

**Review Article**
**Open Access**

## Leveraging AWS SNS/SQS for Real-Time Cache Synchronization in Distributed Systems

**Ananth Majumdar**

USA

**ABSTRACT**

In distributed systems, maintaining consistency across multiple caches presents a significant challenge, often leading to data inconsistencies and degraded application performance. This paper proposes and evaluates a novel approach to cache synchronization using Amazon Web Services (AWS) Simple Notification Service (SNS) and Simple Queue Service (SQS). The proposed system leverages the pub-sub model of SNS combined with the reliable message queuing of SQS to create a scalable and robust cache synchronization mechanism. We present the architecture, implementation details, and performance analysis of this system, demonstrating significant reductions in database load and improvements in response times while ensuring eventual consistency across distributed caches.

Comparative analysis with alternative solutions highlights the advantages of the SNS/SQS approach in terms of scalability, reliability, and ease of implementation. This research contributes to the field of distributed systems by offering a practical, cloud-native solution to the cache synchronization problem, applicable to a wide range of applications requiring consistent data across multiple instances or services.

**\*Corresponding author**

Ananth Majumdar, USA.

**Received:** November 06, 2023; **Accepted:** November 13, 2023; **Published:** November 20, 2023

**Keywords:** Cache Synchronization, Distributed Systems, AWS SNS, AWS SQS, Cloud Computing, Eventual Consistency, Microservices, Scalability

**Introduction**

In the era of cloud computing and microservices, distributed systems have become the backbone of modern software architecture. These systems, while offering unprecedented scalability and resilience, bring their own set of challenges. One of the most persistent issues in distributed computing is maintaining data consistency across multiple instances or services, particularly when it comes to caching.

Caching is a crucial optimization technique used to improve application performance by storing frequently accessed data in memory. However, in a distributed environment where multiple instances of a service are running simultaneously, ensuring that cached data remains consistent across all instances becomes a complex task. When one instance updates its cache, how do we ensure that all other instances are promptly informed and updated? This challenge becomes even more pronounced in systems deployed behind load balancers, where requests from the same client may be routed to different instances over time.

This paper proposes a solution to this problem by leveraging two powerful services provided by Amazon Web Services (AWS): Simple Notification Service (SNS) and Simple Queue Service (SQS). By combining these services, we can create a robust, scalable, and efficient publish-subscribe (pub-sub) system that

keeps caches synchronized across multiple instances in near real-time.

SNS is a fully managed pub-sub messaging service that enables decoupled microservices, distributed systems, and serverless applications to exchange messages. SQS, on the other hand, is a fully managed message queuing service that allows for the decoupling and scaling of microservices, distributed systems, and serverless applications. When used together, these services provide a powerful mechanism for distributing cache update notifications across a distributed system.

In the following sections, we will delve deep into the intricacies of this solution, exploring its architecture, implementation details, performance considerations, and best practices. We will also present a real-world case study and compare this approach with alternative solutions. By the end of this paper, readers will have a comprehensive understanding of how to implement a robust cache synchronization system using AWS SNS and SQS, enabling them to build more efficient and consistent distributed applications.

**Background**
**Distributed Caching in Multi-Instance Environments**

Distributed caching is a technique used to store frequently accessed data in memory across multiple instances of an application. This approach aims to reduce database load and improve response times by serving data from fast, in-memory caches. In a multi-instance environment, such as those commonly found in cloud

deployments, each instance typically maintains its own local cache.

### **Consistency Issues in Distributed Caches**

The primary challenge in distributed caching arises when data is updated. If an update occurs in one instance's cache, other instances may continue to serve stale data from their local caches. This leads to data inconsistency across the system, potentially causing confusion for users or errors in application logic.

Consider an e-commerce platform where product inventory is cached across multiple server instances. If a purchase reduces the inventory count and this update is only reflected in the cache of the instance processing the purchase, other instances may continue to show incorrect inventory levels. This could lead to overselling or lost sales opportunities.

### **Performance Implications of Outdated Caches**

Inconsistent caches not only lead to data accuracy issues but can also negate the performance benefits that caching is intended to provide. If cached data cannot be trusted due to potential inconsistencies, applications may resort to frequent database queries to ensure data accuracy. This increased database load can significantly impact system performance and scalability.

### **Scalability Challenges in Cache Synchronization**

As the number of instances in a distributed system grows, the complexity of keeping all caches synchronized increases exponentially. Direct communication between all instances becomes impractical and does not scale well. Furthermore, in dynamic cloud environments where instances may be added or removed based on load, a flexible and scalable synchronization mechanism is crucial.

## **AWS SNS and SQS: An Overview**

### **Amazon Simple Notification Service (SNS)**

Amazon SNS is a fully managed pub-sub messaging service that enables the decoupling of microservices, distributed systems, and serverless applications [1].

### **Publish-Subscribe Messaging**

SNS implements the pub-sub pattern, where publishers send messages to topics, and subscribers receive messages from topics they are interested in. This decoupling of publishers and subscribers allows for flexible and scalable communication patterns.

