The development of monitoring technologies requires fast and accurate computational methods for neutron transport. We constructed a Reduced Basis Model (RBM) for 2D one-speed steady-state neutron diffusion to reduce the time-to-solution of parametric neutronics simulations. RBMs are systematic methods for Model Order Reduction (MOR) which use an offline/online paradigm, in which training snapshots are produced with a high-fidelity solver as some system parameter varies, and the subsequent training subspace is then reduced through proper orthogonal decomposition (POD). The RBM is created through Galerkin projection of the diffusion equation onto the reduced basis. We tested our RBM with a k-eigenvalue problem in a simplified reactor core quadrant, as the macroscopic absorption cross section in a fuel region was perturbed. The RBM saw improved error relative to perturbation theory with robust extrapolation outside the training subspace with four or fewer principal components. The reduction of the training subspace through POD showed the convergence of the RBM to the high-fidelity solution with each additional principal component. The speedup exceeded 100x. Future work includes the effective precomputation of nonperturbed parameters through affine decomposition, optimized training subspace construction through a greedy algorithm, and preserving angular information by constructing an RBM for the transport equation using a discrete ordinates solver.