaceable as containers fail, are removed and upgraded in flight.

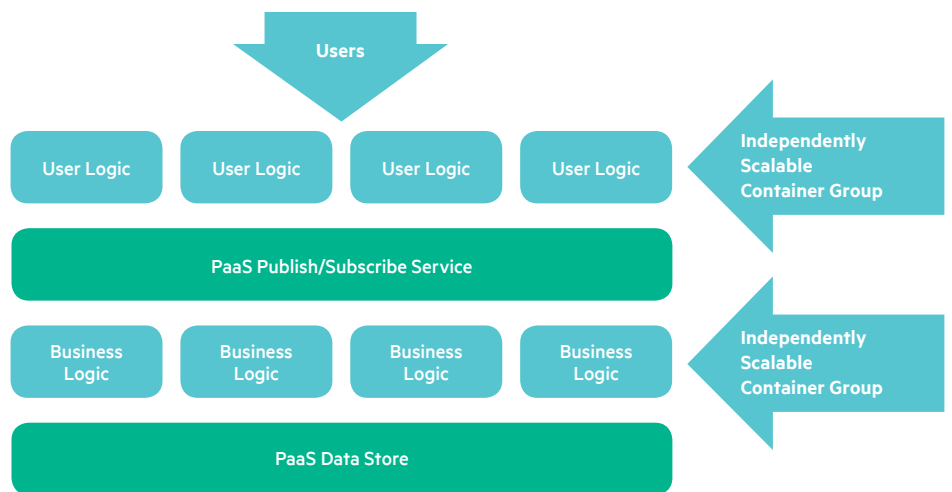


Figure 1: Functionality in a Containerized Application Stack

Lesson 2: Many organizations start their journey by taking traditional RDBMS technologies and putting them into containers. However, this practice creates operational issues and makes the movement of containers difficult due to growing volumes. This approach also risks encouraging embedding of business logic into the database. When using a container approach, the business logic should be outside the database to facilitate scaling independent from other application functions, but also to allow flexibility in the choice of data store to be used. Business logic should ideally be integrated with other aspects of application functionality through a publish/subscribe model, enabling scalability at different tiers of the application as workloads evolve.

Lesson 3: Application state should never be stored in the containers. Containers are transient resources that cannot be relied upon to store persistent information. Instead a data store layer should be used to track state that is necessary to manage workflows and user experience.

Lesson 4: Messaging and workflow should be handled by services outside the containers and application logic. This will ensure that services evolution and scaling is accomplished in a manner that is independent from the application functionality. It also enables more robust and reliable error and retry handling.

These lessons show the importance of respecting two core cloud principles when bringing containers to cloud: separating compute from storage, and the emphasis on decomposing monolithic applications into smaller, loosely-coupled functional modules. When combined with the ephemeral “deploy when needed, destroy when done” capabilities of containers, these lessons enable the provisioning of architectures that feature the best inherent characteristics of both cloud and containers.



By using cloud-based data services to fulfill persistent data storage needs for containers, the major drawbacks of “data in the container” (lack of independent scaling), Docker volumes (tied to specific machines), or creating “data containers” (limited flexibility) are avoided. By using a PaaS data store, the cloud provider manages the operational aspects of data protection and scalability, and ensures containers can focus on application logic, presentation and user service delivery. AWS’ Elastic File System, as an example, enables a shared, mountable, automatically scalable data environment well-suited for persisting various types of container data.

Containers can be used as an effective way to package applications for rapid deployment, upgrades and lifecycle management. However, when used incorrectly, containers can become another bloated virtual machine replacement with little to no value-add. It is important that as part of a container implementation strategy, the persistence of data and the movement of data is put in the proper functional areas to facilitate scalability and operations. By leveraging PaaS services for communication and data stores, containers can be focused on the execution of business logic while ensuring robust persistent connections and interactions with the necessary data.

Written by Joey Jablonski, former CTO of Cloud Technology Partners (CTP) and currently VP of Data Engineering & Analytics at iHeartMedia. and Neal Matthews, Principal Architect, Cloud Technology Partners (CTP)

When used incorrectly, containers can become another bloated virtual machine replacement with little to no value add.



RUNNING HYBRID CO EFFICIENT



CLOUD MAY

- 74** Five Essentials for Managing Hybrid Cloud
- 78** Go Big on Security: Five Steps to Overcome Security Concerns in the Cloud
- 84** Why You Need an Enterprise Cloud Managed Service Provider
- 88** Getting Compliance Under Control
- 92** HPE's Evolution to a Hybrid Cloud Environment

Optimizing cloud operations with holistic visibility and control will drive efficiencies and free resources.



FIVE ESSENTIALS FOR MANAGING HYBRID CLOUD

Hybrid cloud management is more of an art than a science, and the complexity will likely increase over time. This means that our approaches to hybrid cloud management need to evolve as well.

If you use cloud computing, chances are good that you leverage some sort of hybrid cloud as part of your environment. Until recently, this usually meant pairing a private and public cloud, such as OpenStack and Amazon Web Services. Today, the term “hybrid” typically means legacy, or traditional, systems paired with one or more public clouds, a single private cloud that interfaces with two or more public clouds, or any combination thereof.

Things get complex, quickly. As enterprises attempt to figure out the best approaches to security, governance, and management, they’re finding that no single approach or tool can solve all problems. So, if an enterprise IT shop has a hybrid cloud or, more likely, more than one hybrid cloud, what should it do about management?

The first step is to understand the essentials. When defining them, it’s not about just picking the tools and technology you’ll use. Most IT operations managers make the mistake of focusing on tools that may make hybrid cloud management easier, rather than understanding their own requirements, which means they get both the approach and the tool selection wrong. You need to understand the security, data, governance, and end-user dynamics that affect how you approach hybrid cloud management.

