

MOSES CORE

Deliverable D 4.8

Moses MT Market report

Work Package:

WP 4: Industry Outreach

Date (mm/yy):

January 2015

Dissemination level:

Public

Authors:

Achim Ruopp, Jaap van der Meer

Table of Contents

1	Introduction.....	4
2	The History of Moses and Development Principles	4
2.1	History	4
2.1.1	Start of the Moses Open Source Project	4
2.1.2	Funding through Euromatrix and Euromatrix+ and Start of MT Marathons	5
2.1.3	Recent Changes during the MosesCore project	5
2.2	Development Principles	5
3	The Present of Moses: Market Adoption	6
3.1	Types of Players.....	6
3.1.1	Introduction.....	6
3.1.2	MT Suppliers.....	7
3.1.3	Free Machine Translation.....	7
3.1.4	Value-Added Resellers.....	8
3.1.5	In-House MT Users	8
3.2	Industry Adoption – Types of Offerings	9
3.2.1	Introduction.....	9
3.2.2	Paid Licenses.....	9
3.2.3	Machine Translation as a Service	10
3.2.4	Paid Licensed and SaaS Moses-based Solutions in the MT Solutions Market	11
3.2.5	Volume-Based MT Services	12
3.2.6	Professional services	13
3.3	Industry Adoption – Adoption Models.....	14
3.3.1	Different models.....	14
3.3.2	Model 1: Moses out-of-the-box	16
3.3.3	Model 2: Moses with open-source/free add-ons.....	16
3.3.4	Model 3: Commercialized Moses	17
3.3.5	Model 4 – Hybrid MT.....	17
3.3.6	Summary.....	18
4	The Moses MT Market as Part of the Larger MT Market.....	19
4.1	Market Volume and Sizing.....	19
4.2	Moses Impact on the MT Market beyond Solutions	21
4.2.1	Cost Savings	22
4.2.2	Increased Diversity of MT Offerings	22
5	Conclusions.....	22
	Appendix.....	23

Profiles of Moses-based MT Solutions	23
MT Showcase Use Cases.....	24
MosesCore Newsletter Use Cases.....	27
MT Suppliers.....	28

1 Introduction

This report gives an overview of the landscape of Moses Machine Translation: the history, its adoption by industry, estimates of market size and demographics, types of users and usages, types of players in the market of Moses-backed solutions and their offerings and potential future scenarios for Moses use outside academia. It's a valuable source of information for everyone who is interested in MT and specifically in the Moses open-source MT system.

Open-source machine translation systems do not differ from proprietary systems from a technical perspective. Virtually all approaches to machine translation are available as software code and packages licensed under an open-source license, albeit some never developed a large contributor base and their development has slowed down.

The crucial difference of open-source development in comparison to proprietary software is that the development of a complex MT system core including a decoder (the software component that performs the translations) is the shared effort of a loosely organized group of expert developers while still allowing the commercial exploitation of the results.

In the field of open-source machine translation, the developers are often academically funded or the original source code comes from a since-abandoned proprietary software development effort (OpenLogos). In some cases academic institutions insist on academic licenses which prohibit commercial use. This often creates duplicate development efforts in academia and industry for the same core algorithms and abandoned academic software. In our opinion, allowing the use of academically developed code by anybody for any purpose is justified, as the research is most often publicly funded and it strengthens the development effort for any use, including academic use.

This report explores how the availability of Moses as an open-source solution has influenced the MT market, Moses' contribution to making a broad range of MT solutions available to users and the novel uses Moses enabled in the market as well as in internal use by organizations that would otherwise not have been possible.

2 The History of Moses and Development Principles

2.1 History

2.1.1 Start of the Moses Open Source Project

The open-source statistical machine translation toolkit Moses was started in 2005 at the University of Edinburgh by Hieu Hoang under the supervision of Philipp Koehn. It was developed as an open-source successor of the Pharaoh decoder that was developed earlier by Philipp Koehn. While Moses is backwards compatible to Pharaoh, it was designed to be more modular way and to support factored decoding. Factors are a way to inject linguistic information like parts of speech into the statistical process.

