



Online translations versus secure internal network solutions.



● Samuel Läubli
CTO at *textshuttle.ai*

Web giants like Google offer their online translation services ‘for free’. In this article, we explain why companies preferring to keep their data private are well-advised to host translation software in their own network, and how this can be done without any hassle.

Over the last few years, translation companies and departments have realised the added value of partially automated translation workflows – despite initial scepticism. Machine translation (MT) allows them to meet the growing demand for rapid and cost-effective translation.

Take one of our clients, a Swiss financial institute, for example. Headquartered in a multilingual country, they need to translate a large number of factsheets and marketing texts from or to German, French and Italian – day by day and covering all levels of confidentiality. To be able to cope with the big amount of translations, they incorporated a customised MT system into their workflows. It is not least due to the sensitive data that, in the end, they decided against MT solutions from the cloud.

A short history of online machine translation services

Google was one of the first providers that sensed and faced the need for fast and easily accessible translations. Early versions of their freely available Google Translate service were impressive from a technical standpoint: text snippets of almost any domain could be translated into dozens of languages. However, the service was frustratingly imprecise from a linguistic point of view, mostly because it could only ensure fluency for short sequences of words, not entire sentences.

More recently, the advent of neural machine translation (NMT) led to a dramatic reduction of errors in morphology, word choice and word order. Translation services offered by big players like Google and Microsoft, as well as new competitors such as the German start-up DeepL, are all based on NMT by now.

How much are you willing to reveal?

Apart from being freely available, translation services from the cloud come in handy as they require no installation. Just hit a website, paste some text in a language you don’t understand, and your English translation is ready almost instantly. As simple as it seems, this takes a huge effort behind the scenes: companies invest millions to offer their services 24/7 at the highest speed possible.



Besides usage fees, providers of translation services collect data to monetise these investments. Google, for example, states in their Terms of Use that ‘you give Google a perpetual, irrevocable, worldwide, sublicensable, royalty-free, and non-exclusive license to Use content submitted, posted, or displayed to or from the APIs through your API Client. «Use» means use, host, store, modify, communicate, and publish.’

Companies like our client, who want to benefit from the most recent technology in automated translation but prefer to keep their data private, are well-advised to host translation services in their very own network. Customised systems can be trained on internal data and lead to significantly better results than the generic systems offered by cloud providers.

Easy deployment on client site through containers

It is thanks to a relatively new concept that the deployment of an on-site solution does not bring tears to the eyes of administrators anymore: the packaging of software solutions into containers.

At TextShuttle we adopted the use of Docker, a freely available software to build containers. If we develop a machine translation system for internal use on that basis, as it was the case in the example above, the system can be deployed with a single command. Whenever the client requests for changes, we implement and test them on our own servers first and pack the updated version into a new container – to swap it with the previous container on the client side, all at the push of a button.

The principle behind building containers is simple. Conventionally, all components of a software solution are installed on each server intended to host it. For one thing, this is complicated since there are typically various packages that need to be installed and configured. For another, companies often rely on servers with different hardware and software configurations clustered together. This requires IT administrators to adjust installation and configuration steps to each server, resulting in high effort for installing, updating, and recovering software in case of incidents.

When relying on containers, in contrast, only one piece of software – the ‘container ship’, so to speak – needs to be installed on all servers. Irrespective of operating system and hardware, it provides a standardised runtime environment. Having the same runtime environment at their disposal, developers like ourselves can now pre-install complete software solutions into containers.

Moreover, containers on different servers can be clustered with Docker Swarm. This facilitates the distribution of software services across multiple machines, e.g., to allow for rapid response times with increasing numbers of users. Services like machine translation can be hosted locally with little effort in this way – without disclosing data to Google or anyone else.

Short CV of the author: Samuel Läubli. Before joining the Machine Translation group at the UZH Institute of Computational Linguistics as a PhD Student in 2016, Samuel Läubli implemented MT systems geared to post-editing as a Senior Computational Linguist at Autodesk, Inc. (2014–2016). He obtained a Master’s degree in Artificial Intelligence from the University of Edinburgh (2014), and a Bachelor’s degree in Computational Linguistics from the University of Zurich (2012). Samuel is CTO at textshuttle.ai, a spin-off specialising in machine translation and opinion mining.