Here are five concepts that you need to understand before moving to a hybrid cloud management platform.

1. Understand what’s being managed

While this seems like the single most important thing to know, many of those who define a hybrid cloud management strategy fail to understand the profiles of the workloads that will run on public and private cloud(s). You need to understand what the applications do, including how they interact with the end users, manage data, how they handle networking, security patterns, performance, etc.

Specific things need to be understood, as follows:

- Who owns the workload within the organization? Who needs to be contacted when things go south?
- What do the workloads do for the business, per their criticality to the business? This goes to how many resources you spend on managing the workloads on the hybrid cloud, which need to align with the value they bring to the business.
- When do the workloads run? Some run continuously, while others may run during the same hour in the day. Again, this goes to how you approach the management of workloads within the hybrid cloud.
- Where do the workloads run? On the public cloud, private cloud, or in both places?
- Why were the decisions made about where to run the workload? And when may they need to be re-evaluated?

2. Understand security and governance

These days, security and governance are a requirement, whether a mandate from your customers (see “SLAs”) or from your senior management. This means you need to proactively manage security to make it work. You can also leverage new mechanisms such as identity and access management (IAM), which allow assigning of identities to data, people, devices, and servers, to configure who can access what, and when. Finally, information needs to be encrypted at-rest in some cases, and in-flight in others.

Core to this part of hybrid cloud management is how you deal with a few issues:

- **Security and performance.** If the needs of the workload are that information be encrypted at-rest (on the storage systems in the private or public cloud), or in-flight (moving over the network), that may result in the risk of lower overall performance. That needs to be understood and managed, including the use of performance monitoring tools.
- **Policy management.** Governance requires that policies are written and enforced, and this enforcement needs to be understood by those who are managing the hybrid cloud so that they do not conflict or otherwise get in the way of operations.

3. Build a “single pane of glass”

Those who manage hybrid cloud manage complexity, because the private and public clouds all come with their own native APIs and resources. Indeed, they all manage storage, networking, provisioning and security differently. Thus, you can either learn all of the native interfaces for all private and public clouds, or you can instead build a single pane of glass that abstracts you away from that complexity.

There are tools that can manage the cloud services using a single interface to translate what something means on one cloud versus another cloud. For instance, you need to monitor performance on Google Cloud Platform, and OpenStack private cloud, and Amazon Web Services. All provide different approaches and interfaces to manage performance, and the single pane of glass interface deals with the differences on your behalf, translating what’s important to those who manage the hybrid cloud in and between the different clouds that are under management.

4. Understand the SLAs

SLAs, or Service Level Agreements, are a contract with the end users stating that you, the hybrid cloud manager, and the cloud providers themselves, will supply a specific level of service, else there will be penalties. While you can certainly pass the buck to the public cloud provider in living up to their own SLAs, the hybrid cloud itself is your baby, and thus you’ll be held responsible if the system misses the limits outlined in the SLAs you’ve agreed to.

At a high level, what’s defined in the SLA needs to be defined in the management layer as well. It’s not just about providing a baseline of good performance to the end users, but it’s about providing performance that meets specific expectations. For instance, the ability to provide a sub-second response to the sales person leveraging the inventory application that exists within the hybrid cloud.

When it comes to hybrid cloud management, SLAs are not legal tools. But they are a way to define user and business expectations. Thus, it’s easy to leverage these expectations to define the service expectations that need to be managed by the hybrid cloud management layer, and the hybrid cloud managers. Use them as guidelines.

5. Understand the tools

Many charged with hybrid cloud management often focus too much on the management tools that are available. These tools cover areas such as API management, resource management, cloud management platforms, performance management, DevOps management, security management, network management, native platform management, etc.

There can be as many as a dozen or so tools that you'll need if you're going to manage a hybrid cloud effectively. Picking tools is a matter of understanding what we've defined above as the requirements patterns, and then figuring out the solutions patterns that will match. For instance, if your requirement is that you encrypt all data at-rest or in-flight, and you need to manage both the encryption and the performance, then you would look for a tool or tools that provide both types of solutions.

It's then a matter of working through the requirements, defining the common patterns, and then matching up the tools with the pattern. Sometimes you won't find all of the tools that you need for all of the patterns. In those cases, you may consider custom tools. Or, perhaps a hybrid cloud is not in the cards after all, if management of those workloads on those platforms can't be done using automation through management tools. In other words, in some cases, it's OK to leave a workload off a hybrid cloud if they can't be managed properly or at a reasonable cost.

Hybrid Cloud Growth and Adoption Patterns

The issue with hybrid cloud is that the analysts and technology providers are tossing many things into that category. For our purposes, we can call a "hybrid cloud" any collection of systems where you have at least one private cloud and one public cloud working together to support systems for IT.

However, managers also need to understand the emerging notion of the "pragmatic hybrid cloud." This is a traditional set of systems, typically running in a data center, paired with at least one public cloud. This configuration is growing in popularity, as many enterprises continue the migration to cloud, but they do so without leveraging a private cloud. Instead, they pair their traditional systems with public cloud-based systems.

The emerging patterns are ones of increasing complexity. The right way to manage these combinations is not straightforward, because management needs to deal with the increasing complexity as well. This presents a new problem area, because one of the core management reasons for hybrid cloud adoption is to hide complexity behind an abstraction layer. In fact, instead of avoiding the complexities associated with hybrid cloud, managers need to better understand the native features and capabilities of the public or private clouds that make up the hybrid cloud.

Cloud Management: Still More Art Than Science

Hybrid cloud management is still more of an art than a science, considering that we've not been at it for very long. The complexity will likely increase over time, and that will mean that our approaches to hybrid cloud management need to evolve as well.

For now, it's a good time to understand just what you're dealing with in terms of technology. If you're charged with management of these platforms, you have your work cut out.