Moses was first released to the public under the LGPL (Lesser or Library GNU Public License) after further development at a 2006 summer workshop at Johns Hopkins University. The workshop provided a big boost to development and broadened the developer base. This first release included an end-to-end pipeline for the building of statistical phrase-based MT systems with the optional use of factors.

2.1.2 Funding through Euromatrix and Euromatrix+ and Start of MT Marathons

After this promising start the Moses project started to receive funding from the European Union through the Euromatrix program (2006-2009) primarily to provide a common experimentation platform for the field of machine translation. Euromatrix also allowed the organization of the first MT Marathon in 2007, a gathering of students, researchers and the Moses community at large to discuss new, innovative ideas and to integrate some of the most promising ideas into Moses.

The Euromatrix funding period and the follow-on Euromatrix+ funding (2009-2012) saw the addition of new concepts of statistical machine translation like syntactic models, the addition of toolkits like the integrated language model KenLM, the addition of feature functions and many other contributions.

The Euromatrix/Euromatrix+ funding firmly established Moses as the prime experimentation platform for statistical machine translation in the academic community. By having this platform a lot of groundwork is already laid and academic researchers can try new ideas in a short amount of time.

2.1.3 Recent Changes during the MosesCore project

Moses firmly established itself in academia, but thanks to the liberal open-source license also saw increased use in the industry and government sectors in recent years. The funding agencies noticed this development and for the MosesCore funding period starting in 2012, under which this report is funded, added an industry/government focus to enable the use of Moses by organizations small to large, particularly to address multi-linguality challenges in the common European market. This included hiring the original developer Hieu Hoang as the Moses coordinator, defined releases including installable packages developed by the project partner Capita TI and the promotion of the Moses toolkit by the industry organization TAUS, a further project partner, via MT Showcases at industry events and online.

This report focuses specifically on analysing the non-academic use of Moses, which usage patterns have emerged, how these compare to usage patterns of other well-known open-source projects and the impact of the MosesCore project on this use.

From the technical side MosesCore brought faster MT system training with parallelism at all stages, pre-built binaries, pre-trained models for a range of European language pairs, the testing of Moses on Linux, Mac OSX and Windows (with Cygwin) and the release of virtual machine images to get started without having to go through installation. Further the code base was refactored to make improving and expanding Moses easier, which allowed the recent addition of new features like the updateable phrase-table, a feature that allows edits by post-editors to be learned immediately for all future machine translations with the system.

2.2 Development Principles

Moses is developed and maintained mainly by academics across the world under the leadership of the University of Edinburgh. It serves as an experimentation platform for research – new research ideas can be added and evaluated quickly without having to re-implement previous techniques already present in Moses. To make this distributed development as efficient as possible the Moses project follows an open development model with many branches and many committers. To ensure that basic functionality does not get broken in the open development process, regression tests get run frequently.

The open development approach is supported by the extensive documentation, by hosting the project on the popular distributed version control (DVC) service GitHub and the active Moses-support mailing list where everybody from novice to expert is welcome to ask questions. The project coordinator Hieu Hoang, along with the core team of developers keeps users informed about major feature additions and tries to keep the project backwards compatible to allow users to upgrade more easily.

Development is not entirely conducted online – with funding from the European Union the Moses team also conducts a yearly, one-week, in-person Machine Translation Marathon during which new ideas and proven concepts get integrated into the code base

By following these modern open-source development practices the Moses project has gained a large following of contributors and users in academia, government and industry. While continuous development is sufficient for many users, some users and the funding agencies asked for a more defined release cycle – since the start of the MosesCore project the development team publishes annual releases.

The growing use of Moses by a diverse set of users presents the challenge to balance between providing a completely flexible research platform and providing a productive MT decoder for industry/government use.

This report aims to shed some light on the usage of the open-source project by a variety of non-academic users and hopefully can serve as a tool to inform future areas of investment.

