Cloud Economics 101

Turn Your Cloud Into a Profit Center, Not a Cost Center

lightcrest

Intro

The cloud is no longer just a “place” or a “model” - the industry has yielded several modalities labeled “cloud”, not all of them optimal for every type of business. Now that the “honeymoon” period of the cloud has ended, technologists and business executives alike must navigate the cloud landscape with guarded optimism.

The cloud is not a magical place where your data lives forever and you auto-scale with zero problems. It’s an abstraction of compute, storage, and networking infrastructure that is being marked up at a profit - to the cloud provider’s benefit, and ideally to yours as well. Unfortunately, in many cases the relationship is not mutually beneficial - consumers often pay much more than they should for suboptimal performance and have less control over their environment.

Given the dizzying amount of options, executives must ask one key question before making deployment decisions: **what is the the most cost-effective modality for my compute and storage workloads?**

We will analyze various types of businesses and discuss their cost and performance implications in the cloud.

Steady State Workloads

But before we dive into various business types and their associated cloud characteristics, let's define a term: **steady state capacity.**

Steady-state capacity is the minimum capacity you need to run your business in the cloud. It is the capacity you need to run your workloads all the time, regardless of the calendar day or current season. For example - a SaaS business that sells licenses to a healthcare management platform might **never** dip below the required capacity of 18TB of storage and 36 virtual cores. There might be slight hikes in traffic or storage utilization depending on the week, but overall their "baseline" is served by this minimum resource allocation.

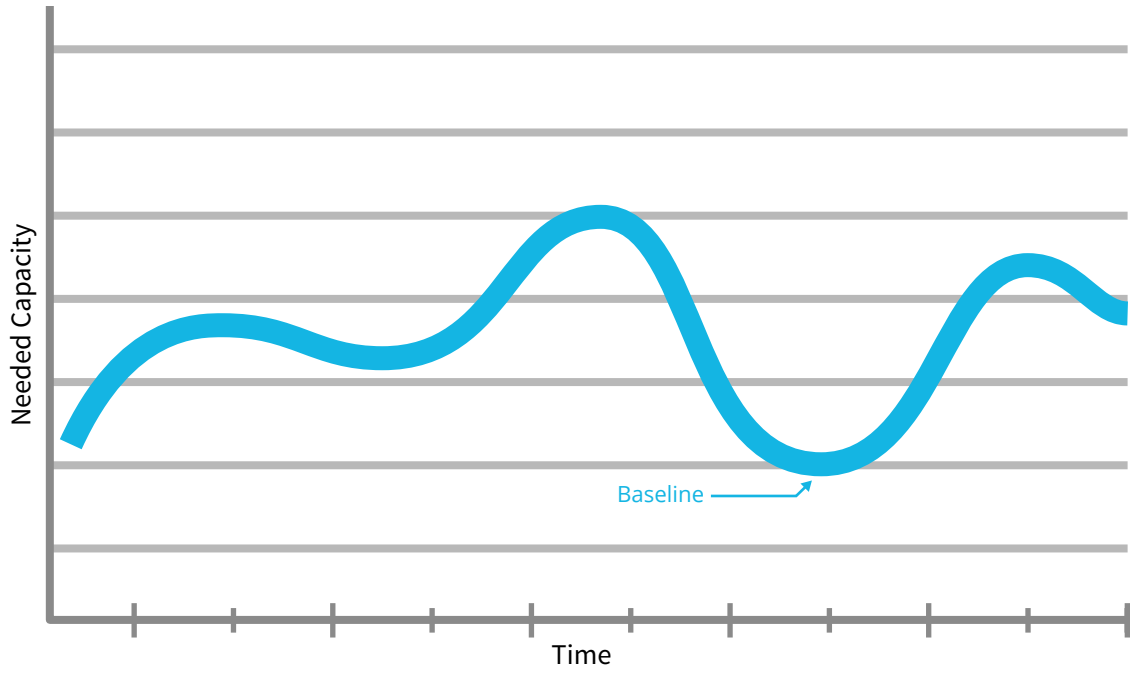

Now that we've defined steady state capacity, let's segment businesses by phase and see how work patterns dictate which cloud platform is optimal.