Written by David Linthicum, former SVP, Cloud Technology Partners (CTP) and currently Chief Cloud Strategy Officer, Deloitte



GO BIG ON SECURITY: 5 STEPS TO OVERCOME SECURITY CONCERNS IN THE CLOUD

With a bit of understanding and planning, there's no reason your data won't be as secure in the public cloud as it is in the most on-premises systems. Actually, your public cloud security can surpass your on-premises security.

Here's the reality: Public cloud providers have better security mechanisms in place and are more paranoid — and attentive — to security risks throughout their entire stack. Considering the paranoia around cloud computing and security, most public cloud-based systems have better thought-out security mechanisms than those in traditional data centers.

So, why do many enterprises still distrust the cloud?

What public clouds bring to the table are better security mechanisms and paranoia as a default, given how juicy they are as targets. The cloud providers are much better at systemic security services, such as looking out for attacks using pattern matching technology and even AI systems. This combination means they have very secure systems.

It should be no surprise that the hackers move on to easier pickings: enterprise data centers.

The on-premises systems that IT manages are typically a mix of technologies from different eras. The aging infrastructure is often less secure – and less securable – than the modern

technology used by cloud providers. This is simply because the old, on-premises technology was designed for an earlier era of less sophisticated threats. The mixture of different technologies in the typical on-premises data center also opens up more gaps for hackers to exploit.

Because on-premises systems continue to age, their intrinsic security can be easily defeated by hackers. Moreover, the number of attacks increase weekly, and defenses need to be proactive – more proactive than most enterprise IT organizations are, and likely more proactive than they can each afford to be.

So, why have you not gone big with cloud security? Perhaps it's just a lack of knowledge about the process to secure your public cloud-based data store or application.

Here are 5 easy steps to cloud computing security.

Step 1: Understand your true requirements.

As we implement cloud-based systems and wrap them with the right security approaches and technology, the largest

issue I find is that few in the enterprise understand the true security requirements. Typically, they have notions about the legal and compliance issues around the protection of corporate and government data that are not based in reality.

Things that need to be reviewed in detail include any laws or regulations that require compliance, and thus what technology is mandated (e.g., encryption levels or location of data). Moreover, existing internal policies around the protection of data, including the existing approaches for evaluating risk, must be identified. These should be written down and approved by leadership so everything is clear and well understood.


Step 2: Consider identity-based security.

The best approach to cloud computing security requires that we deal with all assets, including humans, servers, databases, data, processes, services, etc., as identities. These identities can then be managed, in terms of access to resources, and as resources themselves. The application of identity-based security to cloud computing is quickly emerging. The most successful and useful cloud security systems are able to manage fine-grained identities to control when and how they interact.

Step 3: Create a plan.

Many consider security to be one of those things that gets added in the final hours of deployment or migration. The reality is that approaching security in general – and cloud specifically – requires that a master security plan emerge using the requirements we've gathered in Step 1. Keep in mind, security is systemic to cloud computing. It's a part of every step in the plan.

This drives down to the actual solutions, including solution patterns and candidate technology that should be evaluated as a potential fit. Many in IT approach security technology with a bias toward their favorite or existing solutions. Don't lock yourself into a technology until you've understood the requirements, and tested the technology.



Keep in mind, security is systemic to cloud computing. It's a part of every step in the plan.

Step 4: Select the right security technology.

Goes without saying, right? However, most of those who implement security technologies never test it before the implementation. Many take the vendor or cloud provider's word for things, which is a huge mistake.

POC testing is mandatory. You should go into deployment with no questions unanswered.



Step 5: Deploy, test, monitor.

Deploy the security solution with the understanding that it is not a separate entity from the core system or the data but is bound to it. Many think they can decouple security from the core processes and data, but that is just not the case.

Make sure to test the security. Many firms provide “white hat” penetration testing, and a few good weeks of that type of testing will provide some good insurance that the solution

