

The Cloud Foundation Framework A Blueprint for Building Enterprise Cloud Environments from Scratch

Siri Chandana Chava,
Independent Researcher

Abstract

Most enterprise cloud estates are not designed. They accrete. A first workload is deployed under deadline pressure, a second one copies the first, and within a year the organization is operating a production environment whose identity model, network layout, and cost boundaries were never deliberately chosen. Retrofitting a foundation onto a running estate is one of the most expensive and disruptive projects an infrastructure team can undertake. This article presents the Cloud Foundation Framework, a layered model for establishing the non-negotiable substrate of an enterprise cloud environment before the first production workload is deployed. The framework organizes foundation work into six dependent layers: tenancy and governance, account topology, identity, network, guardrails, and observability. It argues that the order of construction matters as much as the components themselves. Drawing on foundations I have built for regulated and large-scale enterprises across AWS and Azure, I describe each layer, the failure modes that appear when it is skipped, a pragmatic definition of a minimum viable foundation, and a small set of measurable signals that indicate whether a foundation is actually working. The intent is not a product-specific runbook but a durable mental model that survives changes in tooling and cloud provider.

Keywords: *cloud architecture, landing zone, enterprise infrastructure, governance, policy-as-code, identity, network segmentation, FinOps*

1. Introduction

There is a moment, early in the life of almost every cloud program, when someone deploys the thing that works. A small team, given a mandate and a corporate credit card, stands up a virtual network, a handful of servers or a managed database, a pipeline to ship code, and, crucially, an application that customers can reach. It is a genuine accomplishment, and it is usually celebrated. It is also, more often than not, the moment the organization stops thinking about its foundation and starts accumulating one by accident.

The reason this happens is not incompetence. It is incentives. The first workload has a deadline, a sponsor, and a visible payoff. The foundation has none of those things on day one. The account structure that will contain a hundred future workloads, the identity model that will govern who can touch production, the network design that will either contain a breach or let it spread: none of it is visible until it becomes load-bearing, and by the time it is load-bearing it is very hard to change.

I have spent much of my career being called in at exactly the wrong moment. Eighteen months after the first workload, the estate has grown to dozens of accounts or subscriptions, nobody can say with confidence who has administrative access to what, a single flat network means a compromised test server can reach the production database, and the monthly bill has become a mystery that finance escalates and engineering cannot explain. The work of untangling all that is real and it is expensive, and almost all of it could have been avoided by a few weeks of deliberate design at the start.

This article is my attempt to write down what that deliberate design consists of. I call it the Cloud Foundation Framework. It is deliberately provider-neutral. The specific services change (what Azure calls a management group, AWS calls an organizational unit, and they are the same idea wearing different names), but the layering does not. The central claim is simple. A cloud foundation is a stack of six dependent layers, each one resting on the layers beneath it, and the value comes from building them in order, before the workloads arrive, rather than retrofitting them afterward.

2. Why Foundations Cannot Be Added Later

It is tempting to believe that foundation work can be deferred, that you can ship the first few workloads and then, once the value is proven, go back and put the substrate in place. In practice this almost never works, and understanding why is the strongest argument for doing the work up front.

The problem is that foundational decisions are not local. When you decide, on day one, that everything will live in a single account with a flat network and a shared set of administrator credentials, you are not making a decision about day one. You are making a decision about every workload that will ever be deployed into that environment, because each of them will be built on the assumption that the flat network exists and the shared credentials work. Changing the assumption later means changing every workload that depends on it, and by then there are many, owned by teams who have moved on to other things.

Consider identity. Adding single sign-on and least-privilege roles to a greenfield environment is a configuration exercise that takes a few days. Adding them to an environment where forty engineers have been using long-lived access keys for a year is a migration project with a change-management plan, a communication campaign, and a tail of broken automation that nobody documented. The technical work is the same. The surrounding cost is an order of magnitude higher, and it comes with the risk of an outage every time you revoke a credential that something quietly depended on.

