Network Operations and Security Professionals' Guide to Managing Public Cloud Journeys

Version 1.0
Released: March 1, 2020

Securosis, L.L.C. 515 E. Carefree Highway Suite #766 Phoenix, AZ 85085 **T** 602-412-3051 info@securosis.com www.securosis.com

Author's Note

The content in this report was developed independently of any sponsors. It is based on material originally posted on the Securosis blog, but has been enhanced, reviewed, and professionally edited.

Special thanks to Chris Pepper for editing and content support.

Licensed by Gigamon

Gigamon® Gigamon is the first company to deliver complete network visibility and analytics on all information-in-motion, from raw packets to apps, across physical, virtual and cloud infrastructures. We aggregate, transform and analyze network traffic to solve for critical performance and security needs, including rapid threat detection and response, freeing your organization to drive digital innovation. In short, we enable you to run fast, stay secure and innovate. Gigamon has been awarded over 75 technology patents and enjoys industry-leading customer satisfaction with more than 3,000 global organizations, including 80 percent of the Fortune 100. For the full story on how Gigamon can help you, please visit www.gigamon.com.

Copyright

Table of Contents

Your Cloud Journeys is Unique, but Not Unknown

Cloud computing is different, disruptive, and transformative. It has no patience for traditional practices or existing architectures. The cloud requires change, and there is a growing body of documentation on end states you should strive for, but a lack of guidance on how to *get there*. Cloud computing may be a journey, but it's one with many paths to what is all too often a highly nebulous destination.

Although every individual enterprise has different goals, needs, and capabilities for their cloud transition, our experience and research have identified a series of fairly consistent patterns. You can think of moving to cloud as climbing a mountain with a single peak, with everyone starting from the same trailhead. But this simplistic view, which all too often underlies conference presentations and tech articles, fails to capture the unique opportunities and challenges facing you. At the other extreme, we can think of the journey as involving a mountain range with innumerable peaks, starting points, and paths... and a distinct lack of accurate maps. This is the view which tends to produce hands thrown up in the air, expressions of impossibility, and analysis paralysis.

But our research and experience guide us between those extremes. Instead of a single constrained path which doesn't reflect individual needs, or totally individualized paths which require you to build everything and relearn every lesson from scratch, we see a smaller set of common options with consistent characteristics and experiences. Think of it as starting from a few trailheads, landing on a few peaks, and only dealing with a handful of well-marked trails. These won't cover every option, but can be a surprisingly useful way to help structure your journey, move up the hill more gracefully, and avoid falling off some **really** sharp cliff edges.

Introducing Cloud Adoption Patterns

Cloud adoption patterns represent a consolidated set of cloud adoption journeys, compiled through discussions with hundreds of enterprises and dozens of hands-on projects. Less concrete than specific cloud controls, they are a more general way of predicting and understanding the problems facing organizations when moving to cloud, based on starting point and destination. These patterns have different implications across functional teams, and are especially useful for network operations and network security, because they tend to fairly accurately predict many architectural implications, which then map directly to management processes.

For example there are huge differences between a brand-new startup or cloud project without any existing resources, a major data center migration, and a smaller migration of key applications. Even a

straight-up lift and shift migration is extremely different if it's a one-off vs. a smaller project vs. wrapped up in a massive data center move with a hard cutoff deadline (often thanks to a hosting contract which is not being renewed). Each case migrates an existing application stack into the cloud, but the different scope and time constraints dramatically affect the actual migration process.

We'll cover them in detail but the four patterns we have identified are:

- **Developer led:** A development team starts building something in the cloud outside normal processes, and then pulls the rest of the organization behind them.
- **Data center transformation:** An operations-led process defined by an organization planning a methodical migration out of existing data centers and into the cloud, sometimes over a decade or more.
- **Snap migration:** An enterprise is forced out of some or all their data centers on a short timeline, due to contract renewals or other business drivers.
- **Native new build:** The organization plans to build a new application or several completely in the cloud using native technologies.

You likely noticed we didn't mention some common terms like "refactor" and "new to cloud". Those are important concepts but we consider them options on the journey, not defining characteristics. Our four patterns are about the *drivers for your cloud migration and your desired end state.*

Using Cloud Adoption Patterns

The adoption patterns offer a framework for thinking about your upcoming (or in-process) journey, and help identify both strategies for success and potential failure points. These aren't proscriptive like the Cloud Security Maturity Model or the Cloud Controls Matrix — they won't tell you exactly which controls to implement, but are more helpful when choosing a path, defining priorities, mapping architectures, and adjusting processes.

Going back to our mountain-climbing analogy, the cloud adoption patterns point you down the right path and help you decide which gear to take, but it's still up to you to load your pack, know how to use the gear, plan your stops, and apply sunscreen. These patterns represent a set of characteristics we consistently see based on how organizations move to cloud. Any individual organization might experience multiple patterns across different projects. For example a single project might behave more like a startup, even while you concurrently run a larger data center migration.

Next we will detail the patterns with their defining characteristics. You can use them to determine your overall organizational journey, as well as to plan out individual projects with their own characteristics. To help you better internalize these patterns we will offer fictional examples based on real experiences and projects. Once you know which path you are on, our final sections will include top-line recommendations for network operations and security, and tie back to our examples to show how they play out in real life. We will also highlight the common pitfalls and their potential consequences.

This research should help you better understand what approaches will work best for *your project in your organization*. We are focusing this first round on networking, but future work will build on this basis to cover additional operational areas.

Cloud migrations can be tough and confusing. Nothing lets you skip over the hard work, but learning the lessons of those who have already climbed the mountain can save costs, reduce frustration, and increase your odds for success.

Defining the Journey—the Four Cloud Adoption Patterns

Understanding Cloud Adoption Patterns

Cloud adoption patterns represent the most common ways organizations move from traditional operations into cloud computing. They contain hard lessons learned by those who went before. While every journey is distinct, hands-on projects and research have shown us a broad range of consistent experiences, which organizations can use to better manage their own projects. The patterns won't tell you exactly which architectures and controls to put in place, but they can serve as a great resource to point you in the right general direction and help guide decisions.