3 The Present of Moses: Market Adoption

With Moses-based solutions having captured a quarter of the MT solutions market in revenue (see chapter 4), Moses-based solutions represent a complete cross-section of types of MT offerings and are offered by all types of players in the MT market. By providing an open-sourced basis for statistical MT, Moses even enables more diversity in the MT market – because solution suppliers can use the shared development of the core system, they can concentrate their development resources on a diverse set of features ranging from increased usability over easier deployment and improved privacy to higher machine translation quality. This benefits the end users of machine translation with increased availability of MT services with the features that they need at reduced prices. Ultimately this serves the multilingual market, including the common European market with increased market access and increased volumes of translated content at lower costs.

In the following sections we provide our analysis on the types of MT solution market players and types of offerings for the MT market at large and the Moses MT market in specific.

3.1 Types of Players

3.1.1 Introduction

For the translation sector machine translation, like many technologies before, is an enabling technology. It complements and enhances human translation and allows for more content to be translated with faster turnaround times. Machine translation can help to discover more content that needs to be translated. Translation buyers request translations of content they discover through the use of machine translation in cross-lingual search/eDiscovery, content they were previously unaware

of. Machine translation as an enabling technology supports existing business models in the translation sector and unlocks new business models. Beyond the translation industry, machine translation makes it easier to immerse oneself in a foreign language environment and converse with people speaking a different language.

Machine translation is a complex technology to master, implement and provide. As such, it has been offered by a number of commercial suppliers throughout its history. The availability of open-source implementations of state-of-the-art machine translation technologies has broadened the supplier base.

Companies, governments and other institutions that use the enabling technology MT to provide a service to their customers, citizens and users have a complex business decision to make when deciding on a supplier or deciding to go it alone using open-source code.

3.1.2 MT Suppliers

Traditionally MT suppliers often started out with a strong economic basis of large government or institutional buyers and then expanded into a broader market offering for buyers large to small. Companies like SYSTRAN and Language Weaver followed this path. The reason has been the large up-front technology investment necessary to build an MT system, let alone customize it.

Most recently the market has shifted a bit - on one hand large players deemed MT a strategic advantage and decided to invest personnel and resources into the development of in-house MT; on the other hand new suppliers adopted newly available open-source MT systems like Moses allowing them to forgo the core technology investment and innovate on tools surrounding the MT offering. This shift already led to a number of acquisitions of the more traditional MT suppliers: Language Weaver was acquired by SDL in 2010, Systran was acquired by CSLi after 46 years of independent existence and most recently AppTek was acquired by eBay.

We believe that Moses has put market pressure on the existing traditional MT suppliers and accelerated mergers and acquisitions in the MT market. This benefits language services providers like London-based SDL, acquirer of Language Weaver.

3.1.3 Free Machine Translation

In recent years, some variants of machine translation have become available for free in the form of machine translation of web pages and mobile apps for ad-hoc translation. The main suppliers of free machine translation are Google and Microsoft along with Yandex and Baidu. And while MT in this form is not customized, it is available in many language pairs reflecting the international reach of these companies.

Of course free, online machine translation is not truly without cost – it is cross-subsidized by uses that generate revenue for the company offering them. For instance, the automatic translation of web sites generates more clicks for a web advertising company like Google. For a platform company like Microsoft integrating machine translation into different products increases the value of each of the products and leads to more purchases. If translators are using the free services, they are required or at least encouraged to contribute corrections back to the providers in order to improve future MT quality.

Getting plain text machine translation from the online services via an API is a paid service, which is unsurprising given the considerable infrastructure costs that machine translation still incurs. The cost

is still low and small web or app developers use the API service to provide free content machine translations cross-subsidized with advertising.

The large online MT services do not disclose the exact technology powering their services. We know from company publications however that Google¹, Microsoft², Yandex³ and Baidu⁴ develop machine translation in-house. These companies employ large teams of machine translation researchers and have the desire to own the technology end-to-end without having to adapt an open-source technology like Moses.

3.1.4 Value-Added Resellers

The most popular uses for machine translation are gisting (assimilation of content), search and discovery, sentiment analysis of foreign language content, post-editing and speech translation. For each of the uses there is at least one value-added reseller integrating a raw machine translation engine, either purchased from a supplier or self-developed, with the necessary components and user interface to provide the usage in an easily accessible fashion.