Clouds Segmented by Usage Patterns

Cloud environments can be segmented by the phase of your business.

Small Workloads with Unestablished Usage Patterns (less than 8 cores of CPU, 1 TB of storage, 1TB of aggregate bandwidth per month)

Startups often have small computing and storage workloads when they start. They're typically characterized by unpredictable traffic and system load, since the business concept itself is being market tested. Markets and business execution often determine how much usage applications will get, thus public clouds are often good environments for startups since infrastructure can be resized and expanded very rapidly.

Since the infrastructure footprints for startups are typically small, it does not make sense to commit to single-tenant or dedicated infrastructure at this phase. It makes more sense to use small virtual machines, metered storage, and metered bandwidth to keep costs at a minimum. At this scale, the high retail price consumers pay for public clouds like AWS, Azure, and GCE typically do not hurt the startup since there simply isn't that much infrastructure to pay for.

Moderate Workloads with Highly Variable Usage (more than 24 virtual cores, 2TB of storage, 3TB aggregate bandwidth per month)

Young businesses that are either revenue positive or profitable typically have an established baseline for compute, storage, and network utilization. But some are completely variable.

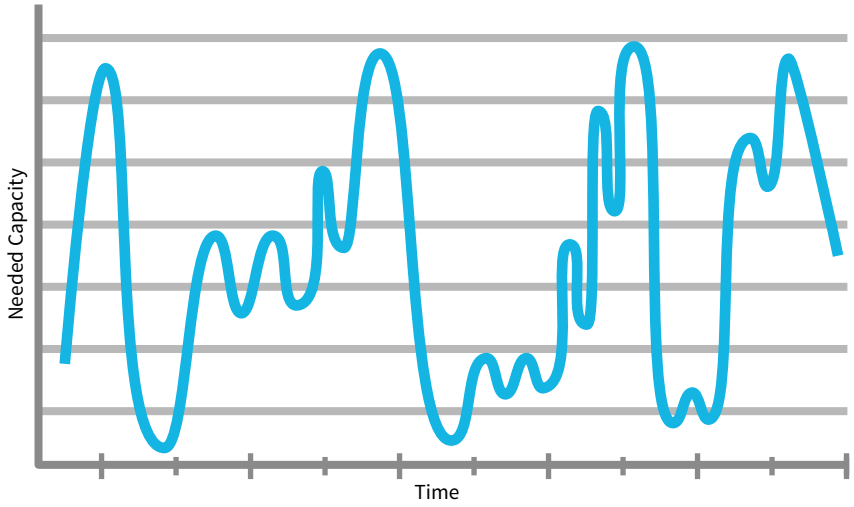

An example would be a software business that models trading strategies for high frequency traders. Their models are complex and their computing requirements for backtesting are highly variable - they could need 2 virtual CPUs and 1TB of storage one day, and 100 virtual CPUs and 50TB the next day. They could go with **no requirement at all for several days.** Here there would be no meaningful baseline, but there is an average capacity that could be used to inform a private cloud build.

If your business looks like this, you want to compute the monthly average for your cloud costs over a 6-12 month period. If your cloud rent is more than $15,000/mo, you have a moderate workload despite minimum steady-stage capacity being completely variable.

For these types of workloads, there are a few options.

1. You can purchase a dedicated cloud infrastructure and **stop paying rent forever**. In this model, you would buy the infrastructure footprint to handle the average. In a financed model on a 24 month term, at month 25 the core cost of the cloud would drop out, and monthly costs would converge to $0.00. Any workloads requiring more than the dedicated footprint would be bursted onto public cloud infrastructure.
2. You can rent forever on a public cloud. The upside is there is less capacity planning up front since the public cloud is entirely elastic and shared. The downside is you will be paying significantly more, and per-VM performance will be less than a dedicated cloud since it is a shared environment.