Another way to think of cloud adoption patterns is as embodying the aggregate experiences of hundreds of organizations. To go back to our analogy of hiking up a mountain, it never hurts to ask people who have already finished the trip what to look out for.

Characteristics of Cloud Adoption Patterns

We will get into more descriptive detail as we walk through each pattern, but we find this grid useful to define the key characteristics.

Characteristics	Developer Led	Data Center Transformation	Snap Migration	Native New Build
Size	Medium/Large	Large	Medium/Large	All (project-only for mid-large)
Vertical	All (except financial and government)	All	Including financial and government	Variable/All
Speed	Fast then slow	Slow (2-3 years or more)	18-24 months	Fast as DevOps
Risk	High	Low(er)	High	Variable
Security	Late	Early	Trailing	Mid to late
Network Ops	Late	Early	Early to mid	Late (developers manage)
Tooling	New, and old when forced	Culturally influenced; old & new	Panic (a lot of old)	New, unless culturally forced to old
Budget Owner	Project based/no one	IT/Ops/Sec	IT or poorly defined	Project-based; some security for shared services

- **Size:** The size of organizations likely to follow this pattern. For example developer-led projects are rarely seen in small startups because they can skip directly to native new builds, but common in large companies.
- **Vertical:** We see these patterns across all verticals, but in highly-regulated ones like financial services and government, certain patterns are less common due to tighter internal controls and compliance requirements.
- **Speed:** The overall velocity of the project, which often varies during the project lifetime. We'll jump into this more, but an example is developer-led, where initial setup and deployment are very fast, but then wrangling in central security and operational control can take years.
- **Risk:** This is a combination of danger to the organization and likelihood of project failure. For example in a snap migration everything tends to move faster than security and operations can keep up, which creates a high chance of configuration error.
- **Security:** When security is engaged and starts influencing the project.
- **Network Ops:** When network operations becomes engaged and starts influencing the project. Security folks are used to being late to the party because developers can build their own networks with a few API calls, but this is often a new and unpleasant experience for networking professionals.
- **Tooling:** The kind of tooling used to support the project. "New" means new, cloud-native tools. "Old" means the tools you already run in your data centers.
- **Budget Owner:** Someone has to pay at some point. This is important because it represents potential impact on your budget, and tends to indicate who has the most control over the project.

The Four Cloud Adoption Patterns

It's time to describe what the patterns look like and identify key risks. Our next section will offer some top-line recommendations to improve your chances of success.

One last point before we jump into the patterns themselves: while they focus on the overall experiences of an organization, patterns also apply at the project level, and an organization may experience multiple patterns at the same time. For example it isn't unusual for a company with a "new to cloud" policy to also migrate existing resources over time as a long-term project. They experience both the *data center transformation* and *native new build* patterns simultaneously.

Developer Led

Mark was eating a mediocre lunch at his desk when a new "priority" ticket dropped into his network ops queue. "Huh, we haven't heard from that team in a while... weird." He set the microwaved leftovers to the side and clicked open the request…

Request for firewall rule change: Allow port 3306 from IP 52.11.33.xxx/32. Mission critical timeline.

*"What the ?!? That's not one of our IPs!" Mark thought as he ran a lookup. "amazonaws.com? You have **got** to be kidding me? We shouldn't have anything up there." Mark fired off emails to his manager and the person who sent the ticket, but he had a bad feeling he was about to get dragged into the kind of mess that would seriously ruin his plans for the next few months.*

Developer-led projects are when a developer or team builds something in a cloud on their own, and central IT is then forced to support it. We sometimes call this "developer tethering", because these often unsanctioned and/or uncoordinated projects anchor an organization to a cloud provider, dragging the rest of the organization in after them. These projects aren't always against policy — this pattern is also common in mergers and acquisitions. This also isn't necessarily a first step into the cloud overall — it can also be a project which pulls an enterprise into a *new cloud provider*, rather than their *existing, preferred cloud provider*.

This creates a series of tough issues. To meet the definition of this pattern we assume you can't just shut the project down, but actually need to support it. The project has been developed and deployed without the input of security or networking, and may access production data.

- **Size:** We mostly see this pattern in medium and large organizations. In smaller enterprises the overall scope of what's going on is easier to understand, whereas larger organizations tend to have an increasing number of teams operating at least semi-independently. Larger organizations

are also more likely to engage in M&A activity which forces them to support new providers. In fact most multi-cloud deployments we run into result directly from acquiring a company using something like Azure when everything else is on AWS.

- **Vertical:** This pattern is everywhere, but less common in highly regulated and tightly governed organizations — particularly financial services and government. Not every financial services organization is well-governed so don't assume you are immune, but controls tend to be tighter on more sensitive data, so when you do hit this pattern the risk might be lower. In government it's actually budgets more than regulations which limit this pattern — few government employees can get away with throwing AWS onto corporate cards.
- **Speed:** In the beginning, at least once security and networking find out about the project, there is a big rush to manage the largest risks and loop the project into some sort of central management. This flurry of activity then slows down into a longer, more methodical wrangling to bring everything up to standard. It starts with stopgaps, such as opening up firewalls to specific IP ranges or throwing in a VPC connection, followed by a longer process to rework both the deployment and internal support, such as by setting up direct connect.
- **Risk:** These are high-risk situations. Security was likely not involved, and we often find a high number of configuration errors when assessing these environments. They can often function as an isolated outpost for a while, but there are still risks of failed integration when the organization tries to pull them back into the fold — especially if they require complicated network connectivity.
- **Security:** Security is typically involved only late in development or after deployment, because the project team was off running on its own.
- **Network Ops:** As with security, networking enters late. If the project doesn't require connectivity back to an existing network they might not be involved at all.
- **Tooling:** Most often these projects leverage integrated tools provided by the cloud service provider. There is rarely budget for security or network specific tooling beyond that, since CSP tool costs are all hidden within basic cloud deployment costs. One problem we sometimes see is that after the project is discovered and there's the mad rush to bring it under central management, a bag of existing tools — which often fit the cloud platform poorly — are forced into place. This is most common with network and endpoint security tools which aren't cloud native — a virtual appliance isn't necessarily a good answer to a cloud problem.
- **Budget Owner:** The project team somehow managed to get budget to start the deployment, which they can use as a cudgel to limit external management. This may fall apart as the project grows and costs increase (as they always do) and the project has to steal budget from someplace else.