Moses, because its code is available for modification, provides a flexible platform for providers of solutions where MT is just a value-added component. Uses where MT is such a value-added component are

- Gisting – the understanding/assimilation of foreign language content using raw machine translation, search and discovery of such content along with automated sentiment analysis. For example in the eDiscovery, news analysis, big data and intelligence fields.
- Speech translation
- Post-editing – the editing of machine translations by human translators to reach human translation quality

Particularly the last use sees increasing adoption in the language industry given proven productivity gains for many types of translated content.

While many value-added resellers choose MT solutions from the MT suppliers described earlier, some see operating their own their own MT systems as a cost effective strategic solution. The resellers often choose Moses as the platform to use. Examples for Moses users are Translated (Matecat), Lexworks, Welocalize and Unbabel on the translation provider side and Bloomberg on the data analysis side. On the speech translation side we do not know whether providers use Moses. Public declarations of Moses use by data analysis and speech translation providers are rare. This does not mean however, that Moses isn't extensively used in these fields.

We do expect further innovation and unexpected use cases where machine translation technology is packaged with complimentary services like search or speech recognition/generation. Moses provides an excellent platform for start-ups in this area.

3.1.5 In-House MT Users

A significant share of machine translation is done in-house at governments, multinational organizations and multinational companies. In recent years more large companies and organizations

¹ <https://research.google.com/pubs/MachineTranslation.html>

² http://research.microsoft.com/en-us/about/our-research/computational_linguistics.aspx

³ <https://company.yandex.com/technologies/translation.xml>

⁴ http://translate.baidu.com/static/help_en.html

decided that it is of strategic importance and cost effective to develop MT knowledge and infrastructure in-house. This is done either with the acquisition of MT suppliers (often as an acquire-hire of assets and talent) or by hiring personnel and then building the infrastructure on top of open-source solutions like Moses.

Examples of in-house users of Moses are the European Commission and WIPO as governmental organizations and SAP, eBay, Oracle and Adobe as multinational companies. These are organizations that adopt Moses directly; others might use Moses-generated translations indirectly by buying solutions from the MT suppliers and value-added resellers described above.

The challenge for in-house Moses users is to build out the MT solution to the desired infrastructure and language pair scale and then maintain this deployment meeting user requirements and integrating innovations in a cost-efficient manner.

Academic users, as individual researchers or research institutions, can be counted in this category as well, as they typically don't buy solutions. They need the flexibility of an in-house deployed and often heavily modified MT solution for research purposes.

3.2 Industry Adoption – Types of Offerings

3.2.1 Introduction

Now that we provided a detailed breakdown of the types of players in the Moses MT market, we want to take a closer look at the types of offering that some of the players provide.

The proliferation of uses of machine translations in the translation industry and beyond has led to a proliferation of machine translation offerings from an ever increasing number of vendors. With a wide variety of needs of machine translation buyers ranging from speed, volume, quality, privacy, language pairs, customizability and last, but not least, pricing, the emergence of a large number of offerings is not surprising.

With the investment necessary to provide comprehensive machine translation solutions we predict that usage clusters will emerge and the consolidation of offerings for each cluster is to be expected. In the realm of free online machine translations, this consolidation already happened with Google, Microsoft, Yandex and Baidu being the only ones providing this at scale.

While some consolidation already happened in the area of large online machine translation services with only four services available at internet scale as well as in the area of proprietary MT suppliers (sales of Language Weaver, Systran, and AppTek), the field in the area of customized machine translation systems for translation providers is still quite open.

With the novel technology, both business and technical consulting services have prospered to guide users through the technology adoption. The need for these services is diminishing as issues around machine translation are better understood by the industry and more out-of-the-box solutions are available.

3.2.2 Paid Licenses

In the past, reflecting the way that software was deployed generally, machine translation systems were almost exclusively available as paid licensed software. Server and desktop variants were available for individual language pairs or language pair bundles. Especially with rule-based machine

translation systems, this customer-site deployment also made sense because the systems needed to be heavily customized by expert users.