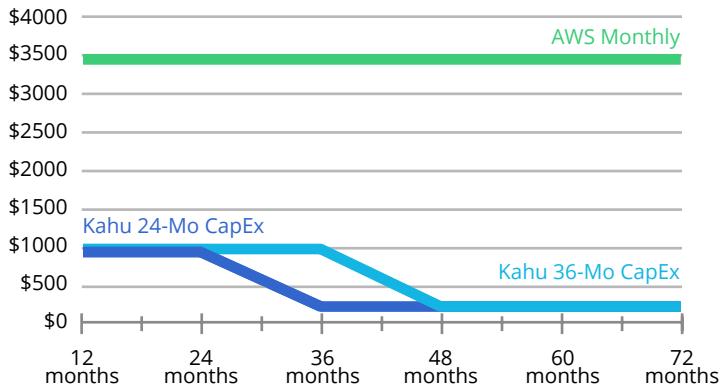

This graphic outlines the cost converging to zero on a capitalized Kahu cloud. Note the significant cost delta between AWS and Kahu. Specs in this comparison: 48vCPU, 256 GB RAM, 3.2 TB storage.

Moderate Workloads with Established Usage (more than 24 virtual cores, 2TB of storage, 3TB aggregate bandwidth per month)

Another usage pattern is an established pattern, with bursts during specific seasons.

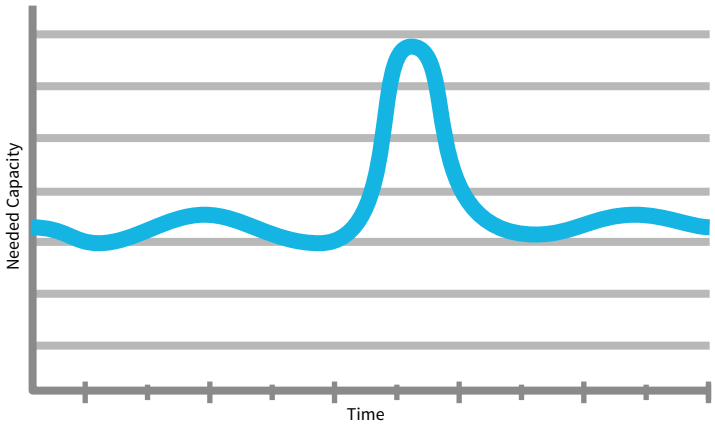

An example would be an e-commerce store that sells mattresses. They might require a baseline of 192 virtual CPU and 6.4TB of storage for the majority of the year, but on Christmas and Black Friday they require 384 virtual CPU and 9.6 TB of storage to accommodate the surge of transactions. They have a defined steady-state capacity in this case.

They may experience traffic bursts across their platform depending on usage patterns unique to their business (i.e. seasonality might result in sudden spikes in traffic for an e-commerce platform, or a healthcare SaaS might have almost no sudden bursts in traffic as seats are sold in a linear fashion), but overall growth tends to be linear for a modestly growing business.

If your business looks like this, **you want to stop paying rent on your steady-state capacity.**

For these types of workloads, there are a few options.

1. You can buy a dedicated cloud for your steady-state capacity and burst seasonal traffic onto the public cloud. This is the most cost-effective approach even if it requires al little capacity planning.
2. You can rent forever on a public cloud. This requires less capacity planning and will end up costing the business significantly more money in the long run.
3. You can over-provision a dedicated cloud on a mixed capex/opex basis. In this scenario, steady-state capacity is serviced by a purchased cloud for maximum cost efficiency, and surges are handled by leased dedicated nodes that augment the cloud in preparation for traffic bursts and other auxiliary workloads.