Key Risks

These should be obvious: you have an unsanctioned application stack running in an unapproved cloud, with which you may have little experience, uncoordinated with security or networking. However many project teams *try to do the right things*. You can't assume the project is an abject failure. Some of these projects are significantly better designed and managed, from the cloud standpoint, than lift and shifts or other cloud initiatives. It all depends on the team. Based on our experiences:

- Security configuration errors are highly likely.
- There may be unapproved and *ad hoc* network connections back to existing resources. At times these are unapproved VPN connections, SSH jump boxes, or similar.
- Deployment environments may be messy, full of cruft and design flaws.
- Development/testing and production are generally intermingled in the same account/subscription/project, which creates a larger blast radius for attacks.

Mark glanced up at the wall of sticky notes layered on top of the whiteboard's innumerable chicken-scratch architectural diagrams. A year of planning and setup, and they were finally about to move the first production application.

"Okay," started Sarah, "we've tested the latency on the direct connect but we're still having problems updating the IPs for the firewall rules in the Dallas DC. The application team says they need more flexibility since they want to deploy with infrastructure as code and use auto scale groups in different availability zones. They claim that trying to manage everything with such restricted IPs doesn't work well in their VNets. Something about web servers and app servers reusing each other's IPs as they scale up and down and the firewall team is too slow to update the rules."

Mark interjected, "What if we drop the firewalls on the direct connects? Then they can use what they want within their CIDR blocks?"

"Isn't that a security risk?"

"I don't think so," replied Mark, "after going through the cloud training last month I'm starting to believe we've been thinking about this all wrong. The cloud isn't the untrusted network — we're just as likely to get breached from the data center or someone compromising an admin's laptop on the corporate network."

Data Center Transformation

Data center transformations are long-term projects, where the migration is methodical and centrally planned. That isn't always beneficial — these projects are often hindered by overanalysis and huge requirements documents, which can result in high costs and slow timelines. They also tend to create their own particular set of design flaws. In particular, there is often a focus on building a perfect landing zone or "minimum viable cloud" which replicates the existing data center, rather than taking advantage of native capabilities of the cloud platform. Existing tooling and knowledge are thrown at the problem, rather than trying to do things the "cloud way".

Not to spoil our recommendations, but treating the migration as a series of incremental projects rather than a monolithic deployment will dramatically improve your chance of success. Culture, silos, politics, and experience all significantly impact how well these projects go.

- **Size:** You need a data center to transform, so this pattern shows up at very large, large, and sometimes mid-sized enterprises.
- **Vertical:** This pattern is common across all verticals which meet the size requirements. Five years ago we weren't seeing it with regulated industries, but cloud computing has long since passed that limitation.
- **Speed:** These projects tend to move at a snail's pace. There are a lot of planning cycles, and building baseline cloud infrastructure, before any production workloads are moved. In some cases we see progressive organizations breaking things into smaller projects rather than shoehorning everything into one (or a small number of) cloud environments, but this is uncommon. Multi-year projects are the norm, although more agile approaches are possible.
- **Risk:** The risk of a security failure is lower due to the slower pace and tighter controls, but there can be high risk of project failure, depending on approach. Large monolithic cloud environments are highly prone to failure within 18-24 months. Compartmentalized deployments (using multiple accounts, subscriptions, and projects) have a lower chances of major failure.
- **Security:** Security is engaged early. The risk is that the security team isn't familiar or experienced with cloud, and may attempt to push traditional techniques and tools which don't work well in cloud.
- **Network Ops:** Like security, networking is involved early. And as with security, the risk is lacking cloud domain knowledge for an effective and appropriate design.
- **Tooling:** Tooling depends on culture, silos, and politics. There is excellent opportunity to use cloud-native tooling, including existing tools with cloud-native capabilities. But we also see, as in our opening story, frequent reliance on existing tools and techniques which aren't well suited to cloud and end up causing problems.

- **Budget Owner:** These projects tend to have a central budget, so shared service teams such as operations, networking, and security may be able to draw on this budget or submit their own requests for additional project funding.

Key Risks

There are two major categories of risks, depending on the overall transformation approach.

- Large, monolithic projects where you set everything up in a small number of cloud environments and try to make them look like your existing data center. These 'monocloud' deployments are slow and prone to breaking horribly in 18-24 months. IAM boundaries and service limits are two of the largest obstacles. Agility is also often reduced, which even pushes some teams to avoid the cloud. Costs are also typically higher. Organizations tend to find themselves on this path if they don't have enough internal cloud knowledge and experience. Both cloud providers and independent consultants usually just nod their heads and say 'yes' when you approach them with a monocloud proposal because they want your business — even when they understand the obstacles you will eventually hit.
- Discreet, project-based deployment transformations leverage some shared services, but application stacks and business units have their own cloud accounts/subscriptions/projects (under the central organization/tenant). This cloud-native approach avoids many problems of monolithic deployment, but brings its own costs and complexities. Managing large numbers of cloud environments (hundreds are typical, and thousands are very real) requires deep expertise and proper tooling. The flexible nature of software defined networks in the cloud is a complex problem, especially when different projects need to talk to each other and back to non-cloud resources, which many enterprises never move to the cloud.

Snap Migration

Sitting in the back of the conference room, Bill whispered to Clarice, "No way. No forking way. This is NOT going to end well."

"Do they have any friggin' idea how bad this is?" she replied.

"How do they possibly expect us to move 3 entire data centers in 18 months up to Amazon... I can't even get a VM for testing in less than 3 months!"

"I get they want out of our crappy contract before the renewal, but this is insane."

"Heh," huffed Bill, "Maybe they'll finally approve those cloud training classes we've been asking for."