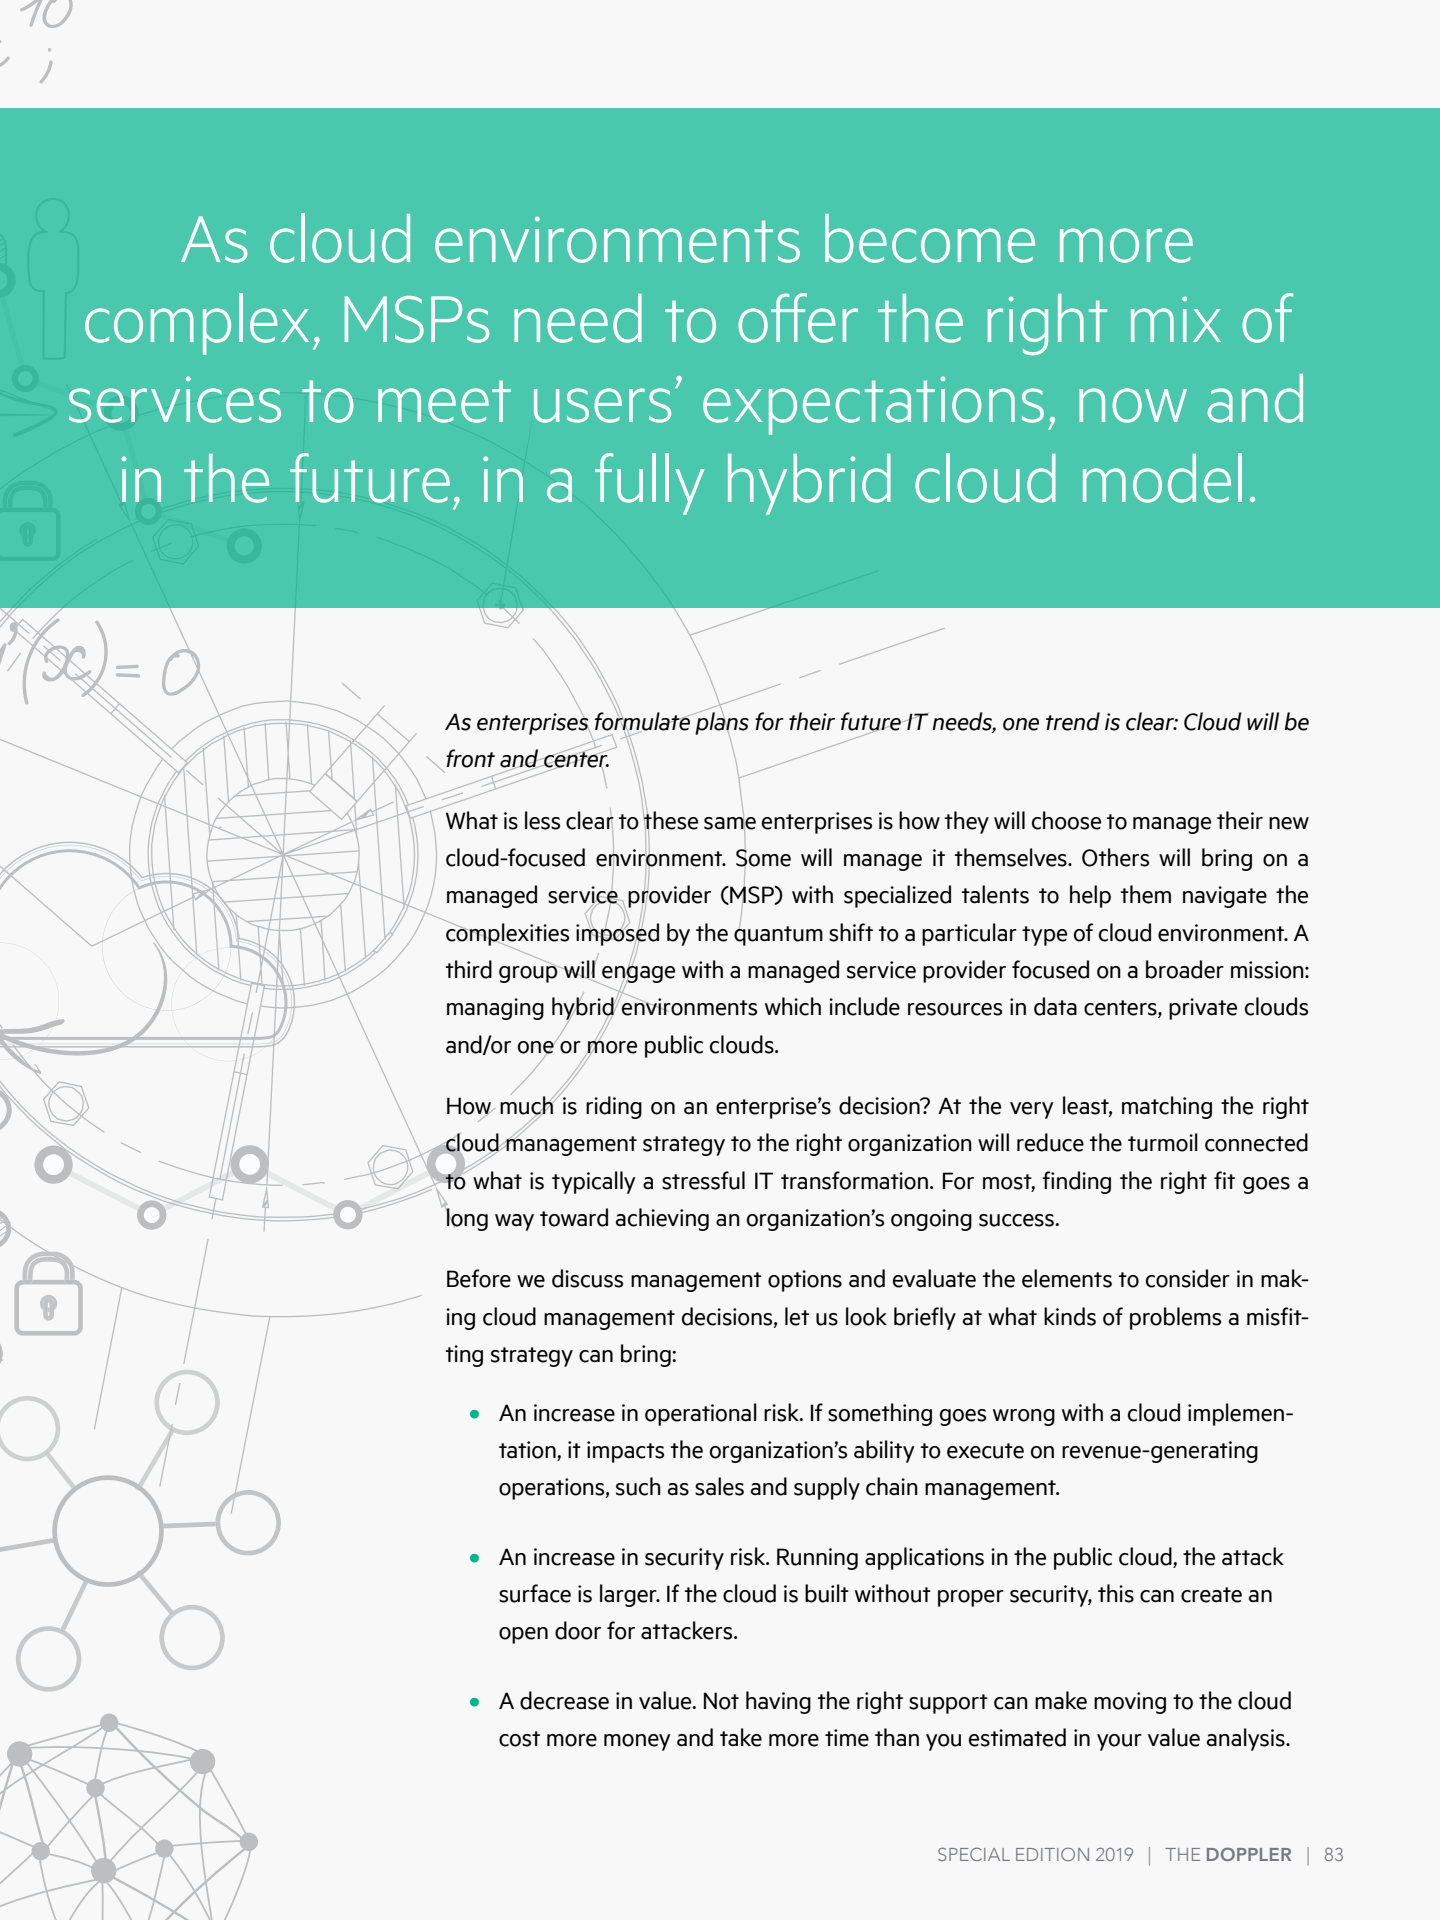
works, or perhaps it will point out the need for some additional configuration. Finally, understand that monitoring is required over time.

