Gartner Research

Enhance Pricing Models to Address 3 Unpredictable Use Situations

By Ron Burns, Jeff Chamberlain, Laurie Wurster, John Santoro

29 September 2021

Gartner

Enhance Pricing Models to Address 3 Unpredictable Use Situations

Published 29 September 2021 - ID G00752943 - 11 min read

By Analyst(s): Ron Burns, Jeff Chamberlain, Laurie Wurster, John Santoro

Initiatives: Product Introduction; Foundations of Product Introduction

Pricing models based on the number of users or transaction volumes often lack flexibility that allows customers to manage spikes in usage. Product managers need to provide transparent options for customers to handle unexpected use and to encourage shared value growth.

Overview

Key Challenges

- Without policies to handle usage tier overages, large unexpected bills can negatively affect customer satisfaction and result in missed growth opportunities.
- Spikes in usage or unexpected use situations, compounded by a lack of pre-established paths or flexible options to handle them, expose cost risks for both customers and providers.
- Large product portfolios encompassing numerous and varied pricing models can limit wallet-share growth when it complicates a buyer's ability to predict individual product usage and budget accordingly.

Recommendations

Product managers forming product introduction plans and looking to enhance pricing models to handle unexpected use, should:

- Define a customer transparent tier strategy by including overage policies, forgiveness terms and the flexibility to move between tiers within a contract period.
- Create paths for customers to follow when unexpected use situations occur by implementing flexible pricing options, such as burst pricing, spot consumption and reserved capacity pricing.

- Reduce complexity, and allow customers greater flexibility and usage variation across products by creating a common denominator pricing model across product portfolios, such as a credit model.

Analysis

Many modern pricing models are based on metrics that are perceived to align with or predict a desired business outcome, benefiting customers' value expectations while allowing providers growth potential. These metrics can be measures of size (for example, number of end users) or volume-driven on the basis of data transactions or data throughput, and are commonly purchased in tiers and via yearly subscriptions.

However, these models often lack the flexibility to manage spikes in usage or temporary needs for additional capacity. Without policies or options to handle unpredictable use situations, technology providers could upset customers with large unexpected bills, which can translate into customer churn and missed growth opportunities.

Pricing model options that allow customers to flexibly accommodate unpredictable use situations will maximize product profitability.

Product managers need to assess the most common situations customers face with unexpected use and the pricing model options and policies that can efficiently meet these situations (see Figure 1). These pricing model options can maximize product profitability (for example, extend product use to new use cases or expedite renewals at a higher tier), in addition to the benefits of customer satisfaction.

This research covers three types of situations in which a customer may unpredictably exceed the usage limit stipulated in their contracts, and provides three examples of pricing solutions for each. These situations are not mutually exclusive, and some product managers may need to accommodate all three to maximize product use and profit opportunities.