"Yeah right," she replied sarcastically, "like they'll train us instead of throwing cash at some consultant."

Snap migrations are the worst of all worlds. Massive projects driven by hard deadlines, they are nearly always doomed to some level of failure. In our experiences the decision-makers behind these projects rarely understand the complexity of migrating to cloud, and are overly influenced by executive sales teams and consultants. Not they are always doomed to complete failure, but the margins are thin and you will be navigating a tightrope of risks.

There is a subset of this pattern for more limited projects which don't encompass absolutely everything. For example imagine the same contract renewal drive, but for a subsidiary or acquisition rather than the entire organization. The smaller the scale the lower the risk.

- **Size:** Mid to large. You need to be big enough to have data centers, but not so big that you own the real estate they sit on. A defining characteristic is that these projects are often driven by contract renewals on hosted or managed data centers. That's why there's a hard deadline... management wants out, as much as possible, before they get locked into the next 7-year renewal.
- **Vertical:** Organizations across all verticals find themselves in hosting contracts they want out of. We even know of projects in highly regulated financial services which you'd think would never accept this level of risk. Government is the least likely, and tends to be driven more by whichever political appointee decides they want to shake things up.
- **Speed:** 18-24 months for the first phase. We rarely see less than 12 months. Sometimes there will be a shorter initial push to get out of at least some data centers, as the contract moves into a month-by-month extension phase.
- **Risk:** As high as it gets in every possible way. Organizations falling into this pattern might have some internal cloud experience, but as a rule not enough people with enough depth to support the needed scale. There is heavy reliance on outside help, but few consulting firms (or cloud providers themselves) have a deep bench of solid experts who can avoid all the common pitfalls.
- **Security:** Security is somewhat engaged but can't do anything to slow things down. They are also typically tasked with building out their own shared services, so likely aren't staffed to evaluate individual projects. They tend to trail behind deployments, trying to assess and clean things up after the fact. They often get to set a few hardline policies up front (typically relating to Internet accessibility), but until they stand up their own monitoring and enforcement capabilities, things slip through the cracks.
- **Network Ops:** There is a bit more variability here, depending on the deployment style. If there is a monocloud (or small number of environments), networking is typically engaged early and plays a very strong role in getting things set up. They are tasked with configuring the fatter pipes needed for such large migrations. The risk is that they often lack cloud experience, and introduce designs which work well in a data center but fit cloud deployments poorly.
- **Tooling:** Panic is the name of the game. The initial focus is on the tools at hand, and vendors already in place, combined with cloud-native tools. We hate to say it, but this can be deeply

influenced not only by culture but by which consultants are already in the door. Eventually the project starts introducing more cloud-native tooling to solve specific problems. For example in projects we've seen, visibility (cloud assessment and mapping) tools tend to be early buys.

- **Budget Owner:** This can be poorly defined, but they often pull from a central IT budget or specially designated project budget. Whoever controls the money has the most influence. The chances of success go up when all teams are properly funded and staffed. Also, water is wet.

Key Risks

They abound but we can categorize them based on project characteristics.

- As any IT pro knows, every project of this scale runs over time and budget.
- There is often a reliance on outside contractors who push things along quickly, but don't know (or care) enough to have a sense of the enterprise risk. Their job is to get things moved — not necessarily to do so the safest way. This can lead to exposure as they accept risks a company employee might avoid.
- Security often lacks general cloud security knowledge, as well as provider and platform experience. They can build this but it takes time, and in the process the organization is likely to accumulate technical security debt. For example two of the most common flaws we find on assessments are overly-privileged IAM and poorly segregated networks.
- Rapidly designing a cloud network at scale is difficult and complex, especially for a team which is still keeping the existing environment running, and (like security) probably lacks deep cloud experience. We often see one or two members of a team tasked as cloud experts but this really isn't enough. Given the time constraints the network often ends up poorly compartmentalized, and projects tend to be shoveled into shared VPCs/VNets in ways which later run up against service limits, performance problems, and other constraints.

Native New Build

John was actually excited when he walked into the meeting room. It had been a long time since he had the chance to stretch his security creativity muscles. He nodded to Maria from networking as he pulled out an open Aeron knockoff chair and dropped his new matte-black Surface on the conference table. He still wasn't sure what stickers to throw on it, but after a little burn-in period during the Azure training he and Maria had just finished, he was getting used to working off an underpowered device. He still had his old desktop for handling all the data center thick clients, but for this Azure project all he needed was a web browser and some PowerShell. Although he kind of envied the consultants with their brand new MacBooks.

"Hey everyone," Wendy started, "we have a tight timeline but we finally have approval for the new Azure subscription. Maria, can you get the network connected up?"

"Actually, I don't think we need to. John signed off on using JIT connections and client VPNs instead of requiring a dedicated backhaul. We went through the architecture and there aren't any dependencies on internal resources. We know we'll need more of a hybrid design for the CRM project but we are free and clear for this one."

Native new build projects are true cloud-native deployments. That doesn't mean they are all brand new projects — this pattern also includes refactoring and rearchitecting, so long as the eventual design is cloud native. These may also include hybrid deployments — the new build may still need connections back to the premises.

- **Size:** All sizes. In a large enterprise this will likely be a designated project or series of projects (especially in a "new to cloud" organization). In a small startup the entire company could be a new build.
- **Vertical:** All verticals. Even government and highly regulated industries. We have worked on these projects with financials, state governments, and even public utilities.
- **Speed:** As "fast as DevOps". We don't mean that facetiously — some teams are faster and some slower, but we nearly always see DevOps techniques used, and they often define the overall project pace. *These are developer-driven projects*.