Easy steps? Yes. Needs some thinking? You bet. However, by following the steps above you’ll find that your cloud-based system is invariably more secure than anything currently in your data center. It’s tough for even those who don’t like the use of public clouds to push back on that fact.

Written by David Linthicum

WHY YOU NEED AN ENTERPRISE CLOUD MSP





As cloud environments become more complex, MSPs need to offer the right mix of services to meet users' expectations, now and in the future, in a fully hybrid cloud model.

As enterprises formulate plans for their future IT needs, one trend is clear: Cloud will be front and center.

What is less clear to these same enterprises is how they will choose to manage their new cloud-focused environment. Some will manage it themselves. Others will bring on a managed service provider (MSP) with specialized talents to help them navigate the complexities imposed by the quantum shift to a particular type of cloud environment. A third group will engage with a managed service provider focused on a broader mission: managing hybrid environments which include resources in data centers, private clouds and/or one or more public clouds.

How much is riding on an enterprise's decision? At the very least, matching the right cloud management strategy to the right organization will reduce the turmoil connected to what is typically a stressful IT transformation. For most, finding the right fit goes a long way toward achieving an organization's ongoing success.

Before we discuss management options and evaluate the elements to consider in making cloud management decisions, let us look briefly at what kinds of problems a misfitting strategy can bring:

- An increase in operational risk. If something goes wrong with a cloud implementation, it impacts the organization's ability to execute on revenue-generating operations, such as sales and supply chain management.
- An increase in security risk. Running applications in the public cloud, the attack surface is larger. If the cloud is built without proper security, this can create an open door for attackers.
- A decrease in value. Not having the right support can make moving to the cloud cost more money and take more time than you estimated in your value analysis.

In today's competitive business environment, these are problems enterprises do not want. Companies therefore need to consult with the right experts to make sure their cloud implementations go as smoothly as possible and position their businesses to meet future IT needs.

Cloud Management Options

Going Solo

Some organizations prefer to manage their cloud environments themselves. This can work, in certain situations. Some large companies insist on going it alone in order to retain control over their environments and their investments. Others will start with an MSP to get the necessary tools, training and methodology, and then manage the environment themselves. However, these arrangements usually take longer to achieve the value the customer wants. Practice makes perfect. An organization embarking on a cloud implementation for the very first time is taking a big chance learning on the job.

Working With a Niche MSP

As more enterprises move to the public cloud, many are engaging with service providers who specialize in migrating to and managing implementations by one of the three main cloud providers: AWS, Microsoft Azure or Google Cloud Platform. These MSPs usually have tooling and expertise in only one of the cloud choices – but what if the implementation involves more than one cloud, or resources in a data center environment? The majority of MSPs are still “boutique” cloud service providers, and many do not have the capabilities to service multiple IT environments.

Working With an Enterprise MSP

Experts roundly agree that the future will be dominated by hybrid IT environments, with hybrid cloud spending expected to nearly triple from 2016 to 2021. Data center environments will not go away entirely, housing workloads it makes no sense to move, while the rest of the workloads move to public and private clouds. This means future environments will become more complex, requiring help from MSPs who possess a broad range of skills, resources and management models.

To guide enterprises in the future, MSPs will have to branch out in a number of ways. They will need to manage on- and off-prem environments, multiple public clouds and enterprise-grade implementations. They will need to exercise flexible processes and methodologies. And they will need to pay close attention to compliance, cost controls and global IT demands.

Here are five elements enterprises should look for in a cloud MSP.

1. A hybrid model, applied consistently

If an MSP proposes multiple models for multiple private clouds, it will be very difficult to unify those operating models. Managing workloads in different clouds – and in data centers – in the same way reduces the risk of opening gaps that create security holes.

2. A focus on enterprise-level concerns in all places

An approach like our Minimum Viable Cloud (MVC) methodology gives an MSP the ability to build environments that are fully enterprise-enabled out of the gate. If an MSP offers low-ball pricing to stand up a simple environment to migrate particular workloads into, that environment will not be designed for enterprise-class workloads. Its security, auditability and compliance functions will not be operable at the level an enterprise needs.

3. A commitment to flexibility

Enterprises may have similar needs, but every implementation is different. You want your MSP to be consistent, but it should not be so locked down it cannot adapt to changing requirements. An enterprise may make tooling choices that have to be incorporated into the environment. If the MSP can accommodate the request, it should. If the customer wants to use its own tools, the MSP should have the resources to bring those tools into the automation framework, so they are not being hand deployed every time they are used. You do not want unlimited variability, but a rigid, cookie-cutter approach is never successful.