While desktop MT systems are playing a diminishing role in the market, server paid licenses are still an important offering. Systems are typically offered with per-server paid license and maintenance contracts.

On-site installed systems typically provide a lot of flexibility and options for customization. This requires staff to be trained on this customization or customization services/consulting being provided by the supplier (which we discuss in detail in the section “Customization” further down). In some cases, the MT software does not come with baseline general domain MT engines for the desired language pair and engines need to be trained – users have to obtain the appropriate training data.

Another major factor for on-site deployment is privacy concerns and requirements. Where disclosure embargos or regional containment of content is required, on-site server deployments of MT systems are a must.

As the number of words that can be machine-translated with an on-site system is typically unlimited, purchasing this offering also makes economic sense in use cases where large volumes of words need to be translated.

Overall, we expect paid licenses to lose market share to cloud offerings as the cloud promises more flexible use, immediate quality updates and lower prices. In fact, in the market of Moses-based MT solutions there are only two providers that exclusively offer a paid licensed on-site solution: Precision Translation Tools with DoMT Desktop, Server and Enterprise and Simple Shift with myMT. Precision Translation Tools is planning the roll-out of a cloud-based SaaS solution in 2015. Other Moses-based suppliers offer on-site paid licensed solutions for customers with specific privacy requirements and/or high volume translation throughput, but their focus is on a hosted SaaS solutions they provide.

The dearth of Moses-based solutions for on-site installation of course also has to do with the fact that interested users can download the open-source code base as a starting point. Many in-house Moses users in fact do this.

3.2.3 Machine Translation as a Service

Machine Translation as Software as a Service (SaaS), in the simplest case, is an MT server provisioned on the MT provider’s infrastructure and made available for customer use. From a technical perspective, there isn’t much difference to an on-premise server, except that the customer doesn’t have to worry about running and maintaining hardware.

This SaaS provisioning has however often very different pricing models: suppliers offer service level agreements and ask for subscriptions and/or maintenance fees to be paid on a regular basis without an upfront license fee.

Subscription pricing can work to the benefit of both parties: the supplier gets a more predictable revenue stream, while MT users do not have to pay large up-front license fees and can subscribe/unsubscribe based on business needs that are often cyclical in the translation business. Providing MT systems as a SaaS solution also allows the supplier to take advantage of economies of scale: multiple users can use the same general domain MT system or MT systems for multiple users can be run on one server or a cluster. This can raise privacy and scalability worries with some users,

but for many users, provided they can negotiate service level agreements they are comfortable with, this is the right size solution. There is no question however, that SaaS means a certain loss of control for the user with long-time provider survival being the biggest concern.

The loss of control for the user, to some extent, also extends to the customizability of the MT system. Given the complexities of MT customization and training, this is for many cases a positive. SaaS providers offer easy to use web-based user interfaces to train, evaluate and optimize MT systems. Users can upload their own data and build MT engines on the fly. The web-based SaaS solutions also offer integration into the most popular translation management systems and computer-aided translation tools. The ease of integration and use is in fact one of the major selling points of SaaS MT suppliers, particularly the ones that share the Moses open-source solution as their core engine.

In addition to the attractive proposition to build an easy-to-use, modern web interface around Moses, which the open-source solution does not provide, deploying Moses as part of a SaaS solution offers two additional advantages: 1) the supplier can manage the complex setup and integration of Moses in an environment that the supplier controls 2) the LGPL license under which Moses is licensed requires that changes to Moses are shared with users of the code once modified code gets distributed, e.g. installed at a customer site – this requirement doesn't apply to SaaS solutions.

3.2.4 Paid Licensed and SaaS Moses-based Solutions in the MT Solutions Market

The MosesCore project period from 2012 to 2014 has seen a large increase in solutions based on Moses available in the marketplace – the number of solutions available either as paid licensed software or SaaS solutions increased from 5 to 15.