- **Risk:** We will talk more about risk in a moment, but here we'll just note that risk is highly variable, dependent on the skills and training of the project team.
- **Security:** Unlike our previous example security may be late to the project. There is usually an inflection point, when the project is getting close to production, at which security gets pulled in. Before that the developers themselves manage most security. This improves over time — the organization is more likely to struggle in this area on early projects, but start integrating security earlier over time, as more and more moves to cloud and skills and staffing improve.
- **Network Ops:** Networking is more likely to be engaged early if there are hybrid connectivity requirements, or might not be involved at all depending on the overall architecture. These days we see a growing number of serverless (or mostly serverless) deployments where there isn't even a project network, and all the components talk to each other within the "metastructure" of the cloud management plane.
- **Tooling:** Typically newer, cloud native, and often CSP provided. Quite a few of these projects start in their own cloud silos and use the provider's tooling, but as more of these projects deploy there is increased demand for central management and shared services (such as security assessment) to be added. We sometimes see development teams forced to use traditional on-premise tools, but this tends to be cultural — it isn't usually the best solution to the problems at hand.
- **Budget Owner:** Project based. Once you do enough of these there will also be shared services budgets for teams like security and networking.

Key Risks

Despite our optimistic opening, these projects bring their own risks. The project may be well-segregated in its deployment environment (an Azure subscription in our example) but that doesn't mean developers won't be over-provisioned. Or that a PaaS endpoint in the VNet won't ignore the expected Network Security Group rules (yes, that happens).

- This pattern can carry all the risks of the developer-led pattern if it is poorly governed. We have seen large organizations running dozens or hundreds of these projects, all poorly governed, each carrying tons of risks. If you read about a big cloud breach at a client who was proudly on stage at their cloud provider's conferences, odds are they are poorly governed internally.
- Cloud native services have different risks which take time to understand and learn to manage. In the data center migration pattern there is less reliance on the latest and greatest "serverless this, AI that", so traditional security and management techniques can be more effective. With native new builds you may be using services the cloud provider itself barely understands.
- Friction between security and the project team can badly impact the final product. Overly proscriptive security pushes teams into workarounds. Early partnering, ideally during initial development of the architecture, with security pros trained on the cloud platform, reduces risk.

- Managing a lot of these projects at scale **is really really hard**. Setting up effective shared services and security and network support (especially when hybrid or peered networks are involved) take deep expertise. Cloud providers are often terrible at helping you plan for this — they just want you to move as many workloads to them as quickly as possible.

Mastering the Journey—Building Network Manageability and Security for your Path

Recommendations for a Safe and Smooth Journey

Learning cloud adoption patterns doesn't just help us identify key problems and risks — we can use them to guide operational decisions to address the issues they consistently raise. This research focuses on managing networks and network security, but the patterns include broad security and operational implications which cover all facets of your cloud journey. Governance issues aside, we find that networking is typically one of the first areas of focus for organizations, so it's a good target for our first focused research. (For the curious, IAM and compliance are two other top areas organizations focus on, and struggle with, early in the process).

Developer Led

Mark sighed with relief and satisfaction as he confirmed the VPN certs were propagated and approved the firewall rule change ticket. The security group was already in good shape, and they managed to avoid having to add any kind of direct connect to the AWS account for the formerly-rogue project.

He pulled up their new cloud assessment dashboard and all the critical issues were clear. It would still take the IAM team and the project's developers a few months to scale down unneeded privileges but… not his problem. The federated identity portal was already hooked up and he would get real-time alerts on any security group changes.

"Now onto the next one," he mumbled after he glanced at his queue and lost his short-lived satisfaction.

"Hey, stop complaining!" remarked Sarah, "We should be clear after this backlog now that accounting is watching the credit cards for cloud charges; just run the assessment and see what we have before you start complaining."

Having your entire organization dragged into the cloud thanks to the efforts of a single team is disconcerting, but not unmanageable. The following steps will help you both wrangle errant projects under control, and also build a base for moving forward. This was the first adoption pattern we started to encounter a decade ago as cloud starting growing, so there are plenty of lessons to extract. Based on our experiences, a few principles really help to manage the situation.

- Remember that to fit this pattern you should be new to either the cloud in general, or to this cloud platform specifically. These are not recommendations for unsanctioned projects covered by your existing experience and footprint.
- Don't be antagonistic. Yes, the team probably knew better and shouldn't have done it... but your goal now is corrective action, not punitive.
- You need to reduce urgent risks while developing a plan to bring the errant project into the fold.
- Don't simply apply existing policies and tooling from other environments to this one. You need tooling and processes appropriate for this cloud provider.
- In our experience, despite the initial angst, these projects are excellent opportunities to learn your initial lessons on this platform, and to start building out for a larger supported program. If you keep one eye on immediate risks and the other on long-term benefits, everything should be fine.

The following recommendations go a long way toward reducing risks and increasing your chance of success. But before the bullet points we have one overarching recommendation: *As you gain control over the unapproved project, use it to learn the particulars of this cloud provider and build out your core cloud management capabilities.* When you assess, set yourself up to support your next ten assessments. When you enable monitoring and visibility, do so in a way which supports your next projects. Wherever possible build a core service rather than a one-off.

- Step one is to figure out what you are dealing with.
    - How many environments are involved? How many accounts, subscriptions, or projects?
    - How are the environments structured? This requires mapping out the application, the provider's PaaS services (such as load balancers and serverless capabilities), IAM, network(s), and data storage.
    - How are the services configured?
    - How are the networks structured and connected? The Software Defined Networks (SDN) used by all major cloud platforms only look the same on the surface — under the hood they are quite a bit different.
    - And, most importantly, *where does this project touch other enterprise resources and data?!?* This is essential for understanding exposure. Are there unknown VPN connections? Did someone peer through an existing dedicated network pipe? Is the project talking to an internal database over the Internet? We've seen all these and plenty more.
- Then prioritize your largest risks.
    - Internet exposures are common and one of the first things to lock down. We commonly see resources such as administrative servers and jump boxes exposed to the Internet. In nearly every single assessment we find at least one instance or container with port 22 exposed to the world. The quick fix for these is to lock them down to your known IP address ranges. Cloud providers' security groups are very effective because they just drop traffic which doesn't meet the rules, so they are an extremely effective security control and a better first step than trying to push everything through an on-premise firewall or virtual appliance.
    - Identity and Access Management is the next big piece to focus on. This research is focused more on networking, so we won't spend much time on these here. But when developers build out environments they almost always over-privilege access to themselves and application components. They also tend to use static credentials because unsanctioned projects are unlikely to integrate into federated identity management. Sweep out static credentials, enable federation, and turn on MFA everywhere you can.