Figure 1. Examples of Unexpected Use Situations and Pricing Options

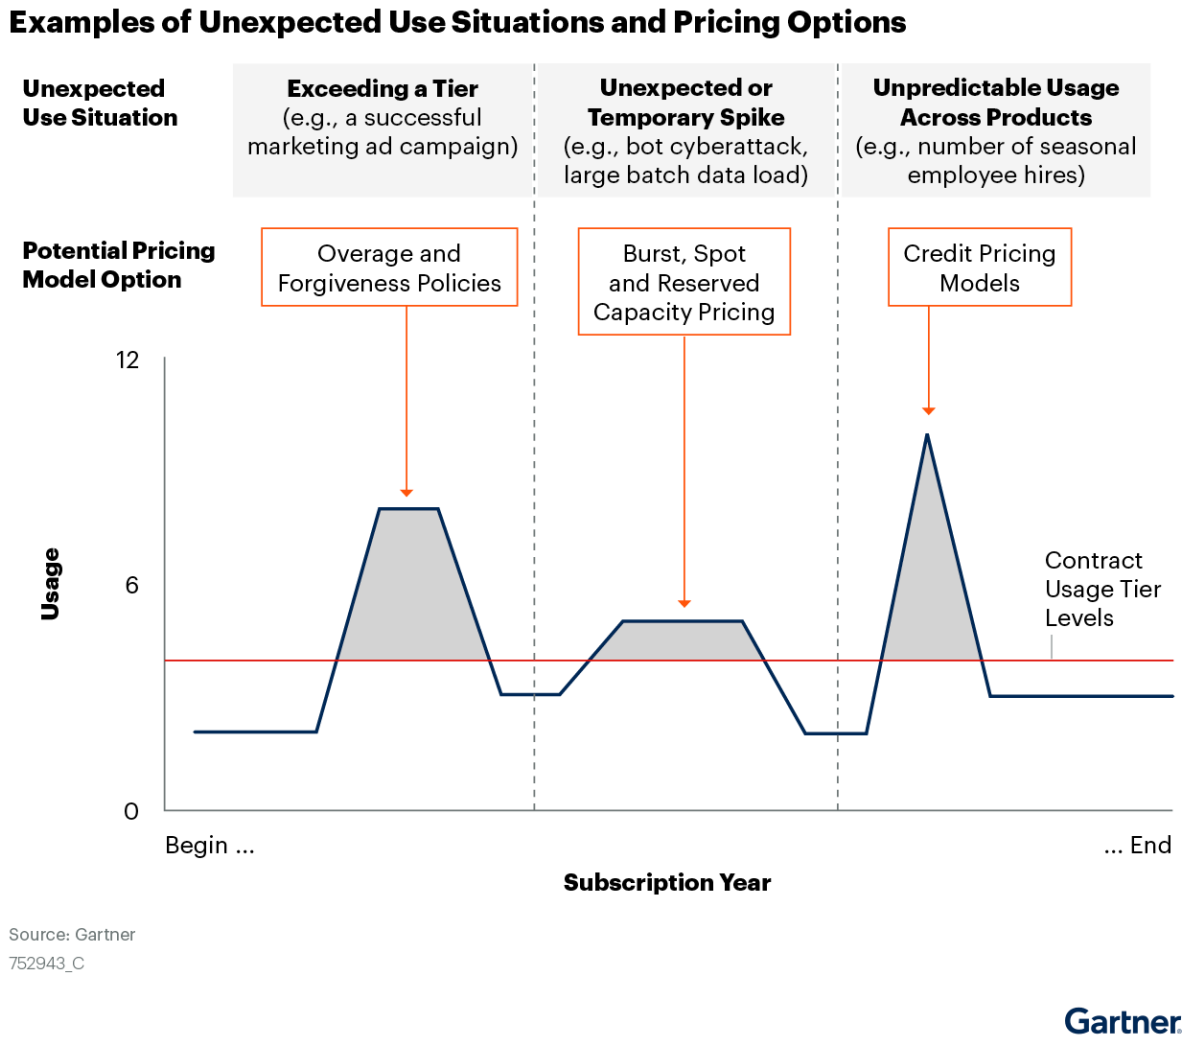

Source: Gartner
752943_C

Gartner

Exceeding a Usage Tier

Tiered pricing models by definition have a scope and ceilings. For example, a buyer of a cloud database subscribes to the 1TB, 10TB or 100TB tier (relating to data at-rest storage volumes). However, there are some controversial market perceptions with tier-based models, overage policies, and approaches for moving customers to higher tiers:

- **Vague definition of tier metrics.** If the pricing metric is not something used in the customer's business language or measurable, then predictability becomes a concern.
- **Unlimited use.** This may seem like an attractive marketing term, but the variable cost management can be a perpetual risk to the provider.

- **Implied per-unit volume discount at higher-level tiers.** Providers need to analyze whether customers have control over increasing usage or if volume discounts would motivate a customer to use more. Tiered pricing can encourage usage growth but can also cause friction and mask poor predictability for customers that unintentionally grow into higher tiers.
- **Excessive penalties for exceeding a tier limit.** This can increase short-term revenue but damage long-term annual recurring revenue (ARR) if perceived as price gouging.

Product managers need to proactively address these perception and operating risks. Clear and fair policies will encourage customers to accurately estimate usage, or even pay for the "insurance" of a higher tier to allow room to do more. Transparency is key with overage policies, and surprises should be minimized. Product managers should plan for making product metrics customer-accessible with configurable administrative dashboards, usage reports, and notifications/alerts that can be predictive of overages. Customer success teams can proactively reach out to give guidance on how to optimize use, or trigger upsell and cross-sell campaigns.

