How to Align Your Applications to the **Optimal Infrastructure Environments**

Avoiding the Pressure to Move Every Workload to the Public Cloud

DATABANK
Data Center Evolved™

Table of Contents

Achieving the Right Balance of IT Infrastructures for **Peak Application Performance**

You're in charge of IT and find yourself caught between the pressure of preserving your budget while maintaining application performance and data integrity. With all the buzz about the big public cloud platforms, your CEO and CFO, likely, assume that's the best and least costly place to run your application workloads. They're all over you, wanting you to move to the public cloud—right now.

You also realize many companies have migrated to the public cloud, thinking it was a silver bullet, only to discover the downsides: Costs eventually spiral out of control, or the public cloud proved too difficult to manage and secure, often requiring additional resources for your already constrained IT organization.

In some cases, repatriation occurs—where enterprises move workloads out of the public cloud and back to on-premises or colocation data centers. As reported by SearchCloudComputing, an IDC survey revealed **85% of respondents have considered some form of repatriation**.

In addition, not everyone has made the jump to the cloud. DevOps.com cited another IDC report that indicates **fewer than 20% of enterprise applications have been moved from onsite data centers to the public cloud**. This has transpired, not because companies do not see value in the cloud, but rather because the process of migrating applications to the cloud is complicated—and managing applications in the public cloud might require more work than anticipated.

These trends come as no surprise, given the difficulty of achieving the right balance of environments across a hybrid infrastructure and avoiding lock-in with public cloud platforms as application requirements change. So the question still remains:

How can enterprises strike the right balance of IT infrastructures while also maintaining an agile footing to keep application workloads running at peak performance?

To help your IT team stay agile in managing your IT infrastructures, this white paper shows you how to determine the optimal alignment of applications to various infrastructures. In many cases, the best solution is a hybrid approach, where the components of an application run in different environments—spanning public cloud, private cloud, on-premises, and colocation data centers.

We explore the key questions to consider, sample scenarios, customer use-case examples, and some tips you can draw upon to drive application performance and reduce infrastructure costs. By collaborating with the right data center partner, your team can leverage multiple infrastructure technologies and platforms to establish an agile IT footing.

Application Questions Help Determine **Infrastructure Requirements**

The first step in aligning your workloads to the right infrastructure environment is to examine your application. If you're feeling overwhelmed by the number of issues to address, you're not alone. The right data center partner can help you work through the key questions to ask so you can begin to find the answers you need. (Key Application Questions to Ask)

If you are already leaning toward a particular environment, be sure to examine and validate your reasons. The benefits that a particular environment delivers today may not hold up over the long term.

For example, if you already operate your infrastructure in the cloud and use an auto-scaling service like Kubernetes, consider your peak requirements for CPU, RAM, and storage IOPS. You can then check with a private cloud provider to see what they charge to handle those requirements. If you use steady-configured resources, a colocation data center provider can advise which resources would function cost-effectively in a private cloud or on dedicated bare-metal servers.

Answers Help Prioritize Needs

By answering the key application questions, you can begin to assess, for example, the priority of cost control relative to the ability to quickly scale up and down as spikes in activity occur and then dissipate. You can also evaluate the balance of the cost against user-response times and the amount of downtime you can afford.

Key Application Questions to Ask

- What is the role of the application workload in your business?
- Is the workload mission-critical to business operations?
- How mature is the application within the development lifecycle?
- What challenges are faced by your application—cost, performance, scalability, security, compliance?
- What are the disaster recovery requirements?
- Does the application integrate with other applications running in different infrastructures?
- Is the application portable to another infrastructure environment?
- Who manages and maintains the application?

The maturity of your application is also a factor. When first building applications, IT determines the infrastructure costs to support the minimum viable product of the application. If the application is built in the cloud, the cost will be low at first because there isn't any demand for the application, and it requires a small amount of compute resources. Then when the application matures and the demand rises, you can model the costs over a long time period to get a better idea of which components belong in a data center and which components belong in the cloud.

The portability attribute is also vital to consider. In some cases, public cloud vendors entice customers to code applications to sync with the provider's proprietary platform services. This may enable the application to run more efficiently on that particular platform, but it also means the application can't be migrated to another environment without extensive recoding. While public cloud providers don't charge you to upload data, data egress fees can be expensive. In other words, public cloud providers make it easy to get in, but not so easy (and costly) to get out.