Network segmentation is worse, because the cost of retrofitting it is paid in risk rather than effort. A flat network does not announce itself as a problem. It works perfectly well right up until the day an attacker compromises a low-value host and discovers that, from there, they can reach everything. Segmenting a running estate means untangling which systems actually talk to which, a dependency map that no one has and that can only be reconstructed by watching traffic and breaking things to see what complains.

The pattern generalizes. Foundational layers are cheap to establish when the estate is empty and expensive to establish once it is full, and the expense grows with the square of the estate rather than linearly, because every pair of workloads that share an assumption has to be reconciled. This is the core justification for the framework. Foundations are not merely nice to have. They are nearly impossible to add once the moment to add them has passed.

3. The Framework

The Cloud Foundation Framework organizes foundation work into six layers, shown in Figure 1. The layers are ordered by dependency. Each one relies on the correctness of the layers beneath it, and building them out of order produces work that has to be redone. You cannot meaningfully design an account topology before you have established a governance root to own it. You cannot design a network before you know what accounts it must connect. You cannot write effective guardrails before you know what identity model they are meant to enforce.



Figure 1. The Cloud Foundation Framework

Figure 1. The Cloud Foundation Framework: six dependent layers, built bottom-up before the first workload lands.

The layers are not phases in a project plan so much as strata in a structure. In practice you iterate across all of them, but the dependency direction is fixed, and information and constraints flow upward. A decision at the tenancy layer, say that the organization will operate a single management tenant with a hard separation between billing authority and workload access, propagates up through every layer above it. A decision at the observability layer, by contrast, changes nothing beneath it. This asymmetry is what makes the ordering real rather than arbitrary.

The sections that follow take each layer in turn, from the bottom up. For each, I describe what belongs in it, the specific failure that appears when it is neglected, and the smallest version of it that is still worth building.

4. The Six Layers

4.1 Tenancy and Governance Root

The foundation of the foundation is the answer to a question most teams never explicitly ask: who owns the top of this environment? In every major cloud there is a root of authority, an organization in AWS, a tenant and root management group in Azure, an organization in Google Cloud, from which billing, policy, and account creation ultimately derive. The mistake I see most often is that this root is created casually, using an individual engineer's identity, with the billing relationship and the policy authority tangled together in the same set of hands.


The governance root should be treated as the most privileged and least-used part of the environment. In practice that means a dedicated management account or tenant that hosts nothing but the organization's control plane, billing ownership held by the organization rather than an individual, a small and named set of humans with break-glass access to the root protected by hardware multi-factor authentication and monitored on every use, and a clear separation between the authority to spend money and the authority to touch workloads. When these are conflated, the person who can approve a budget can also, technically, delete production, and no audit will ever untangle which hat they were wearing.

Skip this layer and the failure is not immediate. It surfaces during the first serious audit or the first security incident, when the organization cannot answer basic questions about who has ultimate control and how root access is governed. That is a bad moment to discover the answer is nobody knows.

4.2 Account and Subscription Topology

With a governance root in place, the next decision is how to carve the estate into accounts or subscriptions. This is the single most consequential structural choice in the entire foundation, because the account boundary is the strongest

isolation boundary the cloud provides. Anything inside an account can, with enough misconfiguration, reach anything else in that account. Between accounts, isolation is the default and connection is deliberate. That asymmetry is a gift, and the topology is how you use it.



Environment isolation is enforced at the account boundary, not by naming convention.
Figure 2. A representative multi-account landing zone.

Figure 2. A representative multi-account landing zone. Environment isolation is enforced at the account boundary, not by naming convention.

Figure 2 shows a topology I return to repeatedly, adjusted for scale. A small number of organizational units group accounts by purpose rather than by team. A Security unit holds the accounts that must remain trustworthy even if everything else is compromised: an immutable log archive and the security tooling that watches the rest of the estate. An Infrastructure unit holds shared network and shared services. A Workloads unit holds the actual applications, separated into distinct accounts for production, staging, and development. A Sandbox unit holds space for experimentation, walled off so that a curious engineer cannot accidentally spend the quarter's budget or expose data.

