HOMOGENEOUS VS. HETEROGENEOUS CLOUDS: HOW TO CHOOSE YOUR HYBRID CLOUD

When it comes to cloud, you've got a lot of decisions to make. Choosing your cloud strategy starts with determining public, private, or hybrid environments. Many organizations opt for both—the hybrid cloud approach.

But that's not the end of the decision. Once you decide to pursue a hybrid cloud strategy, you must also decide whether to go with a homogenous or a heterogenous hybrid cloud environment.

What does this mean? Let's find out.

Hybrid cloud basics

Choosing a homogenous or heterogenous cloud is based on a hybrid cloud strategy. So let's briefly review the three main cloud environments. **Public cloud** is cloud computing that's delivered across the internet. **Private cloud** is cloud computing that's dedicated to your single organization—no others—and it can be on premises or off-site through your vendor.

Hybrid cloud is a mix of these two environments. It's what most organizations opt for in their cloud strategy. Hybrid allows you to take advantage of:

- The affordability and agility of public cloud when you want it—for certain data storage types, many "as a service" offerings.

- The security and privacy of private cloud when you need it—for securing certain types of data or for running heavy compute loads.

Once you opt for a hybrid cloud, you have to make another decision: homogenous or heterogenous.

What is a homogeneous cloud?

Simply put, a homogenous cloud is one where everything is from the same vendor. That single vendor supplies both your public cloud access and offerings and any private cloud you might have, both on-prem or off-site.

More technically, a homogeneous cloud is one where the entire software stack, from the hypervisor (remote cloud provider), through various intermediate management layers, all the way to the end-user portal, is provided by a single vendor.

Many enterprise providers used to offer this approach, but this fell out of the market years ago. Traditionally, hybrid cloud options were simply software stacks that you'd run on premises that also integrated with a public cloud. This has changed. Cloud providers are now releasing physical appliances—hardware. More recent pushes from Azure and AWS, however, are aiming for a homogenous cloud renaissance.

What is a heterogeneous cloud?

A heterogeneous cloud, on the other hand, integrates public and private components from more than one vendor, either at:

- Different levels, such as a management tool from one vendor driving a hypervisor from another
- The same level, where a single management tool drives multiple hypervisors

For example, you'd choose a public cloud provider, like Azure, GCP, or AWS, and then pair it with a private offering like those from VMware, CloudStack, or OpenStack.

Homogenous cloud: pros and cons

The argument for homogeneous environments is that because everything comes pre-integrated they are easier to set up, and if something goes wrong there is only one responsible party—"one neck to wring" as the saying has it. But there are plenty more benefits to a homogenous cloud.

Homogenous cloud benefits

- Essentially turnkey, with "off the shelf" functionality
- Easier to install and set up
- Easier from an operations and management standpoint. Because the public and private portions are from a single provider, they're designed to work together.
- Services like disaster recovery, security, governance, and monitoring are built-in and span both environments
- Cheaper because the on-prem portion is delivered as drop-in hardware or a prebuild rack
- Talent needs skills specific only to that provider

That's a lot of benefits! But there's one significant drawback—vendor lock-in.

Homogenous cloud drawbacks

Easy to use often means harder to leave. By giving so much power to one vendor, users place themselves at the mercy of that vendor's commercial and technical strategy. Leaving that vendor, no matter the reason, becomes risky, expensive, and difficult, especially for security and governance strategies.

In farming, this is known as a monoculture, that is, only a single crop is grown. Superficially, this is an attractive idea, as farmers can specialize and take advantage of economies of scale. The problem is that any disease, frost, drought, or other event that affects that crop can wipe out the entire harvest. To mitigate that risk, farmers try to sow several different types of crop, as insurance against losing the entire harvest.

The same arguments apply in IT. The advantage of an IT monoculture is that everyone can specialize in that one vendor's tools and it is easy for admins to cover for each other.

The downsides are a bit different: on the technical side, features will be available when—and if—the vendor chooses to develop them. The real pain often comes on the commercial side, because once users are "locked in" to a single-vendor strategy, they have no recourse if that vendor decided to change its pricing structure in a way that causes costs to increase.

Heterogenous cloud: lock-in vs management

Heterogeneous architectures attempt to bypass this lock-in effect by introducing components from many different vendors and allocating their use according to a common set of strategies. This gives you significantly more control over your cloud architecture—if that's something your organization needs or that your IT maturity can handle.

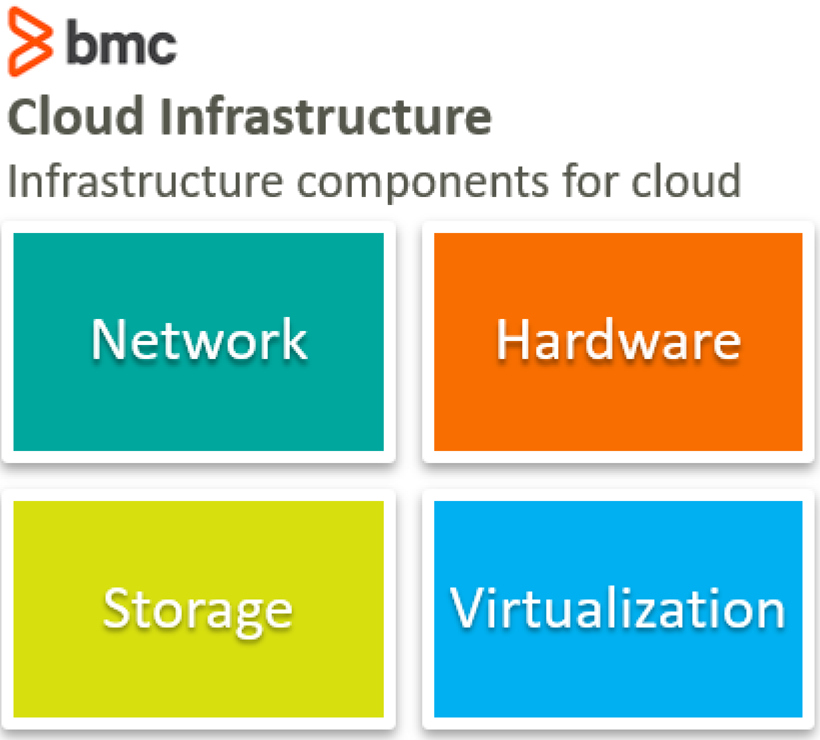

At some point, however, you'll likely need to introduce a single management component.