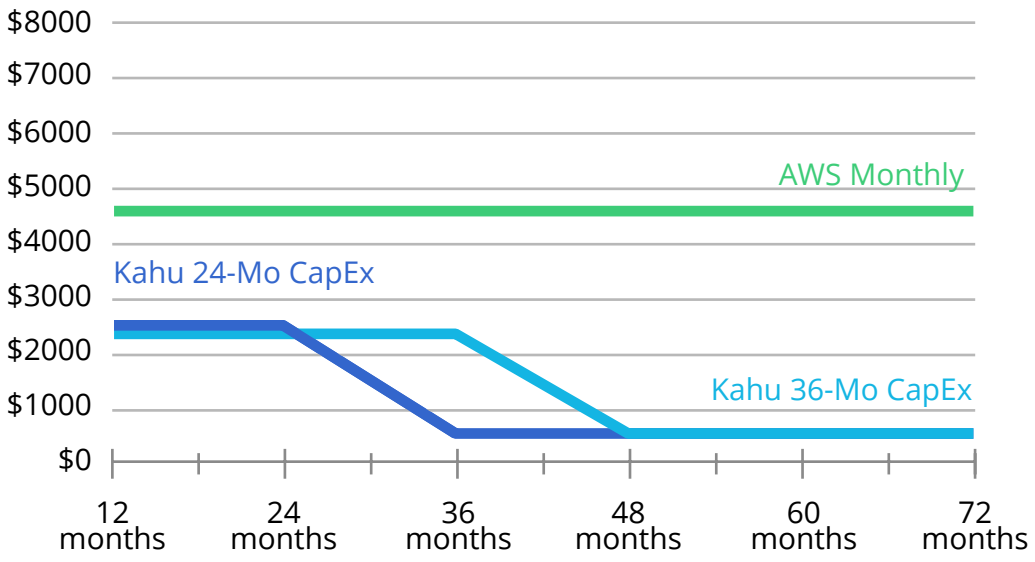

This graphic outlines the cost converging to zero on a capitalized Kahu cloud. Note the significant cost delta between AWS and Kahu. Specs in this comparison: 2 Kahu K4 Models compared with AWS m3.xlarge instances and 3.2 TB of EBS.

Heavy Workloads with Highly Variable Usage (more than 50 virtual cores, 10TB of storage, 4TB aggregate bandwidth per month)

In some cases, businesses may have very significant compute and storage requirements, but have wildly variable usage patterns. An example might be an advertising network that is constantly shrinking and expanding its cloud footprint to accommodate variable advertising click-through across its platform. They may have a minimum of 50 virtual cores and 10TB of storage, but could scale up to as many as 400 virtual cores and 100TB of storage, regardless of seasonality.

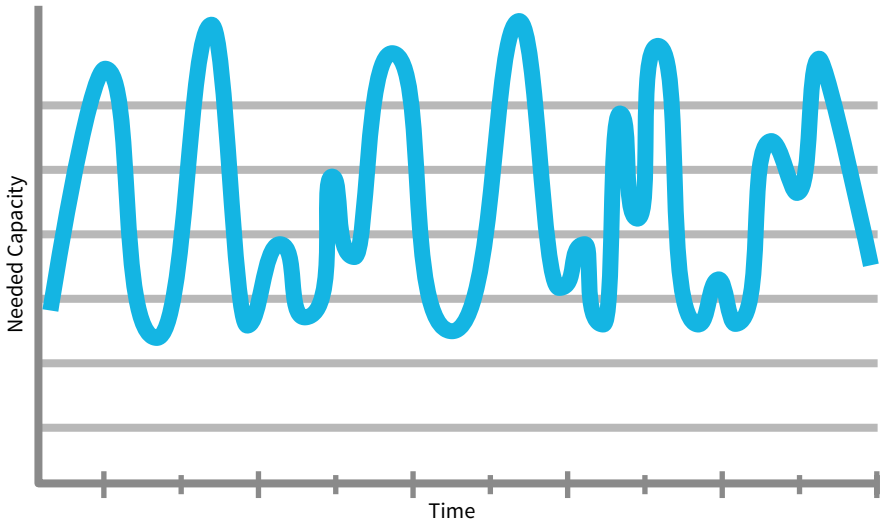

In this case, the business needs to juxtapose the cost of capacity planning overhead and on-going rent on a public cloud.