The principle underneath the picture is that environment separation must be enforced at the account boundary and never by naming convention alone. I have lost count of the environments where production and development lived in the same account, distinguished only by a tag or a prefix, with a shared set of permissions that in theory kept them apart. In theory. The account boundary is enforced by the cloud provider and cannot be bypassed by a fat-fingered role policy. A naming convention is enforced by hope. When the blast radius of a mistake is the difference between one environment and all of them, hope is not an acceptable control.

The smallest useful version of this layer is three accounts (production, non-production, and a security or logging account) plus the management account above them. That is enough to establish the habit of isolation. Everything else is elaboration on the same idea.

4.3 Identity and Access Baseline

If the account topology decides where things are isolated, identity decides who can cross those boundaries. This is the layer where I see the widest gap between organizations that treated it as foundational and those that did not, and the gap compounds faster than any other.

The baseline has a few non-negotiable elements. There is a single source of truth for human identity, a corporate identity provider, federated into every account, so that access is granted to people who exist in one directory and revoked in one place when they leave. Nobody uses long-lived static credentials for interactive access. Humans assume roles, machines use workload identity, and the standing set of long-lived secrets trends toward zero. Permissions are granted least-

privilege and, wherever the tooling allows, just-in-time, so that elevated access is requested, time-boxed, and logged rather than held permanently against the day it might be needed.

There is also a break-glass path, a deliberately uncomfortable and heavily monitored way to obtain emergency access when the normal identity plane is unavailable. It exists because the one thing worse than access that is too easy is discovering, during an incident, that legitimate responders are locked out.

The failure mode of a neglected identity layer is quiet and cumulative. Access is granted and never revoked. Credentials are shared because requesting them is slow. The set of people who can touch production grows and grows, because removing access is scary and granting it is easy. Two years in, the honest answer to who can access production is everyone who has ever needed to, which is to say no one really knows. Establishing least privilege from an empty starting point is straightforward. Recovering it from that state is a multi-quarter program with political friction at every step.

4.4 Network Foundation

The network layer inherits its structure from the account topology and its access rules from the identity layer, which is why it sits above both. Its job is to make the connections that workloads legitimately need while ensuring that everything else is denied by default and that traffic leaving the estate is inspected rather than assumed benign.




Figure 3. Centralized inspection and hybrid connectivity through a shared hub.

Figure 3. A hub-and-spoke network foundation centralizes inspection, shared services, and hybrid connectivity.

The topology I reach for by default is hub-and-spoke, shown in Figure 3. A central hub network hosts the shared, security-sensitive functions: egress inspection, DNS resolution, and the private circuits that connect to on-premises data centers. Each workload environment lives in its own spoke, peered to the hub but not to its siblings. Production cannot reach development at the network layer because there is no path between their spokes. Both reach shared services only through the hub, where traffic can be observed and controlled.

This design does two things that a flat network cannot. First, it centralizes the controls that are expensive to replicate, so you inspect egress in one place rather than configuring it correctly in fifty. Second, it makes lateral movement structurally difficult, so that a compromise in one spoke does not automatically become a compromise of the whole estate. For organizations with on-premises systems, the hub is also the natural place to terminate the private connectivity that makes hybrid operation possible, so that the migration path and the steady state share the same design.

The minimum viable network foundation is a single hub with private connectivity and controlled egress, plus one spoke per environment. Segmentation can grow finer over time. What matters at the start is that the default is deny, and that no workload sits directly on the open internet without an inspection point between it and the outside.

4.5 Guardrails as Policy-as-Code

The first four layers establish structure. The fifth keeps that structure from eroding. Guardrails are the automated controls that prevent the estate from drifting away from its intended configuration. The essential word is automated, because a control that depends on a human remembering to check it is not a control. It is a wish.




Figure 4. Guardrails treated as code and enforced continuously, not audited quarterly.