    - Misconfigurations of cloud services are next. Public storage buckets, unsecured API gateways, and other services which are Internet exposed but *won't show up if you only look at the virtual networks*.

- After cleaning those up it's time to start layering in longer-term remediations and your gameplan. This is a huge topic so we will focus on network management and security.
    - During early discovery of developer-led projects, it is very common to want to tie the errant cloud account back into your on-premise network for connectivity and management. This instinct is usually wrong. Networking wasn't involved at the start, so it is unlikely there is an established network connection, and adding one won't necessarily provide any benefit. If the account is okay on its own, leave it. While outside the scope of this research, a wide range of techniques is available to provide necessary services to disconnected cloud accounts... or cloud-native connections such as service endpoints which achieve the same goals without the heavy lifting on fat pipes and CIDR segmenting.
        - We aren't suggesting you don't manage the network — we are saying you don't need to simply wire it up to your existing infrastructure to manage it or the resources it contains.
        - A big complication for integrating an unplanned SDN is the existing IP addressing (if there's even a virtual network — a real question thanks to new serverless architectures). This may be further motivation to keep it as a separate enclave.
    - We assume you followed our advice above and locked down the perimeter. Now it's important to fully map out all the internal connections, including connections between different virtual networks and accounts which are peered or otherwise connected using cloud-native techniques such as service endpoints.
        - One of the most common networking mistakes in this kind of projects is too-open internal networks. Clouds default to least privilege, but it is still all too easy to just open everything up to reduce friction during development. Use your map to start compartmentalizing internally. This may include network structure changes (routing and subnet modifications), which are fortunately easier to update using API calls and console clicks than stringing wires between routers.
        - Security groups should reference each other (in Azure you need Application Security Groups) instead of relying on IP addressing for internal cloud connections. This is fundamental to cloud networks, but not where people with traditional network security backgrounds tend to start.
    - Virtual security appliances (we are mostly talking about firewalls, IDS, and IPS) should only be used when security groups and native cloud capabilities cannot meet your needs. Virtual appliances are expensive to run because cloud providers charge for their compute cycles, and they create unnecessary chokepoints which affect performance and reliability.

The most common situations you still need them for are FQDN-based outbound filtering and specific blocklists which are hard or impossible to enforce with cloud-native security groups.

- Lastly, once everything is in a known good state, you should implement continuous configuration assessments and guardrails to keep things that way. For example in a production application you should generate an alert on any security group change, creation of new internet gateways, and other structural changes. All providers support monitoring these changes but you will likely need third-party tooling to pull the results together across providers and accounts.

Overall the key to handling this situation is to avoid panic, focus on obvious risks first, and then take your time to sweep through the rest in as cloud and provider specific a way as possible. Use it as a base to build your program, understanding that you will need to make short-term sacrifices to handle any significant exposures.

Data Center Transformation

Sarah snagged an extra chair outside Mark's cubicle as he shoved a pile of office detritus to the side to make space for her laptop.

"Okay, " she started, "I published the PrivateLink endpoint for the log receiver and set the internal domain name, but I need you to open up the security group and approve my PR on the CloudFormation templates to deploy all the service endpoints into the VPCs."

"No problem," he replied, "we got approval from the cloud team last week so we're good to go. Do we need to talk to the server image team to embed the DNS for the log agents?"

"They already have it and are publishing the new base AMIs that need it. We think most teams will just set their agents to save instance and container logs directly to S3, but some of the legacy stuff still needs to push them on the network. We are also letting teams use their own PrivateLink addresses if they want to swap out a local collector."

"Nice, " said Mark, "this will really help us drop some peering connections on the transit gateway. And I'm meeting with the database team next week to see if we can start moving them over."

Large multi-year data center moves are some of the most complex projects in information technology. Moving everything from one physical location to another is a massive undertaking. Doing so while keeping services up and running, without shutting the business down (either planned or unplanned), even more so. Swapping to an entirely different technology foundation at the same time? That can be the definition of insanity, yet every single organization of any size does it at some point.

The most common mistakes we see involve shoehorning traditional architectural and security concepts into the cloud — which can lead to extended timelines, increased costs, and long-term management issues. A few key principles can keep you moving in the right direction.

- If you are bad at network management and security in your existing data center, you will be surprised at how little changes in cloud. Look at cloud as an opportunity to do things better, ideally in a cloud-native way. Don't just bring across your existing practices without change — especially bad habits.
- Time is your friend. Don't rush, and don't let your cloud provider push you into moving faster than you are comfortable with. Their priorities are not yours.
- Don't assume your existing tools and processes will work well in cloud. Many organizations bring things across due to employee familiarity or because they already have licenses. Those aren't great reasons to deploy something in an entirely new operating environment.
    - That said, these days many products offer extensions for the cloud. You should still evaluate them instead of *assuming* they will meet your needs, but they might be a useful bridge.
- Learn first, move second. Take the time (if you have it) to hire and build the skills needed to operate on your new platform. You absolutely cannot expect your existing team to handle both the current environment and cloud if you don't give them the time to learn the skills and do the job.

In the *developer led* pattern we had to balance closing immediate risks against simultaneously building support for an entirely new operating environment and preparing for long-term support. Scary and difficult, but also usually self-constrained to something manageable like a single application stack. In a data center transformation the challenges are **scaling,** transitioning completely to a new environment, and any need to carry over legacy resources not designed to run in cloud.

- Start by building your plan.
    - You **will not** want to run everything in a single huge account/subscription/project on just 1-3 virtual networks. This is all too common and falls apart within 18-24 months due to service limits, differences in how cloud networks work, and cloud-native application requirements.

