



Aeron Cluster Standby

Bring improved availability, disaster recovery and load management to your Aeron applications.

With Aeron Cluster, multiple systems and nodes are connected and work together to provide high availability and fault tolerance. However, Aeron Cluster users often have to develop specific business logic to manage snapshots and synchronize data.

Introducing Aeron Cluster Standby

Aeron Cluster Standby provides elevated capabilities of resilience and redundancy to critical systems, ensuring high availability and fault tolerance at all times, hence minimizing the impact of failures or outages on daily operations:

- + Whilst one cluster serves as the primary/active cluster, handling all workloads, another cluster is in standby mode, typically in another availability zone, region or data center, ready to take over operations in case of a failure or planned downtime.
- + By leveraging live cluster streaming and background snapshots, Aeron Cluster Standby can significantly reduce time to recovery and increase overall system availability.
- + Like an active node, each standby cluster node processes every message, ensuring internal state consistency with the primary cluster as quickly as the logic and network allow. Data loss is limited to information that is in transit at the time of failure.

Aeron Cluster Standby enables significant flexibility, allowing configurations that balance bandwidth, costs and resilience. By routing all traffic through a single node, bandwidth costs can be minimized, while having all nodes receive the same data increases resilience. Configurations can be customized to achieve the optimal balance between these objectives.

Aeron Cluster Standby can also be used for manual node rebuilding in the event of a host failure and to offload functionality that could impact latency if running on critical path nodes.

Technical Details

- + Aeron Cluster is based upon a replicated state machine model.
- + Active cluster nodes must work together to provide a consensus as to what is accepted into the log.
- + Once agreed, nodes replicate the log and create their own live view of the application state.
- + Standby nodes can “run behind” the active cluster without creating backpressure or creating latency issues on the critical path.

By removing the latency impact, Cluster Standby enables those features that would otherwise cause problems, e.g. snapshotting and offloading of slower operations.

Operational Considerations

- Users must be familiar with the network layout, bandwidth and link costs before deploying.
- Available as an Aeron Premium feature.

By  Adaptive

Adaptive builds & operates bespoke trading technology solutions across asset classes for financial services firms wanting to own their tech stack to differentiate and compete in the long-term. Central to Adaptive's offering is Aeron, the global standard for high-throughput, low-latency and fault-tolerant trading systems - the open-source technology supported and sponsored by Adaptive.

AERON GITHUB
COMMUNITY

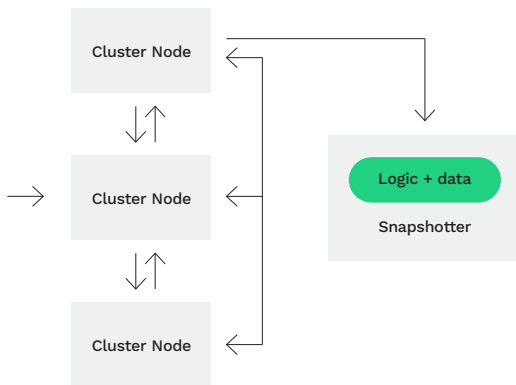


AERON
COOKBOOK



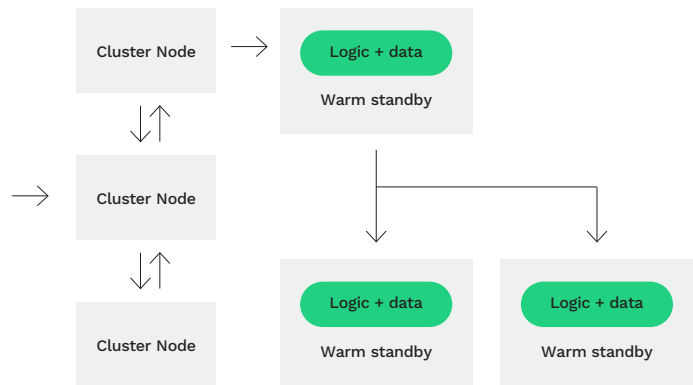


Aeron Cluster Standby - Use Cases



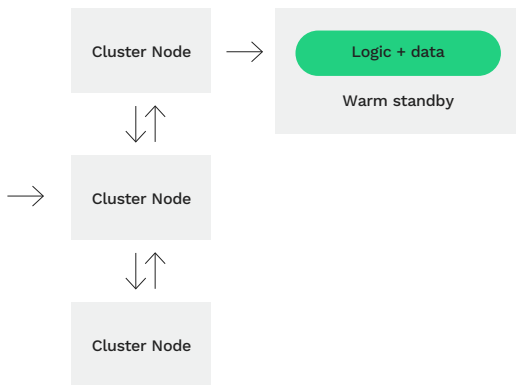
High Availability via Standby Snapshot

Instead of interrupting the cluster, snapshots can happen in the background allowing more frequent snapshots to be taken, improving time to recovery and 100% availability.



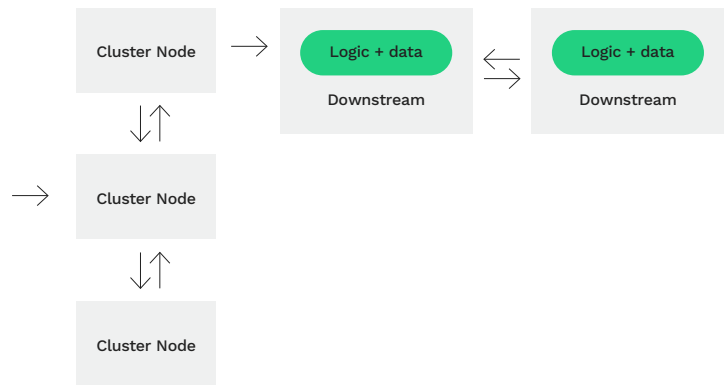
Disaster Recovery via Cluster Backups

- + Replicate cluster log data to standby node(s), ie. in remote locations.
- + Control bandwidth usage by allowing a single node to replicate from the active cluster and other nodes to data chain off that standby.
- + Keep application state live in memory so that failover can happen in seconds when invoked.



High Availability via Node Rebuilding

Allow for faster replacement of failed nodes, when the whole node becomes unreachable: shut down the cluster standby, start a normal cluster node over the same data and update the appropriate naming system (e.g. DNS) with the address of the standby node.



Load Management via Offload Functionality

Keep a full copy of the application state live in memory on the standby, allowing for interesting deployment options, e.g.:

- + Offload queries to a separate node.
- + Write event data to a slower data store for reporting or surveillance.