NetApp Instaclustr

White Paper

Cadence® vs. Temporal: UNDERSTANDING WORKFLOW ORCHESTRATION AND TEMPORAL CLOUD PRICING

Abstract

This white paper provides a detailed cost comparison between the subscription-based model of NetApp Instaclustr for Cadence® and Temporal Cloud's usage-based pricing, hoping to guide decision-makers in selecting a cost-effective and scalable workflow solution that meets their business needs.

Through practical use cases, we demonstrate how Instaclustr's predictable pricing can result in substantial savings, up to ~80% or above when compared to Temporal Cloud. We also explore the core features and development histories of both Cadence and Temporal, highlighting their shared strengths and considerations that should be factored in when ultimately deciding on a solution for your business.

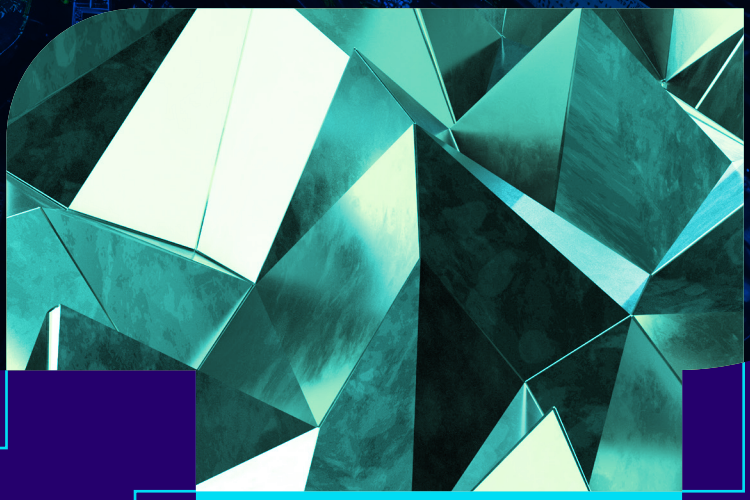

Table of contents

INTRODUCTION

Choosing the right workflow orchestration engine means weighing factors like scalability, fault tolerance, developer productivity, and ultimately cost efficiency. Many solutions simplify creating fault-tolerant business logic, but understanding their cost and pricing models especially at scale can be tricky. That's where NetApp Instaclustr for Cadence® comes in, offering a clear pricing model that can save you ~80% or more when compared to Temporal Cloud.

This white paper breaks down the pricing models of Instaclustr for Cadence and Temporal Cloud, using real examples to show the potential savings available with NetApp Instaclustr. We'll start by providing an overview of Cadence and Temporal, their shared origins, and how they have diverged over time. Then, we'll dive into the specifics of their pricing models, highlighting key differences and their implications for your budget.

We'll also conduct a comparative cost analysis using practical use cases to illustrate how these pricing models impact overall expenses. Finally, we'll explore how these solutions scale with your workload, demonstrating how Instaclustr's model maintains cost efficiency even at higher volumes. By the end of this white paper, you'll have a clear understanding of the differing cloud pricing models available for your workflow orchestration needs.

UNDERSTANDING *Cadence and Temporal*

Cadence is an open source workflow orchestration service designed to enhance scalability, fault tolerance, and developer productivity. It simplifies the creation of business logic and significantly improves operational management. With its built-in **workflow visibility**, Cadence provides your organization with deeper insights and control over workflows, making it an invaluable tool for integrating complex workflows within your distributed services.

In 2016, the Temporal project was forked from Cadence to begin a commercial venture. Since then, Temporal has ceased contributing to the Cadence Project, leading some to mistakenly believe that Cadence has been replaced. However, this is not the case. The Cadence Project has announced a long-term commitment to its development and support, primarily driven by a dedicated team at Uber along with a

strong partnership with NetApp, where we offer Cadence on the NetApp Instaclustr Platform and provide improvements and bug fixes through our dedicated in-house developers.

While Cadence and Temporal have diverged in their development paths, they still share the same core features that made the original project successful. This also makes simple and **easy migrations** to Cadence from Temporal on either self-hosted or cloud a breeze. Both platforms continue to offer robust workflow orchestration capabilities, but each has evolved to cater to its unique users.