Defenders of the homogeneous, single-vendor approach counter charges of lock-in by pointing out that this convergence on a single management layer just moves the lock-in further up the stack—but still leaves users at the mercy of the provider of that one component.

The false equivalence between platform lock-in and supposed management lock-in is a neat rhetorical trick but does not really hold up. Here's why:

- **Management vendors need to keep up with the development pace of the managed platforms.** Otherwise, they risk falling behind the competition from other heterogeneous management vendors. Any attempt at predatory business practices will be nipped in the bud for the same reason.
- **You can more easily migrate away from management suites than change hypervisors or cloud providers.** The reason is that a platform change is almost guaranteed to cause disruption unless it is very carefully managed. Conversely, replacing a management platform, even at very short warning, will certainly be painful for IT and cause delays in delivery of new requests, but will not affect workloads that are already running on the underlying platforms.

So, heterogenous cloud proponents value this direct control over architecture. But more control inherently means two more components you need to decide whether you can handle:

- **Increased complexity.** More control of any shared cloud model means you've got more to deal with: XYZ. Even if your organization requires this control and inherent complexity, you'll also have to evaluate whether your IT organization is mature enough for that complexity.
- **Upskilled talent.** To support a more complex environment more closely, you need talent that has very specific skills. Those skills must also be software agnostic: a specialist in AWS might struggle to support products and services from other vendors.

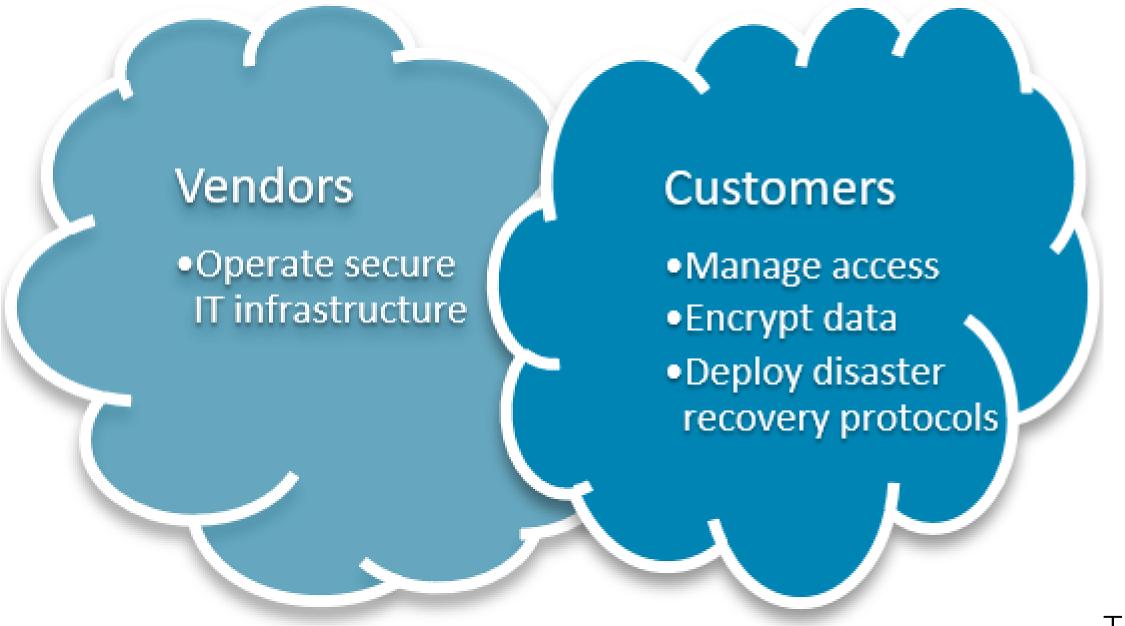

Vendors

- Operate secure IT infrastructure

Customers

- Manage access
- Encrypt data
- Deploy disaster recovery protocols

To handle this complexity and ensure the right talent, you might need to hire more advanced cloud architects, which can take time and cost more due to the IT skills gap. (This is the reason why cloud certifications are among the highest paid.) This might also mean you need to train or certify staff you already have.

How to choose between the two

Cloud migration and management aren't decisions to take lightly. You'll want to consider the specific ins and outs we've outlined in this article, like:

- Your organizational needs (business outcomes) that you expect from cloud today—and in the future
- The cost and resources
- Whether you require more control architecturally
- The maturity of your IT environment
- The talent of your cloud engineers and architects

Additional resources

For more on this topic, explore the BMC Multi-Cloud Blog or read these articles:

- Key Facets of a Smart Cloud Migration Strategy
- What are the Hidden Costs of Cloud Adoption?
- How to Secure Your Public Cloud
- Managing Large-Scale Cloud Migrations in Government Agencies: 6 Keys to Success
- Cloud Growth in 2020: Trends & Outlook
- IT Certifications: A Beginner's Guide

Run & Reinvent Podcast · Episode 7: Vinnie Lima from VVL Shares Best Practices for Creating a Cloud Migration Strategy