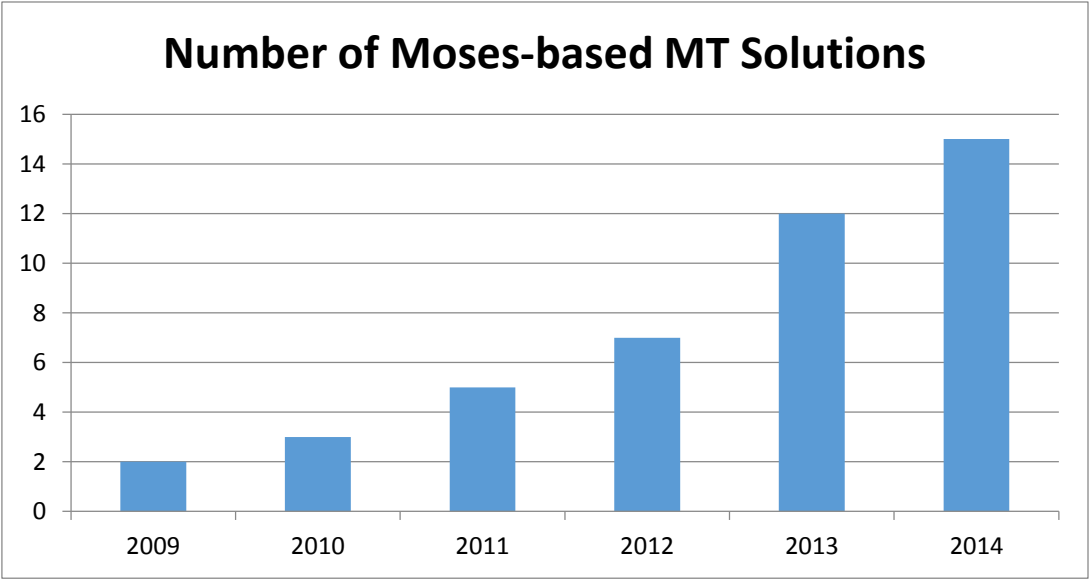


Figure 1: Moses-based solutions available in the market

The increased availability of MT solutions based on Moses provides buyers of MT solutions with a diverse set of options from which they can choose the solution that fits their needs and budget best. We believe that the outreach effort that TAUS conducted since the early days of Moses industry adoption and intensified with support from MosesCore directly contributed to the increased availability of Moses-based solutions in the marketplace.

We compiled a list of Moses-based MT solutions along with the solution release date (official releases, not Beta releases) based on information from the solution provider websites and direct information from the provider obtained in interviews. This list can be found in the appendix.

With 15 available solutions, Moses-based solutions now make up the majority of solutions in the market. Eight of the solution providers are located in the European Union. Based on our provider interviews we estimate about 7 employees per provider, so we estimate the direct creation of about 50 jobs in the European Union.

There are only a handful of solutions based on proprietary MT decoders left in the market:

- Rule-based engines like Lucy LT
- The large online MT providers Google, Microsoft, Yandex and Baidu
- Language Weaver from SDL
- Real-Time Translation Server from IBM

While the Moses-based solutions now outnumber other solutions in the market, they don't bring in the majority of the revenue as our quantitative analysis in chapter 4 shows.

3.2.5 Volume-Based MT Services

Volume-based MT services are cloud-based online MT services that take the SaaS concept one step further – they aim to provide machine translation in many language pairs on a global basis via the internet. There are few of these services at this scale and Google Translate and Microsoft Translator are the most well-known ones. As mentioned in the “Type of Players” chapter, these services are partially cross-subsidized and available for free for certain uses to serve larger company objectives. Google and Microsoft have in recent years built payment systems for various API offerings and also make their machine translation systems available through paid APIs. In April 2014, Yandex also started offering an MT service through the use of an API. The table below shows the latest pricing by monthly volumes for the use of the MT APIs from Yandex and Microsoft.