Product managers should work with sales and revenue operations teams to create flexible terms that allow easy movement from one tier to another during a contract period to support changes in customer needs. This flexibility can often be a competitive differentiator. Tiers can be used as a means to enter into enterprise unlimited use license negotiations. Since it's challenging to predict an upper bound for usage growth, language around reasonable use and contract length can provide some protection until customer relationships mature to a point when more specific levels and terms can be set with confidence.

Examples of tier overage policies:

- **VMware Tanzu Observability** [1] — Customers commit to a tier of telemetry data points per second (PPS), with each tier having a set pricing rate. Spikes or a need for additional capacity above the committed tier during the contract period are handled via a flexible billing period model. For example, assume a commit at Tier 2 with a ceiling of 35,000 PPS. Ten days into a 30-day billing cycle, the customer uses 30,000 PPS and requests to add 10,000 PPS pushing them to Tier 3. In this scenario, the 30-day billing would be split into two billing periods, with the first billing period as 10 days at a 30,000 PPS Tier 2 rate, and the second billing period as 20 days prorated at the Tier 3 rates. This approach captures expected revenue at the foundational pricing tier, while accommodating incremental customer usage at a predefined price point that creates a mutual win for the provider and customer.

- **Google Cloud Storage** [2] — Customers commit to a fixed amount of data storage spend (per month for 12 months) and are allowed to grow with no extra charges for usage over the commitment during those 12 months. The yearly renewal provides a 30% forgiveness threshold option if the customer commits to the next 12 months at the previous peak usage. If that peak usage is within 30% of the original commitment, all of the customer's previous year's overage is free. If usage was more than 30%, the customer would repay the remainder over the next year. This type of forgiveness policy can encourage revenue growth (an upgraded tier) with a specific customer.

Unexpected or Temporary Spikes

Sudden unexpected demands on technology are commonplace, and can be caused by customer business variability/seasonality, cyberattacks, customer workforce adjustments, or providers' marketing campaigns. Various policies and pricing plans can provide attractive paths to handle unexpected use or seasonal spikes, such as:

- **Burst pricing** — Designed to address sudden and relatively brief or unexpected bursts in demand by providing per-unit pricing for specific overage volumes.
- **Spot consumption** — Intended to offer significant savings over standard consumption if the customer is willing to operate in an unpredictable environment. In short, this model provides lower unit-based pricing for the usage of forecast excess capacity. The customer risk is that this model lacks guaranteed availability and includes the risk of instances being reclaimed to support demand variability within traditional pricing tiers. Pricing can be similar to an auction or strictly dependent on resource availability.
- **Reserved capacity** — This approach reserves a fixed inventory for customers — whether used at or below the reserved capacity threshold. It's similar to a savings plan or insurance that can enable lower per-unit costs for unpredictable events in exchange for customer commitment to use a specific tier of volume, instead of making commitments to specific instance configurations or time frames.

Three of the leading cloud service providers — Amazon, Google, and Microsoft — all offer variations of burst, spot or reserved capacity pricing. [3] Product managers should note how these providers use the functionality, policies and price points of these pricing models to differentiate versus their competitors. Specific and other examples include:

- **Blue Prism** [4] — This burst pricing model, which is new for the robotic process automation (RPA) market, is designed to help customers respond to spikes in demand for intelligent automation beyond their digital worker license agreements. Customers can reserve temporary burst digital workers instead of buying additional licenses when they need to satisfy seasonal or unexpected business demands.
- **Azure Batch** [5] — Spot pricing is offered via "low-priority virtual machines (VMs)" to reduce the cost of batch or asynchronous processing workloads that can take advantage of distributed surplus capacity in Azure. The trade-off is that these VMs may not always be available or may be preempted at any time. Low-priority VMs are offered at a significantly reduced price compared with dedicated VMs — for example $0.007 per hour versus $0.036 per hour for normal VMs.
- **AWS Savings Plans** [6] — A reserved capacity option enables lower per-unit costs in exchange for customers' commitment to use a specific amount of compute services, instead of making commitments to specific instances or configurations.

