Node clustering is a powerful tool in the analysis of networks. We introduce a graph neural network framework, named DIGRAC, to obtain node embeddings for directed networks in a self-supervised manner, including a novel probabilistic imbalance loss, which can be used for network clustering. Here, we propose \textit{directed flow imbalance} measures, which are tightly related to directionality, to reveal clusters in the network even when there is no density difference between clusters. In contrast to standard approaches in the literature, in this paper, directionality is not treated as a nuisance, but rather contains the main signal. DIGRAC optimizes directed flow imbalance for clustering without requiring label supervision, unlike existing graph neural network methods, and can naturally incorporate node features, unlike existing spectral methods. Extensive experimental results on synthetic data, in the form of directed stochastic block models, and real-world data at different scales, demonstrate that our method, based on flow imbalance, attains state-of-the-art results on directed graph clustering when compared against 10 state-of-the-art methods from the literature, for a wide range of noise and sparsity levels, graph structures, and topologies, and even outperforms supervised methods.