### **Topics and Subscriptions**

In SNS, a topic serves as a communication channel. Publishers send messages to topics, and subscribers can subscribe to one or more topics to receive messages. This model allows for one-to-many communication, where a single message published to a topic can be distributed to multiple subscribers.

### **Amazon Simple Queue Service (SQS)**

Amazon SQS is a fully managed message queuing service that

enables the decoupling and scaling of microservices, distributed systems, and serverless applications [2].

### **Distributed Queuing System**

SQS provides a reliable, highly scalable, serverless queue for storing messages as they travel between different parts of a distributed system. It ensures that messages are stored durably until they can be processed.

### **Types of Queues**

#### **SQS offers two types of queues:**

- Standard queues: Provide maximum throughput, best-effort ordering, and at-least-once delivery.
- FIFO queues: Provide strict message ordering and exactly-once processing, but with lower throughput compared to standard queues.

### **Integration between SNS and SQS**

SNS and SQS can be seamlessly integrated, combining the pub-sub model of SNS with the reliable message queuing of SQS. This integration allows messages published to an SNS topic to be delivered to multiple SQS queues, providing a powerful foundation for building distributed systems.

In the context of cache synchronization, this integration enables us to broadcast cache update events to multiple service instances efficiently and reliably.

### **Designing the Cache Synchronization System**

#### **System Architecture**

The proposed cache synchronization system leverages the strengths of both SNS and SQS to create a robust, scalable solution. Here's an overview of the key components:

#### **Publisher Services**

These are the service instances that make updates to their local caches. When an update occurs, the instance publishes a message to an SNS topic.

#### **SNS Topics**


An SNS topic serves as the central point for distributing cache update notifications. All instances that need to be informed about cache updates subscribe to this topic.

#### **SQS Queues**

Each service instance has its own SQS queue subscribed to the SNS topic. This queue acts as a buffer, storing update notifications until the instance can process them [3].

#### **Subscriber Services**

These are the same service instances, but in their role as consumers of update notifications. They poll their respective SQS queues for messages and update their local caches based on the received information.



**Figure 1:** Architecture of cache synchronization system with SNS/SQS and the service instances

### Message Flow

The process of synchronizing caches across instances follows these steps:

#### 1. Cache update Event Publication

When a service instance updates its local cache, it publishes a message to the SNS topic. This message contains information about the updated data, such as the cache key, new value, and a timestamp.

#### 2. Message Distribution via SNS

The SNS topic receives the published message and immediately fans it out to all subscribed SQS queues. This ensures that the update notification is sent to all service instances.


#### 3. Message Queuing in SQS

Each instance's SQS queue receives and stores the message. If an instance is temporarily unavailable or busy, the message is safely stored in the queue until it can be processed.

#### 4. Message Consumption and Cache update

Instances periodically poll their SQS queues for new messages. When a cache update message is received, the instance updates its local cache with the new data from the message.

This architecture ensures that all instances eventually receive all cache updates, providing eventual consistency across the distributed system. The use of SQS queues adds resilience to the system, allowing instances to process updates at their own pace and ensuring no updates are lost due to temporary instance failures or network issues.



**Figure 2:** Sequence Diagram Showing the Message flow for Cache Synchronization

### 5. Performance Considerations

#### Message Delivery Guarantees

SNS provides "at-least-once" delivery, meaning that messages may occasionally be delivered more than once. Design your cache update logic to be idempotent, ensuring that processing the same update multiple times does not cause issues.

## Scalability of the Solution

This SNS/SQS-based solution scales well as the number of instances increases:

1. SNS can fan out messages to a large number of SQS queues with low latency.
2. Each instance only needs to interact with its own SQS queue, avoiding contention.
3. New instances can be easily added by creating a new SQS queue and subscribing it to the SNS topic.

## Latency Analysis

While this system provides eventual consistency, it's important to understand the potential latency involved:

1. SNS to SQS delivery is typically very fast (usually < 100ms).
2. The main source of latency is the polling interval of SQS consumers.
3. Using long polling (with 'WaitTimeSeconds' set) can help reduce both latency and cost.

## Best Practices and Optimizations

### Message Batching

To reduce costs and improve efficiency, consider batching cache updates:

1. On the publishing side, batch multiple updates into a single SNS message when possible.
2. On the consuming side, use the 'MaxNumberOfMessages' parameter in 'receive\_message' calls to process multiple updates at once.

### Dead-Letter Queues

Configure dead-letter queues for your SQS queues to capture messages that fail to process

This allows you to investigate and handle problematic messages without blocking the main processing flow.

### Message Filtering

Use SNS message filtering to allow instances to receive only relevant updates:

This can significantly reduce unnecessary message processing and improve efficiency in large-scale systems.

## Monitoring and Alerting

Implement comprehensive monitoring and alerting:

1. Use CloudWatch metrics to monitor SNS and SQS performance.
2. Set up alarms for queue depth, failed deliveries, and message age.
3. Implement application-level logging and tracing to track cache update propagation.

## Security Considerations

### Access Control with IAM

Use IAM roles and policies to control access to SNS and SQS resources:

1. Ensure instances have minimal required permissions to publish to SNS and read from SQS.
2. Use resource-based policies on SNS topics to control which principals can publish messages.

