

Solving KPF's Scaling Challenge with a Custom Solution on AWS

KPF

Executive Summary

[Kohn Pedersen Fox Associates \(KPF\)](#) is a global architecture practice focused on the design of buildings of all types and scales in all contexts across the globe. [KPF Urban Interface](#), an R&D team within the firm, uses a [computational design process](#) to generate thousands of 3D design schemes and simulate their performance across key metrics, including views, daylight, and energy efficiency. This is a process that is computationally expensive: it can take upwards of 10 minutes to generate and evaluate a single 3D massing. For the firm to perform these analyses at the requisite scale, the use of local machines is simply not feasible, as it can take a matter of days or even weeks for thousands of options to be produced and simulated. With the help of ClearScale, KPFui has developed a cost-effective, scalable, and largely automated infrastructure in which to run the software necessary for their computational design workflow. It enables the firm to generate thousands of 3D designs in a matter of hours, a mere fraction of the weeks it may have taken previously.

"With the help of ClearScale we have taken what was a frustratingly time-intensive workflow and turned it into a quick and easy one. We now use it far more often and have made it part of our primary workflow."

Luc Wilson, KPF Urban Interface, Director

The Challenge

KPFui uses [Rhinoceros \(Rhino\)](#), a Windows-based 3D modeling software that employs a mathematical model to render curves and surfaces in conjunction with the plugin [Grasshopper](#), a visual programming language and environment, to generate their computational design models. Both use algorithms and calculations that require considerable processing time and computing resources.

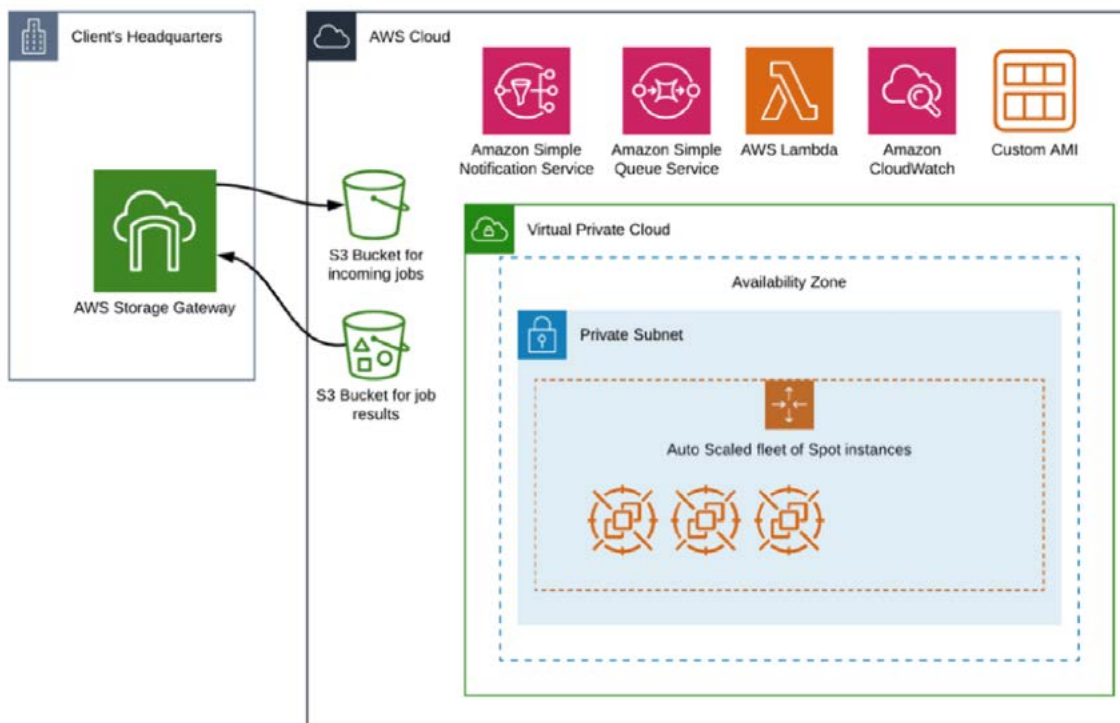
This created a challenge for the team at KPfui, who was running the software on desktop computers that lacked sufficient compute power and elasticity to handle the generation of thousands of 3D models in the window of time necessary to meet fast-paced project schedules. The team would separate the generation of the models into subsets, each team member would manually open four instances of Rhino + Grasshopper, and load a subset of the model to run on their desktop computer. In this manner, running out thousands of models could take a week and required both considerable upfront time to set up and monitoring of the process.

To increase processing efficiency, KPfui wanted to leverage AWS and run jobs in a scalable parallel environment by launching hundreds of Windows instances of Rhino + Grasshopper simultaneously and automate the process of distributing the 3D massing generation and analysis to the individual Rhino instances. They decided to seek out a partner in order to make this happen and found that [ClearScale, an AWS Premier Consulting Partner](#), possessed extensive expertise in developing AWS Cloud-based solutions – including those that involve [automation](#) and [software customization](#).

The ClearScale Solution

ClearScale worked with KPfui to understand the technical requirements of their computational design workflow and the challenge of automating and scaling the process in Rhino, which was not designed to function in this way. To solve this challenge, the ClearScale team wrote scripts that launched Rhino and loaded the KPfui files automatically, eliminating the need for users to open Rhino and load files manually. Next, ClearScale broke the scripts into multiple segments allowing the process to run in parallel environments. Set standards were defined for each chunk of data, and data chunks are passed to the software as a single iteration of an entire job.

Infrastructure Diagram



The job format was split across three stages: geometry preparation, analysis, and visualization. Each stage runs independently on its own scalable fleet of [Amazon Elastic Compute Cloud \(Amazon EC2\)](#) instances running Windows. Spot instances — unused EC2 instances that are available for less than the on-demand price — are employed to lower costs. The number of EC2 worker nodes in the cluster fluctuates based on the total number of source files.

To allow for the automating of tasks and to scale fleets of spot instances, ClearScale developed an AWS infrastructure for EC2 to handle [Amazon Simple Queue Service \(SQS\)](#) queues, [Amazon Simple Notification Service \(SNS\)](#) topics, and the functions of [AWS Lambda](#), which allows for running code without provisioning or managing servers.

The results of the three stages from the scalable cluster of EC2 worker nodes are then stored on separate [Amazon Simple Storage Service \(Amazon S3\)](#) object storage service buckets. An S3 bucket with the results of one stage can be used as a source data for the next stage. Running each stage independently allows the client to select (and pass as metadata) the type of EC2 instance for that stage. In addition, custom software was written in Java to talk to SQS and custom software scripts on the EC2 instances.

The Results

Previously, given both the user and computational expense, KPFui would only use their computational design tools once or twice a month. With the new infrastructure in place, KPFui is now able to use their computational design workflow regularly and keep pace with project deadlines.

Automation of the process has freed up employee time and has allowed them to take on more projects. Using cloud services, the firm's desktop computers are no longer burdened with running time-consuming and resources-intensive process in the background while they work on other projects.

The use of lower-cost spot instances has made the solution more cost-effective. Costs are further reduced because the EC2 fleet of instances are terminated as soon as data-intensive tasks in the computational design process are complete.