4. An enterprise-class focus on bedrock issues



At the enterprise level, you cannot afford to miss a beat on compliance and security. Both have to be in a state where they are being monitored regularly and efficiently. Problems can crop up quickly in the public cloud; environments get stood up and killed off all the time. An MSP cannot just show up and do an audit; it has to be managing all the time. If security and compliance are not in a good place, you will be flying blind.

5. Cost management expertise in a hybrid environment

Optimizing costs in a single cloud is a straight-forward challenge. Doing it in multiple environments is far more complicated and requires an enterprise-grade level of service. Enterprises care about multiple environments. If each environment is being managed individually for cost, the results will not be optimal.

Conclusion

When cloud was bursting on the scene, MSPs quickly tried to align with users' needs – providing services focused on a particular cloud or a particular type of implementation.

That is not enough anymore. Cloud environments are becoming more complex, and data centers still play a role. MSPs need to offer the right mix of services to meet users' expectations, now and in the future, in a fully hybrid cloud model.

Earlier this summer, we announced HPE GreenLake Hybrid Cloud, a comprehensive offering to help simplify our clients' hybrid cloud environments. This solution provides leading-edge operational and management services to span and unify your public and on-premises cloud models.

Regardless of which MSP solution is best for your organization, it needs to provide an integrated story and get you where you need to go by offering a targeted blend of advisory and professional services.

Learn more about HPE Greenlake Hybrid Cloud at www.hpe.com/greenlake

Written by John Treadway, SVP, CTP Strategy and Portfolio, HPE

GETTING CLOUD COMPLIANCE UNDER CONTROL

To gain control over compliance, organizations must first understand their scopes and their ability to handle those scopes.



When companies kept their applications in a data center, compliance was a more straightforward process. It still required energy and diligence, but the tasks were predictable. Servers and software were in the back room, paid for, running on set schedules, year after year. Workers maintained specific legacy systems that they were well trained on, configurations followed established patterns, and workloads were more easily tracked alongside company initiatives. Compliance could be handled as a quarterly or even annual ritual.

Cloud has flipped the compliance process upside down. It's introduced a whole new set of variables – new tools, new configuration and approval processes, new job roles and new rules for companies to follow. The changing environment has turned compliance into a moving target that's harder to control. Compliance can no longer be managed once or twice a year. In the cloud, compliance needs to be managed continuously.

To get cloud compliance under control, organizations must first understand their scopes and their ability to handle those scopes. The scope will vary for each organization, and even within an organization, based on issues such as: the regulations controls themselves; the complexities of requirements demanded by the industry; the geography; the impact to the business if it's out of compliance; and the level of cloud maturity and readiness to take on the job and do it well.

Let's look at these issues in more depth to evaluate how to get your cloud compliance under control.

The Impact of Cloud and Automation on Compliance

Looking more closely at cloud's impact, it's easy to see how challenged organizations are when it comes to maintaining control and, just as importantly, to demonstrating that they're maintaining control.

Above all else, cloud helps organizations improve their agility. They're not hidebound by server policies and schedules, so they make rapid and frequent changes to their environments. Cloud allows them to dial services up and down according to needs and desires, and to create and deploy software rapidly using continuous integration and continu-

ous delivery pipelines. Configurations that wouldn't change for months, perhaps years, in the data center now change in minutes.

The app delivery process used to be concentrated; in the cloud, it's decentralized. Many developers and DevOps personnel play a role in software delivery. Some may not have experience pushing changes to test or to other environments. This adds a layer of risk.

The different cloud environments create a layer of complexity. The trend today is for companies to embrace multiple cloud environments – such as AWS with a combination of Microsoft Azure and/or Google Cloud Platform, or other combinations. Each new tool and new environment, increases the learning curve for a staff that's already struggling to stay current in their training. Plus, the cloud providers themselves are constantly innovating, adding new services and new techniques.

Here's the bottom line. Cloud engagements are so dynamic, they require new, updated compliance programs just to keep up with the commonplace changes in their environments. You can't check every six months and hope for the best. You need to check continuously that the programs in place are robust and happening continuously. Therefore, you need a continuous monitoring and remediation program to ensure that those services running in the cloud are compliant.

Cloud engagements are so dynamic, they require new and continuous compliance programs just to keep up with the commonplace changes in their environments.



Continuous Compliance delivers a holistic, program-based approach to both technical and process-oriented compliance.

The Impact of Cloud on Highly Regulated Industries

Compliance challenges, of course, vary by industry. Moving to the cloud exacerbates the impacts of already complex, interrelated regulations and oversight in highly regulated industries such as financial services and healthcare.

In any industry, the penalties for noncompliance are stiff. Companies face potential fines, loss of business, loss of clients, firings, suspensions – even potential jail terms in certain circumstances. In retail, for example, companies are grappling with the effects of the new PCI regulation requiring a business to protect credit card data and customers' identities. Companies that don't comply may have to pay more for credit card transactions – or lose the ability to use credit cards at all. Noncompliance is clearly not worth the risk.

Geography Plays a Role

It would be one thing if companies were able to rework compliance processes globally, just based on the changes imposed by the cloud. But compliance rules in one locale don't always mesh with those in another. Take GDPR and FCA, for instance. These are a pair of new regulations created in the UK that require businesses to protect the pri-

vacy of individual data. They were created to govern individual data in Europe, but they apply to every global business that touches European consumers.

These are just the latest examples of geography-specific regulations that tilt the playing field for companies preparing compliance plans. As cloud adoption increases, expect to see more government actions to ensure that data is accounted for and protected.

Compliance Needs to be Monitored and Updated

There's a misconception that monitoring for security and for compliance amount to the same thing. Security is a big part of compliance, for sure, and having tools that produce reports about threat detection and security preparedness are critical to the survival of any business.