From the perspective of managing your infrastructure environments, do you have internal resources to handle the task? Are you already working with a managed service partner, or are you looking for a new partner? The infrastructures you select for your application may be influenced by who has access to manage the environments and if they can do so effectively. Sometimes you have the flexibility to manage an environment yourself; in other cases, you are limited to working with the hosting provider, whether it's a public cloud, private cloud, or colocation data center.

Giving Context to Your **Alignment Scenario**

As you work your way through the answers to the questions and sidebar presented earlier, it helps to put your situation into context by considering the scenarios of other businesses that have tried to determine the right mix of infrastructures within their hybrid IT environments. Here's a rundown of several situations that outline primary alignment issues and the approach to address the concern.

Predictable vs. Unpredictable Workloads

If your application workload requires roughly the same compute resource capacity every day, then the cost to run it in the public cloud vs. a data center will be significantly more expensive. You may also run into issues if you expect the performance of servers in the cloud will be consistent. Conversely, dedicated servers in an on-premises or colocation data center that you control deliver performance that is much more predictable.

Also consider that when sharing servers in the public cloud, the performance of virtual machines can change drastically, depending on whether you have co-tenant "noisy neighbors" that use up extra network bandwidth and storage I/O, or memory and CPU resources on the same server. This means the resources you have access to will decrease. It's also difficult to know in advance that such resources are not available until your users (or your customers) try to access the servers. Performance could suddenly decrease by 50% or more.

Another workload scenario involves static, predictable workloads—such as databases and back-office systems that run all day long. From a cost standpoint, they are best served in a colocation data center. Compared to the public cloud, it's less expensive to buy and deploy gear in a colocation rack, let it run for three to five years, and then replace it with new technology.

High Throughput Requirements

The big public cloud providers charge high fees to push bandwidth out to the Internet. In some cases, you will also be hit with a charge when transferring data from one of their data centers to another.

Therefore, if your application requires high throughput at a consistent and predictable level, an on-premises or colocation data center is your least expensive option. You might also consider a Content Delivery Network provider, which focuses on throughput delivery and costs much less than a public cloud platform.

Application Scalability vs. Resource Efficiency

If scalability is the primary driver for your workloads, you may head down a path of applications that run inefficiently. That's because developers tend to write code for cloud applications as though there will be unlimited compute resources. They sometimes build inefficient software that automatically scales the workload demand without considering the required resources. If the application is used heavily, the resources will continue to scale, and costs will spin out of control.

However, if the application is built under the premise of eventually moving to a data center, the compute resources can scale only if IT makes the conscious decision to buy and deploy another server. This approach forces developers to build applications to run more efficiently so they consume only available compute resources.

When it's time to develop an application, consider that a colocation data center can be expensive. The compute resources might cost $500K to stand up the infrastructure and then develop and test the application. It's less expensive to rent the workload from a public cloud provider where development and testing environments can more easily be connected, and there's someone in the middle to orchestrate the development lifecycle.

Security and Compliance

Mission-critical systems and sensitive data are not always appropriate to run in the public cloud. At the same time, from a cost standpoint, a company may not want to build a data center where it needs to own the gear and provide resources to manage the data center.

The best alternative is a private cloud managed by a partner with expertise across the entire technology stack—including the operating system, disaster recovery, security, and compliance. For developing the software that runs in a private cloud, be sure to connect to a public cloud that can seamlessly upload updates to the private cloud.

The best alternative is a private cloud managed by a partner with expertise across the entire technology stack—including the operating system, disaster recovery, security, and compliance.

Misconfigured Environment

In some cases, a new environment is not needed; the current environment just needs to be reconfigured. If one of your applications does not perform at expected levels, it's critical to diagnose the cause before adding compute resources that will increase costs dramatically.

In some cases, companies using a public cloud provider keep adding CPUs and memory to their environment because their application is not working. After the compute costs skyrocket, they finally take a step back and realize the infrastructure was not configured properly. It turned out the issue was load balancing, which, with properly managed services, could have been avoided.

Cost Optimization vs. Peace-of-Mind

If your application is subject to unexpected and sudden spikes in activity, worrying whether your servers will crash can keep you up at night. This is especially true if the application is used by customers to purchase products and services.

In cases like these, where cost is a secondary consideration, the public cloud or a private cloud are usually the best options. Access to additional compute resources is virtually unlimited, and infrastructures can be designed to scale and balance workloads automatically.

Private Clouds: The Best of Both Worlds

When determining the best course of action for your application scenario, it's important to realize the public cloud isn't a complex concept—it ultimately just boils down to your server running in someone else's data center. While the cloud can simplify your infrastructure operations, it obfuscates control over issues like these, and while it may generate cost savings, it could also increase costs.

