We extend the theory of low-rank matrix recovery and completion to the case when Poisson observations for a linear combination or a subset of the entries of a matrix are available, which arises in various applications with count data. We consider the usual matrix recovery formulation through maximum likelihood with proper constraints on the matrix M of size d_1-by-d_2, and establish theoretical upper and lower bounds on the recovery error. Our bounds for matrix completion are nearly optimal up to a factor on the order of $O(\log(d_1 d_2))$. These bounds are obtained by combining techniques for recovering sparse vectors with compressed measurements in Poisson noise, those for analyzing low-rank matrices, as well as those for one-bit matrix completion [Davenport , "1-bit Matrix Completion, Information and Inference," Information and Inference, vol. 3, no. 3, pp. 189-223, Sep. 2014] (although these two problems are different in nature). The adaptation requires new techniques exploiting properties of the Poisson likelihood function and tackling the difficulties posed by the locally sub-Gaussian characteristic of the Poisson distribution. Our results highlight a few important distinctions of the Poisson case compared to the prior work including having to impose a minimum signal-to-noise requirement on each observed entry and a gap in the upper and lower bounds. We also develop a set of efficient iterative algorithms and demonstrate their good performance on synthetic examples and real data.