1. On a public cloud, this business would rarely have to capacity plan and would simply consume/release resources as necessary. However, it would be exceedingly expensive to pay for this much infrastructure on a rental basis forever.
2. On a dedicated cloud on a cap ex basis, this business would drive the majority of their cloud costs out of their platform. However, it would require a reasonable amount of capacity planning to nail down the appropriate steady-state capacity.

Heavy Workloads with Established Usage (more than 50 virtual cores, 10 TB of storage, 4TB aggregate bandwidth per month)

Heavy workloads, whether they require a significant amount of CPU capacity or a ton of Storage, are generally associated with established businesses and enterprises. There are many businesses that do less than $6M/yr (i.e. a small business) that still have very significant compute and storage capacity requirements. Unfortunately, public clouds tend not to be the best solution for these environments if they have an established workload.

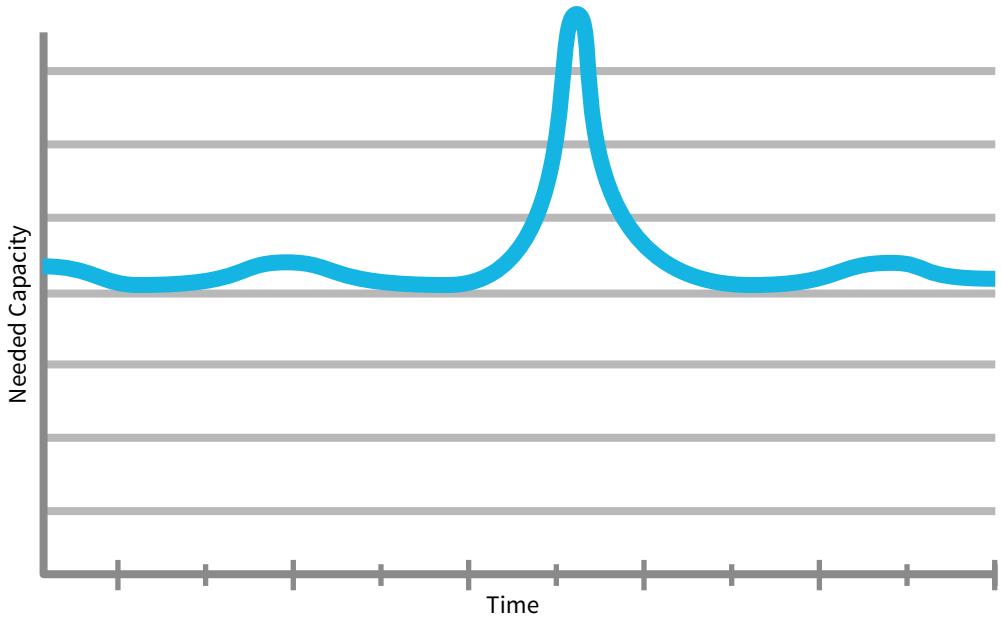

Since heavy workloads are inherently in need of many CPU cores and storage volumes, running them on public clouds at retail prices tends to be exceedingly expensive. Add high availability for underlying storage to this mix and costs double, especially for storage like EBS on AWS.

In this case, unless management simply has no desire to ever see a capacity planning sheet, it makes no financial sense to be in a public cloud. As an example, here is a comparison of a dedicated private cloud on a cap ex basis vs. AWS and Azure.

Click here to see the comparisons with Microsoft Azure and Google Cloud as well.

