News is a popular and influential source of privacy information, and so, an important information source to analyze towards understanding privacyrelated policy, product development and user perceptions. We provide the first large-scale text mining of privacy news using nearly 1700 articles from the New York Times and the Guardian over the years 2010-2016. Results of four independently-trained sentiment classifiers, show that New York Times privacy news is predominantly negative in sentiment, and more negative than randomly sampled news, and some issues of global concern like the Syrian refugee crisis and the 2015 Paris terrorist attacks. We also provide a frequency analysis of entities involved (e.g. companies) over our entire data set and within certain privacy categories. Our analysis demonstrates that news text can be mined to facilitate both the tracking of important privacy events and efforts to build expressive taxonomies of privacy.