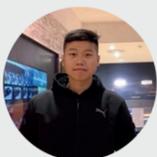


OPTIMIZING EC2 PERFORMANCE WITH PLACEMENT GROUPS

Swipe to know >>>



Aiden Tran
@aidentran

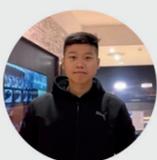




CLUSTER PLACEMENT GROUPS

“Cluster placement groups pack instances closely together into a low-latency group in a single Availability Zone.”

Use Cases: Best for applications requiring extremely low latency and high network throughput, such as large scale Big Data processing.



Aiden Tran
@aidentran





CLUSTER PLACEMENT GROUPS

- **Pros:** Ideal for network-bound applications, offering up to 10 Gbps bandwidth between instances with enhanced networking enabled
- **Cons:** Higher risk of simultaneous failure if the entire AZ is disrupted.



Aiden Tran
@aidentran

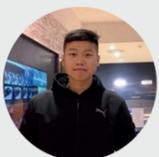




SPREAD PLACEMENT GROUPS

“Spread placement groups distribute instances across distinct hardware, limiting each group to 7 instances per AZ.”

- **Use Cases:** Suitable for critical applications where individual instance reliability is paramount.



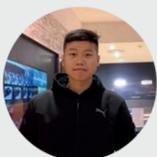
Aiden Tran
@aidentran





SPREAD PLACEMENT GROUPS

- **Pros:** Reduces the risk of simultaneous failures, enhancing fault tolerance by isolating instances from one another.
- **Cons:** Capacity limit can restrict scaling opportunities within a single group.



Aiden Tran
@aidentran





PARTITION PLACEMENT GROUPS

“ Partition placement groups organize instances into logical segments called partitions, which do not share underlying hardware.”

Use Cases: Ideal for distributed data stores like HDFS, HBase, Cassandra, and Kafka, where data resilience is crucial.



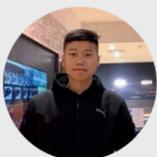
Aiden Tran
@aidentran





PARTITION PLACEMENT GROUPS

- **Pros:** Allows scaling to hundreds of instances while containing failures within a single partition.
- **Cons:** A single partition's failure affects all its instances, though not those in other partitions.



Aiden Tran
@aidentran