Cadence remains a powerful, fully open source solution with a strong community and ongoing support, primarily led by Uber, who uses Cadence to deliver its platform services, but does not have a direct commercial interest in the project. Uber's primary motivation is to enhance and maintain Cadence for its operational needs, efficiency, and robustness, rather than for profit. This non-commercial stance reduces the likelihood of sudden changes in licensing or business models that could impact end users. **Recent efforts** to make the project more community-led, further reinforce this stability and commitment to open source principles.

In contrast, Temporal is developed with a direct commercial interest. As a for-profit venture this increases the chance of license changes or the introduction of new pricing models which we've seen with **similar vendors in the past.**

Overall, the choice between Cadence and Temporal will hinge on factors like the potential for future changes in licensing and business models as well as predictable costs. Organizations seeking a reliable, community-driven, and open source solution will find Cadence to be the better fit.

Pricing model DIFFERENCES

Understanding cloud pricing models, especially when it comes to managed services, can be tricky but is crucial for making informed decisions that meet both your budget and operational needs. NetApp Instaclustr for Cadence and Temporal Cloud offer vastly different approaches to pricing, each with its own advantages and considerations.

Instaclustr employs a subscription-based pricing model that leverages the cost of underlying compute resources, providing predictability and control over expenses. This model is particularly beneficial for consistent usage patterns and allows for fine-grained cost management and **continuous improvements in cost performance** as well as the freedom to **scale at any time.** Additionally, NetApp Instaclustr supports the "Run In Your Own Account" model, which enables you to run the resources in your cloud provider account. This approach maximizes potential discounts by utilizing your existing savings plans.

In contrast, Temporal Cloud adopts a usage-based pricing model, where charges are determined by the actual usage, metered through specific operations known as "actions." This model can offer flexibility for varying workloads but may introduce variability in costs. The operations below on Temporal Cloud count as "actions":

- Starting, resetting, updating, or exporting a workflow
- Starting a timer
- Requesting a search attribute 'upsert'
- Sending a signal
- Receiving a query
- Recording a side effect
- Spawning or executing a child workflow
- Starting, retrying, or heartbeat of an activity

These actions can accumulate significantly over the course of a month, especially when operating at scale, and can vary greatly with changing demand. In addition to action-based pricing, Temporal Cloud also includes additional costs for storage and support, which we will explore further in our comparison.

COMPARATIVE *cost analysis*

A practical approach to contrasting the strengths and weaknesses of the two pricing models is to examine their differences through a specific use case. Since Cadence does not provide metrics that directly correlate to the "actions" used by Temporal for pricing, we will need to approximate these metrics. This can be achieved by using the "requests per second" (requests/s) across Cadence's three internal services: Front End, History, and Matching. These services collectively encompass all, plus more, of the actions that Temporal uses for its pricing model, giving us a good basis for comparison.

For this example, we've chosen to outline a modest deployment and one that fits into Temporal Clouds publicly available pricing data as volumes above 250M actions are handled on a per request basis.

- 100M actions/month
- 50GB "Active Storage" over 1 month
- 1TB of "Retained Storage" over 1 month
- A support package

Temporal has a sliding pricing scale which awards discounts on per 1 million actions on a stacking basis as your usage increases. For our given example, the table below outlines how these charges would be applied.

The Temporal Cloud pricing model categorizes customers into specific plans. For this example, the outlined usage would place us in the Business Plan category that has a minimum cost of $500 a month or 10% of monthly usage, whichever is greater. To support the use case outlined above, the cost for the plan would be **$500** which includes 2.5M actions, 2.5GB Active Storage and 100GB of Retained Storage and business hours support only.

Actions (millions)	**Rate per 1M actions**	**Cost ($USD)**	**Total/month**
2.5M	NA	$698.75	$698.75
First 5M	$50.00	$250.00	$948.75
Next 5M, up to 10M	$45.00	$225.00	$1,173.75
Next 10M, up to 20M	$40.00	$400.00	$1,573.75
Next 30M, up to 50M	$35.00	$1,050.00	$2,623.75
Next 50M, up to 100M	$30.00	$1,500.00	**$4,123.75**

In addition to the above costs, there are additional storage costs to consider when deploying on the Temporal Cloud platform. For our given example, storage is charged in per hour units and would be applied like the below when accounting for our Business Plan inclusions:

Storage type	Total usage	GBh consumed	Cost per GB hour	Cost	Total/ month
Active	47.5GB	$698.75	$698.75	$1456.35	$1456.35
Passive	900GB	$250.00	$948.75	$689.85	**$2146.20**

Based on the publicly available pricing data provided on Temporal's pricing page, this brings the total cost of ownership for the outlined deployment on Temporal Cloud to a grand total of **$6,769.95 per month** or an average cost per 1M actions of **$67.69**.

