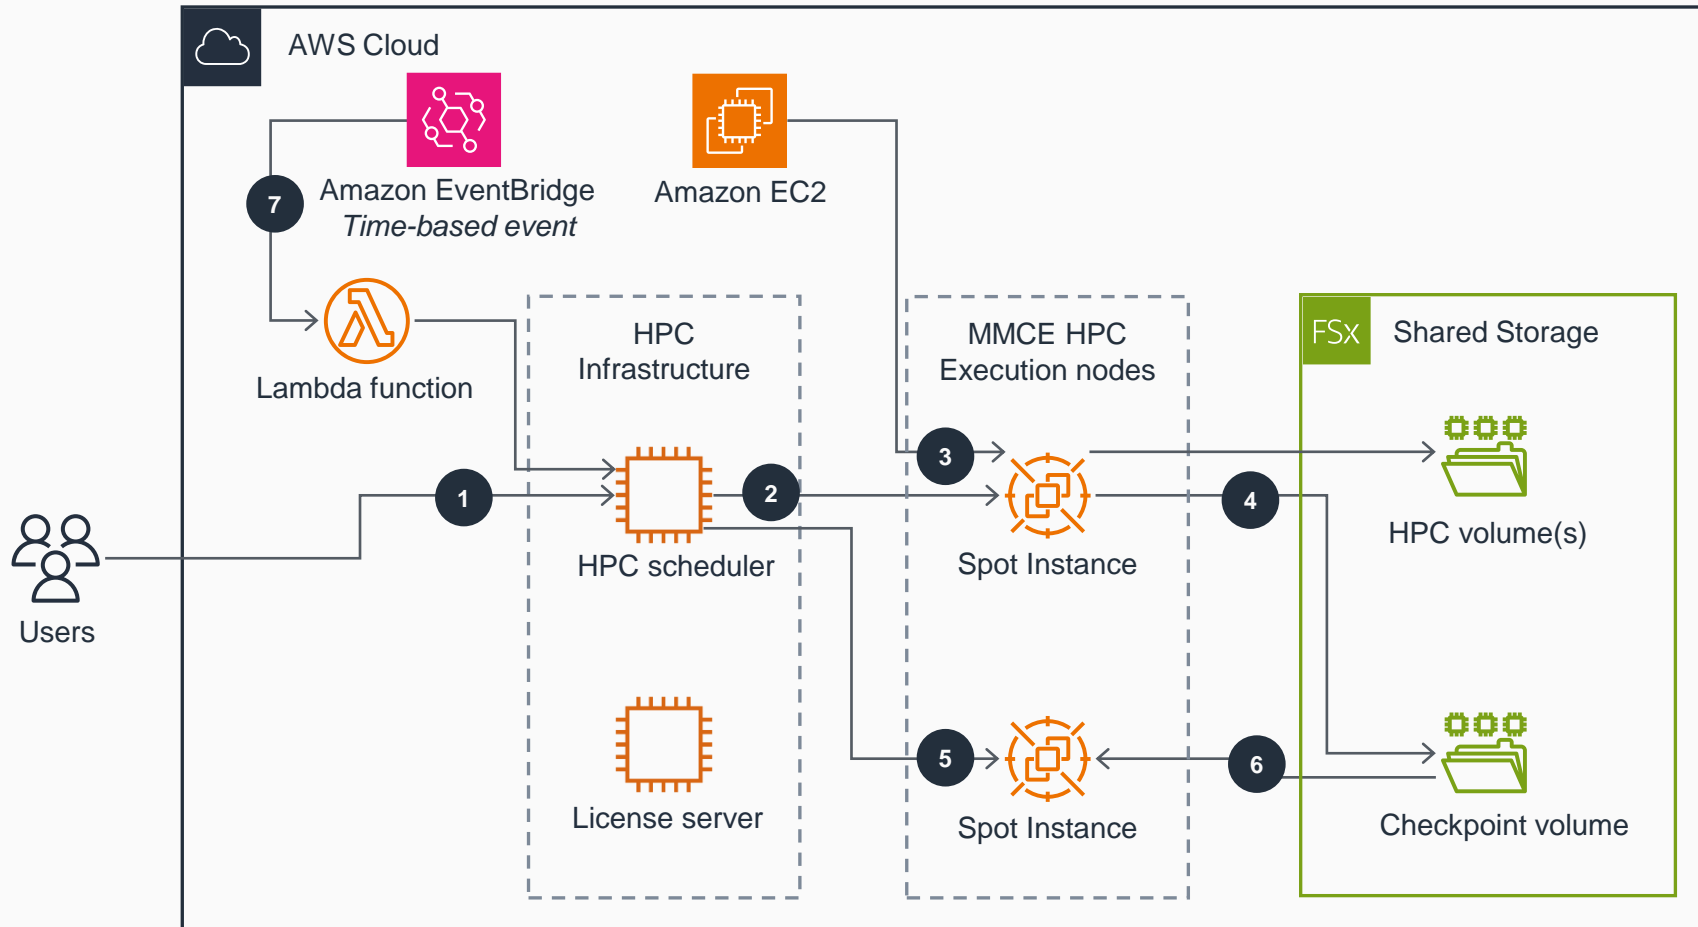


Guidance for Improving High Performance Computing Resiliency on AWS

This architecture diagram shows how to migrate long-running high performance computing (HPC) jobs to Amazon Elastic Compute Cloud (Amazon EC2) Spot Instances, reducing compute costs by up to 90%.



- 1 Users submit long-running HPC jobs to the scheduler.
- 2 The scheduler launches an **Amazon Elastic Compute Cloud (Amazon EC2) Spot Instance** to run the job. The MMCE checkpointing agent is pre-installed to generate periodic incremental checkpoints.
- 3 The **Amazon EC2** service provides a 2-minute advance notice of Spot Instance interruption.
- 4 When a spot interruption is detected, the MMCE agent creates an incremental checkpoint, capturing the process state, memory, and working directory, and stores it in a separate shared storage volume.
- 5 The scheduler launches a new Spot Instance to restart the job. This process is transparent to the scheduler, as no integration is required to handle the checkpoint.
- 6 Before launching the tool, a check for existing checkpoints is made. If a checkpoint is found, the MMCE is called to restore the checkpoint. This allows the job to resume from the last checkpoint instead of restarting the job from the beginning.
- 7 An **Amazon EventBridge** scheduled event triggers an **AWS Lambda** function, which submits an HPC job to clean up old checkpoints and reduce storage space.