But there's more to compliance monitoring than keeping track of security threats. Regular monitoring provides continuous updates and assessments of issues – in the cloud and beyond – that are evolving more quickly and unexpectedly than ever before. It provides the domain-specific data that companies need to successfully manage their compliance programs.

A number of banks are embracing this continuous, real-time monitoring trend. According to a May 2017 report by McKinsey & Co., as the scope of regulation widens some financial institutions have “chosen to be ‘constantly materially compliant,’ a status just shy of full compliance, because of ongoing long-term remediation programs.”

The ability to review stats daily instead of, for example, four weeks before an audit, allows teams to spend more time moving their businesses forward rather than reacting to urgent issues. And real-time analysis pays significant dividends when it leads to early detection of trends. You gain the ability to take proactive steps to remediate and prevent minor issues from becoming critical issues.

Ultimately, real-time analysis of compliance readiness is the catalyst for creating a data-driven, fact-based approach to an organization’s compliance in the cloud program.

Continuous Compliance

CTP’s Continuous Compliance delivers a holistic, program-based approach to both technical and process-oriented compliance. With continuous assessments of cloud environments run against key regulatory frameworks like PCI, NIST and others, Continuous Compliance delivers real-time data to your business and to CTP to drive remediation programs forward.

As regulations and standards evolve, CTP identifies those changes rapidly and adapts client policies to remain in compliance. Clients benefit from reduced risk of gaps in compliance, less time and fewer resources required to constantly research and implement controls, and faster and less labor-intensive audit preparation.

Instead of continually addressing urgent issues, your development teams have more time to focus on the work at hand. With the right tools (and the right insights), they can be more productive and execute more compliant software builds. You experience fewer drills, tighter operational secu-

rity and better visibility into risks before they become critical issues.

Continuous Compliance also enables more focused and informed program-level oversight and governance to help you successfully steer your business forward.

Gaining Control

How do you get your cloud compliance under control? Here is a list of the priorities everyone should consider:

- Continuously assess and monitor activities to identify risks and potential sources of compliance exposure
- Have a well understood process for remediation of control failures and identified risks
- Take proactive steps to review cloud application architectures and corresponding controls to ensure compliance readiness
- Ensure that those responsible and accountable for compliance and remediation within the organization have access to real-time data about control failures
- Regularly update your implementation of regulatory or IT control frameworks; the rules can and do change
- Ensure compliance readiness is a key priority of the CIO, CISO and business unit leaders, in addition to audit staff

These are the key ingredients to a solution that puts you in control of compliance in the cloud.

Written by Bob Krygowski, former Director of Product Management, CTP, and currently Product Lead, Managed Cloud ,at 2nd Watch

**Learn more about Continuous
Compliance solutions at
cloudtp.com/continuous-compliance.**

HPE'S EVOLUTION TO A HYBRID CLOUD ENVIRONMENT

Hewlett Packard Enterprise's former VP of Information Technology shares his lessons learned after migrating 90% of the organization's applications to the cloud.

By conventional measures, the cloud-first transformation Hewlett Packard Enterprise (HPE) executed would have to be labeled a success. After all, we shrunk our data center footprint from nearly 500 to just four, took \$2 billion in costs out of the business, and migrated 90 percent of our apps to the cloud.

But, as Arianna Huffington once said, the path to success is “not a straight line.” She described it as “much more of a dance, and being open to possibilities,” and we can relate to that. We had to do our share of dancing to move past obstacles we inadvertently put in our own path. We of course were open to all possibilities as we moved through the process. But we didn’t take advantage of all the best practices at our disposal – such as the ones HPE and Cloud Technology Partners (CTP) deploy today: creating a Cloud Business Office (CBO), or following a plan that connects the impacts of specific events in a more holistic way.

Here’s a rundown of how we executed our cloud-first transformation, and the valuable lessons we learned along the way.

Embarking on a Journey

In the beginning, we focused our plan on cutting costs and improving agility. We had too many data centers – 85 large properties and more than 400 small sites (under 10,000 square feet) in 29 countries. We also had about 7,000 applications in our portfolio, including more than 100 instances of SAP.

This was the result of 35 years of inadequately integrating acquisitions. When we acquired a company, at least one data center came along with the deal. To keep peace with the acquired company, HPE management told IT to leave them alone for three years, and then go in and integrate the systems. Three years would go by, and the will and the funding would not be there, so many of the integrations just never got done.

We arrived at a point where we were spending just north of \$4 billion a year on IT – about 4 percent of the old HP’s \$115 billion in annual revenues. For years, that seemed like a reasonable percentage. Companies we saw as our peers – Google, Twitter, Facebook (which we could see from our offices in Palo Alto) – hovered around the 4 percent mark. But they were different. When Google spends 4 percent on IT, it is using its IT budget to build products. We were an IT company, but we were spending 4 percent on IT just to manage our business. We should have been comparing ourselves to engineering and manufacturing companies like IBM, Cisco and GE, which spend about 2 percent of revenue on IT. So we had to put together a program to get down to 2 percent.

We did just that, going through a grueling exercise of consolidation, virtualization and automation. We shrank from 465 data centers to six, saving about \$200 million. Fewer

data centers requires less networking, so we shaved \$600 million off our carrier fees. Eliminating 100 instances of SAP saved about \$100 million in licenses. A remaining \$1 billion-plus of savings came from rationalizing our application portfolio from 7,000 to 1,800, and reducing IT headcount from 20,000 to 2,000.

Just Getting Started

At this point, we were pretty happy with our progress. We were running the company's IT with six shiny new data centers, and costs seemed like they were at a manageable level. We didn't realize it at the time, but we were only getting started.