As previously explored, the pricing model on the Instaclustr Platform differs from Temporal's and when working through pricing on our platform it's important to get an idea of how much momentary capacity you'll need to service your application. For the outlined 100M actions per month, we could expect approximately 38 actions/s across your cluster. Accounting for workloads that do not occur in a steady stream, which is common, we can assume that 50% of monthly actions occur within the busiest 8 hours of each day, meaning approximately 57 Actions/s would be peak load during those hours.

$$\text{Actions per Second} = \frac{50\text{MActions}}{8\text{ hours} \times 3600\text{ seconds} \times 30\text{ days}} \approx 57\text{ (Actions per Second)}$$

You may see traffic like this for user driven workflows that have daily trends to them, for example 2 hours in the morning, 2 hours after midday, and 4 hours in the evening.

As this is a mission critical production deployment, we've chosen our largest recommended production configuration using our **Light Production Workloads** offering. This model uses a set of shared Cadence dependency clusters. The deployment model gives you full access to the extensive features of Cadence and is ready for your mission critical workloads as it comes backed by our full production SLA. The charges for this are as follows:

Nodes	$/node/month	Total/month
2 x CAD-SI-PRD-m7g.large-50	$718.02	$1,436.04

Our shared infrastructure offering includes a reasonable amount of "Active" storage which can cover the required 50GB of usage from our example. For additional archival storage you can leverage the **Cadence Archival** feature which utilizes a Bring-Your-Own (BYO) S3 bucket, the costs for this would be the same rate as what Amazon charges for its S3 service at 0.023 per GB, accounting for a total additional cost of $23 a month for 1TB. This size deployment can also serve approximately 160 actions/s, well above the peak workload of ~57 actions/s, allowing for additional bursts when required.

The total cost for a fully managed Cadence deployment on the Instaclustr Platform servicing the same requirements as on Temporal Cloud is **$1,436.04** representing a **saving of 78%**. This price also includes 24x7 production support and access to our world-class Cadence engineers and developers.

If you were to maximize the usage of your above subscription to NetApp Instaclustr to the full ~414M actions per month that would bring your total cost per 1M actions to **$3.46**. If your workload does increase, and you find yourself needing significantly higher capacity in your cadence workloads, then Instaclustr can help you increase the size of your deployment, as we have seen on some of the node sizes we currently support within our fleet the $/1M actions figure to drop below $1.

OPERATING *at scale*

Moving away from the smaller scale example using our shared infrastructure offering we thought it would be important to demonstrate how our instance-based pricing model scales with your workload, maintaining an 80% or lower cost reduction when compared to Temporal.

Technology	Cluster details	Cost ($USD)	Total/month
Cadence®	3 x CAD-PRD-m7g.large-50	$1,344.06	$1,344.06
Cassandra®	6 x CAS-PRD-r7g.xlarge-400	$2,378.34	$3722.40
OpenSearch®	3 x SRH-PRD-r6g.large-800	$1,178.28	$4900.68
Kafka®	3 x KFK-PRD-r6g.large-400	$1,095.55	$5996.23

We expect this size deployment to be able to service ~1500 actions per second per node at 70% utilization which we consider a healthy margin. This would equate to a total of 11.6B actions per month. At this usage level you can expect to pay **$0.57 per 1M actions**.

As previously explored, for the shared infrastructure example we understand that workloads aren't always perfectly predictable and will also sometimes have hourly or even weekly variation to them. For our example above, demonstrating the value NetApp can deliver with our Managed Cadence offering, we'll work through some additional workloads.