Figure 4. Guardrails treated as a control loop: defined as code, validated in CI/CD, enforced preventively, monitored continuously, and remediated automatically.

The right mental model, shown in Figure 4, is a control loop rather than a checklist. Policy is defined as code and version-controlled like any other artifact. It is validated in the delivery pipeline, so that a configuration which violates policy is rejected before it is ever deployed. Where the platform allows, controls are preventive, and the non-compliant action simply cannot be taken. Where prevention is not possible, detective monitoring catches drift after the fact and, increasingly, remediates it automatically. The findings then feed back into the policy definitions, so the baseline tightens over time instead of decaying.

There is a meaningful distinction between preventive and detective controls, and mature foundations use both deliberately. A preventive control, say a policy that refuses to create a storage bucket that is publicly readable, is strictly better when it is feasible, because it makes the bad state impossible rather than merely visible. But not everything can be prevented without blocking legitimate work, and for those cases a detective control that flags the deviation and either alerts or corrects it is the pragmatic choice. The art is in knowing which to use where, and in resisting the temptation to make every control preventive, which produces an environment so rigid that teams route around it.

A foundation without this layer looks fine for a while and then slowly rots. Someone opens a security group to the world to debug a problem and forgets to close it. A public bucket is created for a legitimate reason and outlives it. Encryption is disabled once and becomes the new normal. None of these is catastrophic alone. Collectively they are how well-designed estates become the messes that people like me are eventually called to fix. Policy-as-code is the difference between a foundation that holds its shape and one that merely started out the right shape.

4.6 Observability and Cost Baseline

The top layer is the one that lets you see the estate you have built, both what it is doing and what it is costing. I group operational observability and cost together deliberately, because they answer the same kind of question, which is what is actually happening here, and because both are far easier to establish before workloads exist than after.

On the observability side, the foundation provides a centralized destination for logs, metrics, and traces, with the security-relevant subset flowing into the immutable log archive established back at the account layer. The goal is not to instrument every application from day one, since teams will do that as they build. The goal is to guarantee that when

they do, there is somewhere trustworthy for the signal to go, and that the platform's own control-plane activity is captured whether or not any application opts in.

On the cost side, the foundation establishes the machinery of financial accountability before there is a bill large enough to fight about. That means a tagging or labeling policy enforced by the guardrail layer, so that every resource can be attributed to an owner and a purpose. It means budgets and anomaly alerts wired up per account. And it means a cost-visibility surface that finance and engineering look at together, rather than through the distorting lens of a single monthly invoice. Organizations that add this after the fact spend months in archaeology, trying to attribute historical spend to teams that never tagged anything. Organizations that make attribution a precondition of deployment never have that problem.

5. The Minimum Viable Foundation

A fair objection to everything above is that it sounds like a mandate to spend six months building infrastructure before delivering any business value, precisely the kind of ivory-tower architecture that gives architects a bad name. That is not the argument. The framework describes the shape of a complete foundation. It does not require that the complete foundation exist before workload one.

What must exist before the first production workload is a minimum viable foundation: the thin version of each layer that establishes the right defaults and, more importantly, the right habits. Concretely, that is a governance root with separated billing and break-glass access, three or four accounts that enforce environment isolation, federated identity with no standing long-lived human credentials, a single hub with controlled egress and one spoke per environment, a handful of preventive guardrails covering the highest-risk misconfigurations, and cost attribution enforced from the first resource. This is a matter of weeks, not months, and every hour of it is repaid many times over.

The distinction that matters is between the layers, which are all present from the start even if thin, and the elaboration within each layer, which grows as the estate does. You do not need fifty guardrails on day one, but you do need the guardrail layer to exist, because adding the fiftieth guardrail to an estate that already has the mechanism is trivial, and adding the first guardrail to an estate that has none is a project. Build all six layers thin, then thicken them over time.

6. Anti-Patterns Worth Naming

A few failure modes recur often enough across the environments I have inherited that they deserve names.