Dense Storage

Kahu Model	Role	vCPU	RAMS	torage (GB)
K3	Dense Storage	48	256	19200

	12 months	24 months	36 months	48 months	60 months	72 months
Kahu 24-Mo CapEx	$2,874.66	$2,874.66	$500.00	$500.00	$500.00	$500.00
Kahu 36-Mo CapEx	$2,686.39	$2,686.39	$2,686.39	$500.00	$500.00	$500.00
AWS Monthly	$6,986.30	$6,986.30	$6,986.30	$6,986.30	$6,986.30	$6,986.30

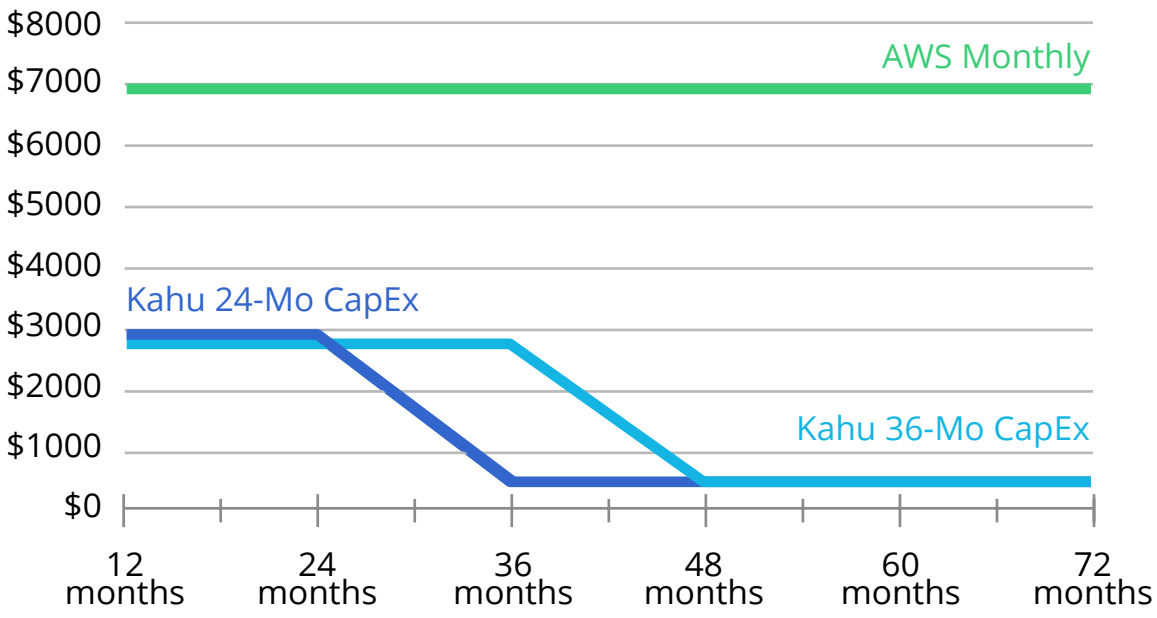

Selecting a Cloud Provider

There are many variables to consider when selecting a cloud provider - the top 2 for any CTO/CEO should be:

Cost

The cloud promised to bring efficiencies to the modern business. This isn't just operational efficiency - it's meant also to provide cost efficiencies. If your current cloud platform costs significantly more than alternatives, your cloud roadmap needs reconsideration.

Performance

Elasticity and operational efficiency benefits of the cloud should not trump performance. Study after study shows that high performing web applications retain and gain more users than their slower counterparts. Do not sacrifice performance when selecting a cloud architecture to save a few bucks.

Selecting the Right Modality

Selecting the right modality for your workload is critical for long-term cost savings and application performance. The models we exhibited in this guide elaborate on what we believe to be 80% of the use cases out there - however, no two businesses are alike and it is critical that your unique requirements are mapped to the correct cloud modality.

Lightcrest can assist you in ensuring that you're getting the best possible performance for your business applications at the lowest possible cost. Selecting the wrong cloud model could nearly **double your cloud costs on a year-on-year basis.**

Talk to our cloud experts to ensure you select the right cloud model for your business.

Request More Information

lightcrest