Private cloud environments offer the best of both worlds. This includes the convenience of cloud features like auto-scaling and the cost savings of better-utilized resources by sharing different workloads on the same infrastructure. Private clouds also give you the ability to cost-optimize for consistent performance by avoiding the "noisy neighbor" phenomenon entirely.

Of course, if your workload experiences very large annual or seasonal spikes (e.g., one busy month every four years during the Olympics, one busy day of the year for the Super Bowl, one busy hour of the day for live streaming popular events, or a few busy days for the launch of a popular new game), the public cloud will offer the best economics. However, if the workload is relatively consistent, a private cloud or a colocation data center, most likely, wins out for reducing the long-term costs. The graphic below conveys the value of workload agility.

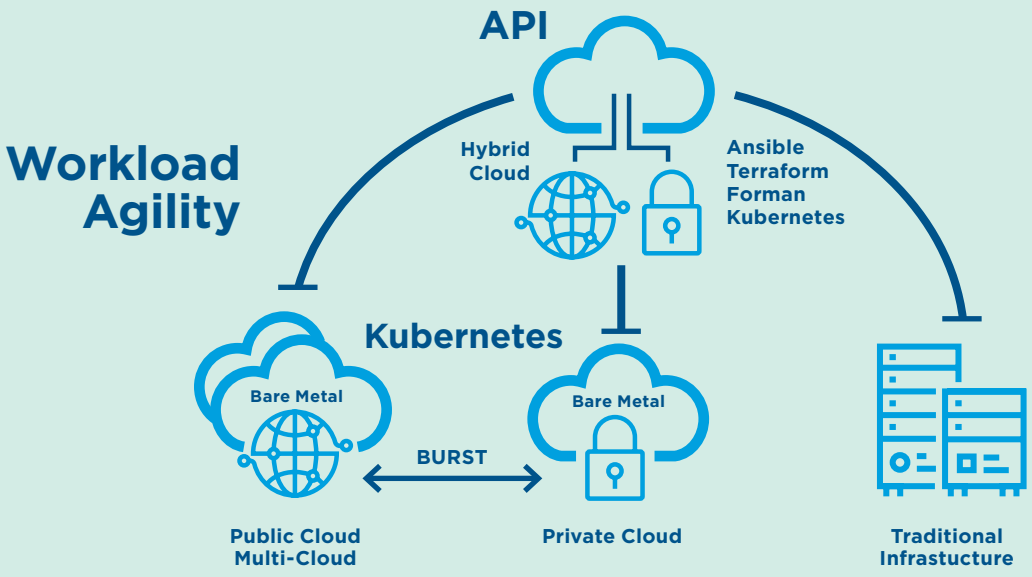

Lessons Learned from **Alignment Use-Cases**

In this section, we present three use-case examples of how companies solved the challenges of aligning applications to the appropriate infrastructure:

Video Encoding

When uploading videos, this online training company needs to convert the files from one raw format into different formats—such as high resolution, low resolution, MPEG-4, and H.265. The company also needs to encode the video as fast as possible. This requires significant compute resources within a short time window each day—approximately 20 minutes.

For this type of workload, it isn't appropriate to build a 10,000 square foot data center. Uploading the videos and encoding them in a public or private cloud is more cost-efficient. Once the encoding is done, the company can scale the cloud resources down, and then download the video to their data center. The data center stores the video on storage arrays that are under their control and much more cost-effective.

Plus, if the training company needs to serve a lot of bandwidth, it's less expensive to serve that bandwidth through a Content Delivery Network (CDN) that can clear the cache and download all the content quickly. Being able to push many gigabits per second for a few days is much less expensive through a CDN than it would be from the public cloud.

The company also manages a workload that serves videos to customers. The workload for running the database servers and credit card processing is relatively constant, so it makes sense to run these services and maintain control from the corporate data center. IT knows exactly how the reliable and redundant services are built, and they run in a way that's less expensive than in the public cloud.

Streaming Video

Even for the largest enterprises with deep financial pockets, there are components of applications that make sense to run in the cloud and other components that should run on-premises. Take, for example, a large streaming video service that stores a massive amount of video origin files on the platform of a large public cloud provider. The company has the capability and the budget to build such a system themselves.

However, they determined it's less costly to store these files on the public cloud, partly because they don't have to allocate resources to manage the files. The public cloud provider does a good job of managing files, and the streaming video company trusts the public cloud provider.

