We present first massively parallel (MPC) algorithms and hardness of approximation results for computing Single-Linkage Clustering of n input d-dimensional vectors under Hamming, ℓ_1, ℓ_2 and ℓ_∞ distances. All our algorithms run in $O(\log n)$ rounds of MPC for any fixed d and achieve $(1+\epsilon)$-approximation for all distances (except Hamming for which we show an exact algorithm). We also show constant-factor inapproximability results for $o(\log n)$-round algorithms under standard MPC hardness assumptions (for sufficiently large dimension depending on the distance used). Efficiency of implementation of our algorithms in Apache Spark is demonstrated through experiments on the largest available vector datasets from the UCI machine learning repository exhibiting speedups of several orders of magnitude.