Unpredictable Usage Across Products

Credit (or token) pricing models require customers to commit to a usage level for the licensed term (typically per year) by purchasing a set number of credits that can be applied across the use of several products in providers' portfolios. Overages are not a factor, and some providers offer a rollover of unused amounts up to contract anniversary. Scenarios that may make credit models advantageous include:

- Large product and service portfolios that have many diverse pricing models and complicate buyers' ability to buy more from the provider (that is, extending budgeting, evaluation and approval processes).
- Difficult (or impossible) for customers to predict usage demands across multiple product and service options, and thus, there is a need to allocate budget flexibly.
- Desire to incentivize customer "stickiness" by accommodating cloud migration credits and bringing your own license, or securing budget for use on training and professional services.

Consensus on the success of credit models is still emerging, and a counterpoint is that these models may discourage usage, lead to hoarding behavior, or cause customers with unused credits to demand a smaller renewal. Prepaid credits should free customers from having to preemptively select potentially needed products/services versus instinctively leveraging the right product/service for the task at hand. Credit models have been rolled out by a number of technology providers, and examples include:

- **Sumo Logic** [7] — Cloud Flex Credits negate spike or increased usage penalties and the need to provision for peaks. Credits can be used to flexibly pay against use of any products such as logs, metrics or storage. A management console reflects real-time utilization by product, and each product tracks to a predefined number of credits per event (for example, 20 credits per gigabyte ingested for search). Credits are licensed annually, and customers choose how best to configure the platform and optimize ingest and use-case patterns.
- **Oracle Cloud** [8] — Customers commit to an amount of Oracle Annual Universal Credits that can be applied toward the future usage of infrastructure as a service (IaaS) and platform as a service (PaaS) cloud services. Savings comes from "combining cost reduction and a predictable monthly spend with a ramp up period as you onboard your workloads," and the flexibility of starting and stopping services similar to a pay-as-you-go model. Credits are purchased via a 12-month minimum increment in advance. Credits are debited monthly based on the customer's actual usage and must be used within the applicable 12-month period.
- **Snowflake** [9] — A credit is a consumption measure, such as when a virtual warehouse is running, the cloud service layer is performing work, or serverless features are used. For example, credits are charged based on the number of virtual warehouses a customer uses, how long they run, and their size. If credits are not used in a given time frame, they can be rolled over if the customer purchases the same amount, or greater, in the next time period.

Evidence

The evidence for this research was derived from inquiries with Gartner clients, research by the Gartner Secondary Research Services team and primary research conducted by the authors.

[1] VMware Tanzu Observability by Wavefront Service Description, VMware.

[2] How Does Your Cloud Storage Grow? With a Scalable Plan and a Price Drop, Google Cloud.

[3] IaaS Cloud Contracting: Comparing Amazon Web Services, Microsoft Azure and Google Cloud Platform.

[4] Blue Prism Introduces Consumption Based Pricing for Automation, PR Newswire.

[5] Use Low-Priority VMs With Batch, Microsoft.

[6] Savings Plans, Amazon Web Services (AWS).

[7] Pricing, Sumo Logic.

[8] Frequently Asked Questions: Oracle Universal Credit Pricing, Oracle.

[9] Understanding Snowflake Virtual Warehouse, Storage, and Cloud Services Usage, Snowflake.

Recommended by the Authors

Some documents may not be available as part of your current Gartner subscription.

Enable Growth With a Tiered Value-Based Pricing Model

How to Understand Cloud DBMS Pricing Models

How Product Leaders Should Select the Right Software Pricing Model

Actionable, objective insight

Explore these additional complimentary resources and tools for tech product management leaders:

eBook

3 Lessons From High-Growth Companies to Build a Successful Product Strategy

Avoid common product strategy pitfalls.

Download eBook

Research

Top Trends for Tech Providers for 2022

Learn the top trends impacting tech providers.

Download Research

eBook

2022 Leadership Vision for Technology Product Managers

Learn the top 3 strategic priorities for tech product managers.

Download eBook

Tool

Gartner Product Decisions

Inform your product strategy and roadmap.

Learn More

Already a client?
Get access to even more resources in your client portal. Log In

Connect With Us

Get actionable, objective insight to deliver on your mission-critical priorities. Our expert guidance and tools enable faster, smarter decisions and stronger performance. Contact us to become a client:

U.S.: 1 844 466 7915

International: +44 (0) 3330 603 501

Become a Client

Learn more about Gartner for Product Teams
gartner.com/en/industries/high-tech

Stay connected to the latest insights

Gartner®