    - You will want multiple cloud environments (accounts/subscriptions/projects are the terms used by different providers) and very likely multiple virtual networks in each environment. These are needed for blast radius control, managing service limits, and limiting IAM scopes.
    - Map out your existing applications and environments (networks, cross-app connectivity, associated security controls, and related supporting services such as DNS and logging). Create a registry and then prioritize and sequence your moves.
    - Map out your application dependencies. You might have 50 applications which all connect back to a shared customer database. This directly impacts how you structure your accounts, virtual networks, and connectivity options.
    - Design a flexible architecture. Think of it as a scaffold to build on as you pull project by project into the cloud. You don't need to definitively plan out every piece of the migration before you move, unless you really like spending massive amounts of money on project managers and consultants.
- Then start building your scaffold.
    - Start with foundational shared services you will need across all your cloud environments: logging/monitoring receivers, cloud assessment (cloud security posture management), cloud automation (including cloud detection and response), other visibility/monitoring tools, and IAM.
    - Your will likely need at least one transit network (a central virtual network used to peer your other virtual networks, even across cloud environments). Design this network (in its own account) for transit only — not to contain any actual resources (except possibly some shared services).
    - Many shared services work better as "endpoint services", which are published within the cloud provider but don't require network peering outside. Implementation is quite different at each cloud provider, so we can't get more specific in this research, but endpoint services really enable you to take advantage of cloud software defined networks, and reduce reliance on fixed IP addresses and traditional network segmentation.
    - Build infrastructure as code templates for "landing zones" for the new accounts you will create for various projects. These can and should embed foundational security controls, such as links to transit networks and endpoint services (as appropriate), baseline network security controls, and implementation of the assessment, monitoring/logging, visibility, automation, IAM, and other core tools you use to track each of your environments.
    - Don't forget, these are just pointers to get you started — we aren't trying to downplay the complexity of these projects.
- With the scaffold in place, it's time to start migrating workloads.

- This is an iterative process. Just as you build a scaffold and smaller environments, move your projects over in prioritized order to learn as you go.
- As you move each project over, try to refactor and rearchitect to the best of your ability. For example you should "fit the network to the application" — you can now have multiple software designed networks, each containing the bare minimum to support one project. This really helps reduce attack surface and compartmentalize.
- Keep up with continuous assurance. Mistakes happen and your shared monitoring, visibility, and remediation tools will help reduce exposure. Don't wait until the end for one big assessment.

These migrations and transformations can be overwhelming if you try to plan everything out as one giant project. If you think in terms of building central services and a scaffold, then migrating projects iteratively, you reduce risk while increasing your chance of success.

Snap Migration

> *Clarice clicked to swap in the new security group and closed out the last (for now) high priority ticket. She checked her queue and the latest assessment results and everything looked okay.*
>
> *"Well," she thought to herself, "I guess it's time to start hitting the internal groups."*
>
> *She Slacked Bill, "I think I have the dev teams locked down to our sanctioned CIDRs — how's the service endpoint project going?"*
>
> *"Pretty good... the log receiver is set up and we are close to cutting over the customer DB. We still need to peer the CRM stack's network, but I think we can start weaning off some of the marketing apps and shut down those direct network connections."*
>
> *"Cool. Paul is assessing the rest of the spaghetti mess. It will take a bit to break out most of the apps into their own accounts, but at least we have a good base for the new projects."*

Snap migration can be the riskiest of all adoption patterns. Short timelines, critical resources, and rarely the skills and staff needed. They combine the messiness of the developer-led pattern with the scale of data center transformation. In our experience these projects often include a heavy dose of cloud provider or consultant pressure to move fast and gloss over complexity.

Let's start with our principles:

- Your primary objective is to minimize immediate risk while creating a baseline to use as you clean things up over time after the cutover.
    - Get the right people with the right skills. This includes training, hiring, and consulting. Make sure you really vet the people you are bringing in — even your cloud provider's experts may be fresh out of school with little real-world experience.
    - Don't just copy and paste your existing network into the cloud. This approach *always* fails within 18-24 months, for many already-cited reasons.
    - Constantly look for opportunities to control blast radius. Use multiple virtual networks and accounts, and only connect them where needed.
    - You typically won't have time in a snap migration for any serious refactoring or rearchitecting. Instead focus on a strong scaffold and management controls, with the expectation that you can start making things a little more cloud native once the main cutover is complete.

These are simply bad situations, which you need to manage as well as possible. Some smart decisions early on will go a long way to helping you set yourself up for iterative cleanup after the mad rush is over.

- Start by building a scaffold — not a parking lot.
    - Follow our recommendations for the *data center transformation* pattern.
    - While you might need to replicate your current network, *nothing says you have to do that in a single virtual network*. With peering and transit networks, you can architect your new cloud network with subnets in separate virtual networks and accounts based on projects, then connect them together with your cloud provider's peering capabilities. For example you can create the 10.0.1.0/24 subnet in one virtual network in one cloud account, and the 10.0.2.0/24 subnet in an entirely different virtual network and account, then peer them together.
    - This improves your long-term security because account segregation, even across peered networks, helps manage the service limit and IAM issues which cause so many problems when everything is in one account. For example if different projects share the same virtual network, it is hard to designate IAM privileges so the various administrators cannot affect each other's resources.
    - Knowing your subnets and connectivity requirements are key factors for success.
- As with our data center transformation pattern, build your shared services after (or concurrently with) your network scaffold.

- Be cautious and judicious about allowing Internet access. Controlling the public perimeter early is crucial. Quite a bit can be accidentally opened up during data migration, as teams rush to throw assets into the cloud, so make sure you keep a continuous eye on things.
    - Also track network connections to your on-premise environments. At some point many of these openings should be shut down, as projects complete migration and no longer need to call back to the doomed data center.
- To the best of your ability, also implement in-cloud network segregation with security groups. Another issue we often see is excessive security group openings within the network — ops, devs, or even security may not know all the right port and protocol combinations for a given application. There is literally zero cost to more security groups, which are effectively firewalls around every resource. Use them to your advantage and dial down permissions.
- In the long term you will want to sweep through and refactor and rearchitect where you can. This is much easier if you migrated into multiple accounts and virtual networks.