With our large-scale production deployment being able to service 11.6B actions per month or (~4500 actions/s) we'll assume you're utilizing 65% of that total capacity (7.7B actions per month).

$$\text{Actions per Second} = \frac{3.85\text{BActions}}{8\text{ hours} \times 3600\text{ seconds} \times 30\text{ days}} \approx 4{,}456\text{ (Actions per Second)}$$

In the example above, half of your entire workload might occur over 8 hours each day, and we'd be able to offer you a value of **$0.77 per 1M actions.**

As another possibly more unrealistic, but none the less valid example, we can demonstrate the value of the subscription model NetApp Instaclustr offers over the previously worked example of Temporal Cloud's pricing by using the same plan to run 100M actions in the same 8-hour period. Doing this, we'd be able to offer you a value of **$59.96 per 1M actions, representing a cost reduction over Temporal of 11%** whilst still being provisioned to handle a load spike 116 times greater per second

Service	Actions run	Action capacity	Cost per 1M actions
Instaclustr	100M	11.6B	$59.96
Temporal Cloud	100M	100M	$67.69

Note: Action Capacity is the maximum number of actions you would be able to run per month on a given plan. Effective cost per 1M actions will reduce as you achieve higher utilization levels.

MAXIMIZING *value*

As outlined by the comparison, there are significant advantages to choosing Instaclustr for Cadence, especially at higher volumes. Along with this we can offer additional savings through a variety of deployment models which our **sales team** are available to chat about at any time.

As an example, taking the above larger scale deployment and running in your own AWS account using the RIYOA model you could expect to pay $4,925.70/month total. This cost is made up of two components, the management fee that you pay to NetApp Instaclustr and the cloud resource cost you pay to AWS which is managed by you and may be subject to any savings plans you have in place further reducing cost. Including the Advanced Visibility feature, which requires supporting Kafka and OpenSearch infrastructure, for the prior example when using the RIYOA model the cost breakdown is as follows:

Actions (millions)	Management cost	AWS cost	Total/month
Cadence + Advanced Visibility	$2,643.75	$2,281.95	$4,925.70

Like the previous example, if half of your workload occurs over 8 hours each day, we'd be able to provide a rate of **$0.63 per 1M actions** when utilizing RIYOA pricing.

CONCLUSION

Choosing the right workflow orchestration engine involves evaluating various factors, including scalability, fault tolerance, and importantly cost efficiency which can be difficult to navigate due to the various pricing models available. Throughout this white paper, we've explored the differences between Cadence and Temporal, their pricing models, and the potential cost savings associated with each.

NetApp Instaclustr for Cadence provides a subscription-based pricing model that offers month to month predictability and fine-tuned control and optimization of your cloud expenses. This model is particularly advantageous for businesses with consistent or known usage patterns, allowing for significant savings. Additionally, the "Run In Your Own Account" model enables organizations to leverage their existing savings plans further maximizing discounts.

On the other hand, Temporal Cloud's usage-based pricing model offers flexibility for scaling. However, this flexibility often comes with higher overall costs, especially as usage increases. The commercial nature of Temporal also introduces potential variables such as future changes in licensing or pricing models that may factor into your decision.

Cadence is supported by a strong community and remains a robust open source solution to your workflow orchestration needs. The community-led approach ensures that Cadence continues to meet your business requirements without the risk of sudden commercial changes.

In summary, if your business can predict its load and prefers a stable, community-driven solution, the subscription-based instance pricing model of NetApp Instaclustr for Cadence offers substantial ongoing savings and operational benefits. By carefully considering your specific needs and usage patterns, you can make an informed decision that optimizes both performance and cost-efficiency for your business.

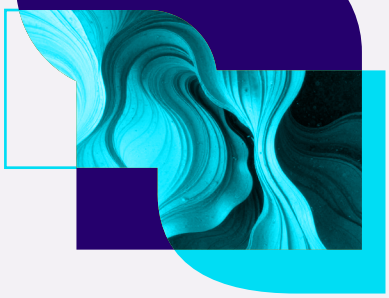

NetApp® Instaclustr specializes in open source technologies for enterprises. Our managed platform streamlines data infrastructure management, backed by experts who ensure ongoing performance, scalability, and optimization. This enables companies to focus on building cutting edge applications at lower costs.

2-17Jul25