The first is the single-account estate, where the entire organization operates inside one account or subscription and relies on internal permissions to keep environments apart. It is the most common starting point and the most expensive to unwind, because unwinding it means moving live workloads across an account boundary that was never meant to be crossed.

The second is convention-as-control, the belief that a naming scheme, a tagging standard, or a wiki page describing the rules constitutes governance. Conventions describe intent. Only enforced policy constrains behavior. The gap between them is where incidents live.

The third is the privileged bootstrap that never gets cleaned up. This is the powerful identity used to set the environment up, convenient during construction and then quietly retained, becoming a permanent skeleton key that appears in no access review because it predates the access-review process. Every environment has one. Mature environments have found and removed it.

The fourth is foundation-by-copy, where the second workload is built by copying the first, inheriting all of its undocumented assumptions, and the third copies the second, so that the estate's design becomes an accident propagated by imitation. This is what happens by default in the absence of a deliberate foundation, and it is why the absence of design is itself a design. Just a bad one.

7. Knowing Whether the Foundation Works

Because the foundation is invisible when it succeeds, it needs its own signals. Otherwise the only feedback it ever generates is the sound of it failing. A small number of questions, asked honestly and periodically, tell you most of what you need to know.

Can you answer, in minutes and with evidence, who has administrative access to production? Can you attribute every dollar of last month's cloud spend to an owning team? If a workload environment were fully compromised, can you state with confidence which other environments would remain unaffected, and defend the claim with the network design? When a policy violation occurs, is it prevented, or at least detected and remediated, without a human noticing first? Can a new team deploy into the estate through a paved path that makes the compliant choice the easy one, rather than by negotiating with a central team for weeks?

An environment where these questions have clean answers has a working foundation, whatever its specific tooling. An environment where they provoke uncomfortable silences has an accreted one, and the silence is the first symptom of costs that have not yet arrived but certainly will. The value of the framework is that it turns those eventual costs into a few weeks of deliberate work, at the one moment when that work is cheap.

8. Conclusion

Cloud foundations are not glamorous. They produce no feature, win no customer, and appear in no demo. Their entire value is negative: the incidents that never happen, the audits that pass without drama, the migrations that are never needed because the thing was built correctly the first time. That invisibility is exactly why they are skipped, and why skipping them is so reliably expensive.

The Cloud Foundation Framework is an argument for spending a small, bounded amount of deliberate effort at the one moment when a foundation is cheap to build, which is before the estate fills up. Its six layers (tenancy, accounts, identity, network, guardrails, and observability) are ordered by dependency, built bottom-up, and present from the start even in their thinnest form. None of the individual practices is novel, and competent architects have used all of them. What the framework offers is the discipline of sequence and the insistence that the substrate be chosen rather than inherited. In a decade of being called in to repair estates that grew without one, I have never regretted building a foundation first, and I have never seen an organization that skipped it avoid paying for the omission. The bill usually arrives at the worst possible time, and always at many times the original cost.

References

- Amazon Web Services, "AWS Well-Architected Framework," AWS Whitepaper, 2020.
- Microsoft, "Cloud Adoption Framework for Azure," Microsoft documentation, 2020.
- S. Rose, O. Borchert, S. Mitchell, and S. Connelly, "Zero Trust Architecture," NIST Special Publication 800-207, National Institute of Standards and Technology, 2020.
- P. Mell and T. Grance, "The NIST Definition of Cloud Computing," NIST Special Publication 800-145, 2011.
- B. Beyer, C. Jones, J. Petoff, and N. R. Murphy, Eds., Site Reliability Engineering: How Google Runs Production Systems. Sebastopol, CA: O'Reilly Media, 2016.
- N. Forsgren, J. Humble, and G. Kim, Accelerate: The Science of Lean Software and DevOps. Portland, OR: IT Revolution Press, 2018.
- J. Humble and D. Farley, Continuous Delivery: Reliable Software Releases through Build, Test, and Deployment Automation. Boston, MA: Addison-Wesley, 2010.