Native New Build

> *Maria checked the assessment results from Dev and everything looked good. The Internet facing bit was just a single-page app hosted in S3, but the Lambda functions needed network access to hit the Elasticsearch cluster. The security groups were locked down tight and the logging all hooked in using S3 and SNS so they didn't need to link back using the logging PrivateLink. The security and networking IAM roles had the right permissions for the monitoring tools and the IR team could escalate to write access as needed.*
>
> *"Hey, John, do you know what org unit we are dropping this marketing app into? I want to check the SCPs to make sure nothing will break."*
>
> *"Yep, let me check..." he replied, "looks like the default marketing one."*
>
> *"Cool, I'll go approve it for prod and promote the Terraform build."*

Cloud native doesn't mean a project is inherently secure, but it does completely shift the security and networking focus. The key principles are:

- *Cloud security and operations start with architecture and end with automation.* A well-designed architecture will reduce most risks. Automation maintains a strong and safe posture over time.
- Serverless, containers, and other emerging technologies are the norm. You may or may not have networks, but any networks you do have will be quite different from traditional infrastructure.

- Your public-facing perimeter is more than just what your virtual networks expose. Many services in cloud providers are (potentially) directly public-facing, so must be managed at the configuration level.
    - Subdomain takeovers in cloud are very common due to these services. Make sure you are monitoring at the DNS level — not just IP addresses.
- The biggest issues we see for this pattern are mostly related to governance. Dev teams are allowed to move fast and break things, and while there is nothing inherently wrong with that, it becomes a problem when they move faster than security can contain risk. Early engagement, architectural support, continuous monitoring, and strong team relations are essential for success.
- Fit networks to applications. This is a core philosophy: start with the application's needs and build the network to fit them.

As your organization becomes more and more cloud native, you will want to start with people and a secure foundation for individual projects to execute on.

- Invest in people. Hire smart, train them, and allow them to become experts on your deployment platforms. When you transition employees with traditional skills to build cloud-native projects, don't force them to split their time. Let them focus.
- Your scaffold will be similar to the ones we recommend for data center transformation, but you should plan on different network and security architectures. In many cloud-native deployments there is no customer-managed network.
    - Rely more on object storage (such as S3), service endpoints, API gateways, and other tools which don't require managing IP addresses for shared services. That said, you will always still need some virtual networks and a transit gateway.
- Set standards for your container networks and integrate them into your overall network management. Publish guidelines and even templates to build an easy path for independent teams to follow. Container networks can be easy to lose track of, especially when they are self-contained.
- Continuous integration and infrastructure as code are your friends. Develop supported templates for different patterns (*e.g.,* serverless, containers, standard virtual networks) which integrate your monitoring, logging, management, and security tools. Project teams can build these into their own templates; offering an easy path again helps encourage compliance.
- You will need to continuously monitor and enforce standards across hundreds or even thousands of cloud accounts. Build this early and automate provisioning through infrastructure as code and other automation capabilities.

As a final reminder, cloud-native architectures and operations are very different. Your core skills and objectives are the same, but the implementation details are incredibly different and often don't even translate between cloud providers. Providers launch new features and services on a daily basis, further challenging overworked security and operations teams.

Learn, take your time, work well with project teams, be nimble, and if you are in management... give your people time to keep up with the rapid rate of change.

If you have any questions on this topic, or want to discuss your situation specifically, feel free to send us a note at info@securosis.com.

About the Analyst

Rich Mogull, Analyst and CEO

Rich has twenty years experience in information security, physical security, and risk management. These days he specializes in cloud security and DevSecOps, having starting working hands-on in cloud nearly 10 years ago. He is also the principle course designer of the Cloud Security Alliance training class, primary author of the latest version of the CSA Security Guidance, and actively works on developing hands-on cloud security techniques. Prior to founding Securosis, Rich was a Research Vice President at Gartner on the security team. Prior to his seven years at Gartner, Rich worked as an independent consultant, web application developer, software development manager at the University of Colorado, and systems and network administrator.

Rich is the Security Editor of TidBITS and a frequent contributor to industry publications. He is a frequent industry speaker at events including the RSA Security Conference, Black Hat, and DefCon, and has spoken on every continent except Antarctica (where he's happy to speak for free — assuming travel is covered).

About Securosis

Securosis, LLC is an independent research and analysis firm dedicated to thought leadership, objectivity, and transparency. Our analysts have all held executive level positions and are dedicated to providing high-value, pragmatic advisory services. Our services include:

- **Primary research publishing**: We currently release the vast majority of our research for free through our blog, and archive it in our Research Library. Most of these research documents can be sponsored for distribution on an annual basis. All published materials and presentations meet our strict objectivity requirements and conform to our Totally Transparent Research policy.
- **Cloud Security Training and Advisory Services**: Securosis built the Cloud Security Alliance's CCSK and Advanced Cloud Security Practitioner training programs. We also provide custom trainings tuned to your needs. Securosis is a premier Cloud Security Alliance Training Partner and provides in-depth strategic and technical assessments and advisory services. We also support a limited number of members in our Cloud Security Coaching program.
- **Retainer services for vendors**: Although we will accept briefings from anyone, some vendors opt for a tighter, ongoing relationship. We offer a number of flexible retainer packages. Services available as part of a retainer package include market and product analysis and strategy, technology guidance, product evaluation, and merger and acquisition assessment. Even with paid clients, we maintain our strict objectivity and confidentiality requirements. More information on our retainer services (PDF) is available.
- **External speaking and editorial**: Securosis analysts frequently speak at industry events, give online presentations, and write and/or speak for a variety of publications and media.
- **Other expert services**: Securosis analysts are available for other services as well, including Strategic Advisory Days, Strategy Consulting engagements, and Investor Services. These tend to be customized to meet a client's particular requirements.

Our clients range from stealth startups to some of the best known technology vendors and end users. Clients include large financial institutions, institutional investors, mid-sized enterprises, and major security vendors.

Additionally, Securosis partners with security testing labs to provide unique product evaluations that combine in-depth technical analysis with high-level product, architecture, and market analysis. For more information about Securosis, visit our website: <http://securosis.com/>.