Node embeddings are a powerful tool in the analysis of networks; yet, their full potential for the important task of node clustering has not been fully exploited. In particular, most state-of-the-art methods generating node embeddings of signed networks focus on link sign prediction, and those that pertain to node clustering are usually not graph neural network (GNN) methods. Here, we introduce a novel probabilistic balanced normalized cut loss for training nodes in a GNN framework for semi-supervised signed network clustering, called SSSNET. The method is end-to-end in combining embedding generation and clustering without an intermediate step; it has node clustering as main focus, with an emphasis on polarization effects arising in networks. The main novelty of our approach is a new take on the role of social balance theory for signed network embeddings. The standard heuristic for justifying the criteria for the embeddings hinges on the assumption that "an enemy's enemy is a friend". Here, instead, a neutral stance is assumed on whether or not the enemy of an enemy is a friend. Experimental results on various data sets, including a synthetic signed stochastic block model, a polarized version of it, and real-world data at different scales, demonstrate that SSSNET can achieve comparable or better results than state-of-the-art spectral clustering methods, for a wide range of noise and sparsity levels. SSSNET complements existing methods through the possibility of including exogenous information, in the form of node-level features or labels.