Within 18 months of building the new data centers, we began running out of capacity. We decided to plug the gap with EcoPODs, modular, containerized data centers we manufacture, which each generate a megawatt of power. We set up an EcoPOD next to each data center and planned to add two a year for the next five years. They were an expense, but far less costly than new data centers.

Our team was happy with our strategy, but our leader was not. Meg Whitman, HPE's CEO at the time, questioned us about our plan and its metrics, such as CPU utilization. At the time our utilization was about 10 percent, just below the industry average. She said she would like to see it in the 80 percent range.

We had work to do. When we looked at our environment, we saw that we had about 10,000 virtual machines (VMs) that essentially were not being used. The reason? Developers were hoarding them. On average, it took 21 days to get a VM approved. Developers didn't want to wait that long. They wanted to have capacity on hand to start developing immediately, as they can today on a PaaS. So, they ordered extras – dozens of them.

This got us thinking in terms of a larger transformation.

Taking the Next Step

The first thing we did was set up a cloud-like system we referred to as highly automated platform provisioning. It was not really a cloud. There were no APIs, just automation.

Developers could go to a portal and order up cores, storage, memory, an operating system, middleware databases and load balancing. Twenty minutes later, they would have an environment.

This helped us to do a better job managing our IT environment. We identified VMs that were overprovisioned, used automation tools and drove our utilization up to 30 percent. We were able to eliminate the use of EcoPODs and shrink the number of data centers down to four.


The next step was to move to the cloud. We started by creating an OpenStack cloud for cloud native development projects and then started brokering workloads to Azure. The positive response was immediate. People were tired of the old way of relying on on-premises resources, so we put together a project to move the majority of our workloads to the cloud.

Our plan called for the dissemination of workloads into four main buckets. The first would house about 10 percent of our applications – traditional IT resources, such as SAP HANA appliances and IBM mainframes, which would have to remain on premises. The remainder would go to the OpenStack VM (10 percent), to the public cloud (60 percent) and to SaaS applications (20 percent).

In the end, we moved far more workloads to OpenStack (about 50 percent) and far fewer to the public cloud (10 percent). The problem was we didn't have a good plan in place to manage costs throughout the process. Public cloud costs were swelling, and we weren't shutting off VMs quickly enough to harvest the savings so we could move fast on public cloud deployments. We got scared and scaled back our cloud efforts.

Understanding the “Why”

This is something CTP's business model could have helped with. HPE had set a goal to move 60 percent of our workloads to the public cloud, but we did not consider the many factors involved in making migration decisions. CTP helps customers understand the “why.” Our problem was, we had no clear idea of why we should move workloads into certain buckets at certain times. We just wanted to do it.

An aerial photograph of a winding asphalt road on a lush green hillside. The road curves through the landscape, with several cars visible at different points along its path. The surrounding terrain is covered in dense green vegetation, and there are patches of snow or light-colored rocks scattered across the hills. The overall scene is bright and clear, suggesting a sunny day.

Our move to Azure had some hiccups. We had very little understanding of the time it would take to do the migrations. Our culture didn't put a premium on scheduling – doing little things like spinning instances down on the weekends when nobody in the organization would need the resources. That cost us a lot of money because the controls weren't there.

We also did not instill a direct relationship between our on-ramp to the cloud and our using on-premises gear. These were seen as two separate processes that were not connected. Control and transparency between environments is essential when you are making the transition between VMs and the public cloud.

If we were to do it all over again, having a Cloud Business Office to direct traffic and manage those relationships would be critical to the success of the project. A CBO gives structure to the business owners and makes sure all functions are handled – everything from security to procurement, to finance, to executive reporting. The CBO holds people accountable and ensures the project is done right.

Conclusion

These days enterprises face stiff competition from cloud-borne companies that are not saddled with legacy costs and processes. Cloud-first transformations can help enterprises close competitive gaps and get ahead. But they can be a challenge to pull off, given the number of obstacles companies can encounter along the way. At HPE we learned a lot of lessons on our sometimes circuitous journey to the cloud. With a little more planning and a little more understanding of the migration process, other companies can follow more of a straight line to get to their own cloud destinations.

Written by Scott Anderson, Vice President of Global Cloud Advisory, HPE.



Accelerate your cloud strategy with HPE
Learn more @ hpe.com/info/hybridcloud



Get the latest news and best practices delivered to your inbox.

Join thousands of IT professionals who get The Doppler Report every week.

hpe.com/info/doppler

Check out these articles in the Fear No Cloud issue

READ

Cloud Kryptonite: 7 Disruptive Forces Draining Your Cloud Program

Early detection of cloud blockers is critical. Pull back the covers and search for these seven signs of cloud kryptonite before they weaken your program.

READ

When Technology Isn't Enough: Cloud Platform Services Don't Run Themselves

When companies are so focused on the very real challenge of building out the technology, they forget that the internal customers for whom these platforms are built do not automatically understand the platform and what it can do for them.

READ

Engaging Talent: The Cloud Transition Model for Your People

Your move to cloud has a cutover date. Your people do not. Learn about the steps managers should take to reduce uncertainty throughout their cloud transformation.

READ

Not All Cloud Superheroes Wear Capes: The Integral Roles Behind Your Cloud Transformation

Achieving quality outcomes does not have to involve all-nighters, last-minute rescue missions and other similar heroics. Learn who is keeping your program on track and on time.

READ

Cultural Transformation in the Cloud: Tales of Missteps and Success

Cloud initiatives are hard to pull off, especially when an organization's culture is affected. But with a little ingenuity, the right people and proven methodology, they can succeed.

CLOUD WITH **CONFIDENCE**[®]



Hewlett Packard
Enterprise

THE DOPPLER is published quarterly by Cloud Technology Partners, a Hewlett Packard Enterprise company.

Join thousands of IT professionals who get our weekly email reports by visiting

hpe.com/info/doppler