## Encryption in Transit and at Rest

Enable encryption to protect sensitive data

1. Use HTTPS endpoints for SNS and SQS API calls.
2. Enable server-side encryption for SQS queues to protect messages at rest.

## VPC Endpoints for SNS and SQS

If your instances run in a VPC, consider using VPC endpoints for SNS and SQS:

1. This allows instances to communicate with SNS and SQS without going over the public internet.
2. Enhances security and can reduce data transfer costs.

## Comparison with Alternative Solutions

To provide context for the effectiveness of the SNS/SQS solution, let's compare it with alternative approaches to cache synchronization in distributed systems [4, 5].

Solution	Pros	Cons
<b>Direct Database Queries</b>	<ul style="list-style-type: none"> <li>• Always provides the most up-to-date data</li> <li>• Simple to implement</li> </ul>	<ul style="list-style-type: none"> <li>• High database load, especially for frequently accessed data</li> <li>• Increased latency for each request</li> <li>• Potential for database bottlenecks during high traffic periods</li> </ul>
<b>Distributed Caching Systems (e.g., Redis)</b>	<ul style="list-style-type: none"> <li>• Provides a centralized cache for all instances</li> <li>• Built-in support for complex data structures and operations</li> </ul>	<ul style="list-style-type: none"> <li>• Introduces a new component to maintain and scale</li> <li>• Potential single point of failure if not properly configured for high availability</li> <li>• May require significant network bandwidth for large datasets</li> </ul>
<b>Custom Pub-Sub Implementations</b>	<ul style="list-style-type: none"> <li>• Can be tailored to specific application needs</li> <li>• Potentially lower cost for small-scale applications</li> </ul>	<ul style="list-style-type: none"> <li>• Requires significant development and maintenance effort</li> <li>• May struggle with scalability for large systems</li> <li>• Lack of built-in features like message persistence and dead-letter queues</li> </ul>
<b>AWS SNS/SQS Solution</b>	<ul style="list-style-type: none"> <li>• Highly scalable and reliable</li> <li>• Managed service, reducing operational overhead</li> <li>• Built-in features like message persistence and dead-letter queues</li> <li>• Cost-effective for various scales of operation</li> <li>• Easy integration with other AWS services</li> </ul>	<ul style="list-style-type: none"> <li>• Eventual consistency model (may not be suitable for all use cases)</li> <li>• Requires familiarity with AWS services</li> <li>• Potential for increased complexity in multi-region setups</li> <li>• Vendor lock-in to AWS ecosystem</li> </ul>

**Table 1:** Comparison of SNS/SQS with Alternative Solutions

In comparison, the SNS/SQS solution offers a balance of scalability, reliability, and ease of implementation. It leverages managed AWS services, reducing operational overhead, while providing the flexibility to handle complex distributed systems.

## Conclusion

The implementation of a cache synchronization system using AWS SNS and SQS offers a robust, scalable, and efficient solution to the challenge of maintaining consistent data across distributed systems. By leveraging these managed services, organizations can significantly improve the performance and reliability of their applications while reducing the operational overhead associated with custom synchronization mechanisms.

## Key Takeaways from this Study Include

1. The SNS/SQS system provides eventual consistency with low latency, suitable for a wide range of applications.
2. The solution scales effectively, making it appropriate for both small applications and large-scale distributed systems.
3. Built-in features of SNS and SQS, such as message persistence and dead-letter queues, enhance the system's reliability and error handling capabilities.
4. The pay-per-use pricing model of these AWS services allows for cost-effective implementation, with costs scaling linearly with usage.

As distributed systems continue to grow in complexity and scale, effective cache synchronization becomes increasingly critical. The SNS/SQS approach presented in this paper offers a powerful tool for developers and architects to address this challenge, enabling the creation of more efficient, reliable, and scalable applications.

## References

1. What is Amazon Simple Queue Service - Amazon Simple Queue Service. (n.d.). <https://docs.aws.amazon.com/AWSSimpleQueueService/latest/SQSDeveloperGuide/welcome.html>
2. What is Amazon SNS? - Amazon Simple Notification Service. (n.d.). <https://docs.aws.amazon.com/sns/latest/dg/welcome.html>
3. Subscribing a queue to an Amazon SNS topic using the Amazon SQS console - Amazon Simple Queue Service. (n.d.). <https://docs.aws.amazon.com/AWSSimpleQueueService/latest/SQSDeveloperGuide/sqs-configure-subscribe-queue-sns-topic.html>
4. FitzpatrickBrad (2004) Distributed caching with memcached. *Linux Journal.* <https://doi.org/10.5555/1012889.1012894>
5. Candan K S, Li W, Luo Q, Hsiung W, Agrawal D (2001) Enabling dynamic content caching for database-driven web sites. *SIGMOD '01: Proceedings of the 2001 ACM SIGMOD International Conference on Management of Data.* <https://doi.org/10.1145/375663.375736>.

**Copyright:** ©2023 Ananth Majumdar. This is an open-access article distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.