This streaming video company also realizes that to serve content to the public via the Internet efficiently and cost-effectively, it doesn't make sense to provision the content from the public cloud. It would also not be cost-effective to use a Content Delivery Network, most of which manage less traffic than this streaming video company.

As a result, the company built its own physical data centers to serve content to the customer base. Although this example is an extremely rare use-case, it demonstrates how it's vital to ask questions to determine the specifics of a workload. It's best to not base decisions on what other companies are doing if their workload or scale are vastly different.

Financial Services Cloud Costs Out of Control

A financial services company developed an application in the public cloud because it was very inexpensive to buy a server. They thought it made logical sense to then move the application into production in the same public cloud.

Suddenly, the workload spiked as the customer base grew, using more and more compute resources. The CIO woke up one day and realized the public cloud bill had ballooned to $80,000 per month. It would have been much less expensive if, after developing the application, it was deployed in a managed colocation data center.

Application-Infrastructure **Alignment Tips**

- **Act now, don't procrastinate:** The pain you feel today will likely get worse.
- **Partner with hosting companies that offer infrastructure portability:** Make sure it's easy and inexpensive to move your workloads across and/or out of their environments.
- **Avoid cloud lock-in:** Resist offers from public cloud providers to recode your software to work only with their proprietary platform technologies.
- **Give yourself options:** Collaborate with a partner that can offer more than one environment among public cloud, private cloud, colocation, and modular data centers.
- **Outsource infrastructure management:** Focus internal resources on strategic initiatives while turning to experts to maintain and monitor your infrastructure and providing connectivity to other infrastructures.
- **Tap into a broad skillset:** Find a partner with a team that can help with migrations, architecture design, database administration, security, compliance, failover, and disaster recovery.
- **Gain visibility:** Leverage a portal that lets you monitor the performance of all your infrastructures and helps you determine when it's time to expand an environment or move a workload to a new environment.

Collaborating with an Infrastructure Partner **Who Gives You Hybrid Options**

The infrastructure solution for your company will, likely, be a hybrid scenario that includes a combination of on-premises, colocation, private cloud, or public cloud data centers. You may need just one or two environments, or you may need all four. If you're hosting a large entertainment event or conference—where Internet bandwidth is limited and large amounts of data need to be processed with ultra-low latency—you might even need a modular data center deployed at the base of a cell tower. Another use-case for a modular data center is to add compute capacity to an existing office building or factory.

In addition to providing experts who advise you on the ideal hybrid infrastructure for your workloads, DataBank offers a range of IT environments:

- Colocation data centers across the US
- Multi-tenant and private cloud environments
- On-ramp interconnects to public cloud platforms
- Edge computing modular data centers

In addition, our managed services engineers help plan and migrate applications from one environment to another. From there, they ensure a strong security posture for your environment and compliance with any regulations pertaining to your industry. Our engineers monitor your environment 24x7 and maintain server performance through regular updates and patches.

We also offer contract portability—there's no fee to move workloads among our cloud and colocation data centers so you're never locked into an environment. Our services feature a portal that provides visibility of performance indicators and control. This allows you to see how each of your infrastructures is performing and when the hybrid balance may need to be adjusted, or when it's time to move a workload from one environment to another.

Advancing Your **IT Roadmap Journey**

To align IT workloads to infrastructures in a way that ensures application performance and controls infrastructure costs, the answer is almost always a hybrid environment. Rarely does it make sense to run all applications on one infrastructure. Different applications and application components are best served by splitting them across public clouds, private clouds, colocation data centers, and on-premises data centers.

And to bring end-users as close to the edge of your network as possible, you may need to deploy workloads across multiple data centers or even a modular data center pod. Additional keys include linking all environments together and having the agility to move workloads back and forth among environments as requirements change.

That's where the value of the **DataBank Data Center Evolved™** approach pays off.

If you're already in the public cloud and find it too expensive, unreliable, or unable to scale, and you don't know why, our team guides you through the process to diagnose and solve your issue. We analyze your scenario to understand what is currently going on with your applications and provide visibility into each of your infrastructure environments. This helps us work together to determine the right mix of infrastructures and where to move workloads today as well as when to adjust in the future.

The access we provide to your hybrid environment also lowers the risk of your infrastructure deployments, even as technology continues to move at break-neck speeds. So instead of procrastinating because you're worried about making the wrong decision, we give you the confidence that you're making the right decisions to advance your IT roadmap journey.

For more information on how to align your workloads to the appropriate IT infrastructures to ensure application performance and to control infrastructure costs, contact ***DataBank today.***