	Yandex	Microsoft
2M characters / month	\$30	Free
4M characters / month	\$60	\$40
6M characters / month	\$90	\$60
8M characters / month	\$120	\$80
16M characters / month	\$240	\$160
32M characters / month	\$480	\$300
64M characters / month	\$768	\$576
128M characters / month	\$1.280	\$1.058
250M characters / month	\$2.000	\$2.055
500M characters per month	\$4.000	\$3.995
635M characters / month	\$3.810	\$5.000
1000M characters per month	\$6.000	\$6.000

Table 1: Pricing of volume-based online MT services

Google charges \$20 per one million characters with a default limit of two million characters per day. Developers and integrators who want to increase the daily volume to fifty million characters per day or even more can contact Google.

Whether the online MT services provided through APIs are profitable and how they are prioritized in comparison to other MT uses by the companies is hard to determine for these large public companies. Priorities might also shift with larger strategic changes. For example, in 2011 Google announced its Translate API would be retired due to abuse. Google only relented after legitimate users campaigned keep the API open and offered to pay for the machine translations. The Google Translate API is now available on as a paid API.

Nevertheless, these online MT services are immensely popular and API access is integrated into virtually every translation tool. Google provides its own online CAT tool, the Google Translate Toolkit, of course with tight integration of Google MT. With their low per-word prices these services set a price floor for the other MT suppliers.

By default, the machine translations from these services are general domain and not customized, although Microsoft offers the Translator Hub for customization. Machine translations are available for many language pairs and for some language pairs these online systems are the only easily available option. The provided general domain machine translations set a high quality floor for other machine translation suppliers – Google and Microsoft employ some of the best MT researchers and developers and the systems are continuously improved.

In order to improve the systems Google and Microsoft also pool data, either training data or post-edited translations. For users with strict privacy requirements this can be a concern – Microsoft recently offered the option for increased privacy for large volume users of their system. Microsoft and Google also use web data to train their systems – by now these systems have become so ubiquitous that they start to train on their own translations. Both providers are taking steps to address this issue.

As of now we know of no volume-based MT services that use Moses. On one hand this has to do with internal decisions of the four large providers, but also indicates that the Moses code base is maybe not yet up to the task of providing an internet scale MT service. After initial work on clustering with the Sun Grid Engine the focus in the performance area has been the optimization for single machines with ever larger RAM and number of processor cores. This still favours small- to mid-size deployments and more investment is needed to make Moses run efficiently on compute clusters.

3.2.6 Professional services

3.2.6.1 Customization

With the technical and business complexities of adopting a new technology like MT, professional services are often offered around the different offerings. Some MT suppliers pride themselves in providing a full end-to-end service to their customers, from analysing customer needs, customizing the MT systems with customer data, to provisioning and maintaining the systems to integrating the systems into the customers processes. Cleaning and preparation of customer data for the purpose of training MT systems is part of this service.

This model serves customers very well who see MT as an enabling technology allowing them to provide their existing core services more efficiently or develop new services that enhance their existing business without seeing in-depth MT knowledge as a core competency. In the context of the translation business, this means for example providing human quality translations with faster turn-around times at lower costs or generating additional demand in human translation by presenting machine translated content for gisting to the end customer. Most commonly hands-on professional

services are provided in combination with the licensed MT/SaaS offerings described above in B2B settings.

This is also true for all Moses-based licensed MT/SaaS offerings – all providers offer consulting around their offering, even if for some, like KantanMT, the focus is primarily on customer self-service. The providers have to find the right balance between personnel-intensive consulting and self-service product offerings that scale easily to a large number of users. With technically complex systems like machine translation this balance is often hard to find without turning away potential customers.

Sometimes even internal users of Moses or Moses-based MT suppliers themselves engaged external expert consultants, mostly on a short-term basis in order to ramp up the adoption of MT technology faster. For internal users and MT suppliers, it is important to bring all knowledge to customize and use MT in-house in order to achieve the desired institutional competency and independence. This kind of consulting was popular in the early days of Moses industry adoption, but has diminished with these users having built up sufficient knowledge and/or teams to use Moses internally or provide a Moses-based solution to the market.

3.2.6.2 Business consulting

New technology like MT presents opportunities to streamline processes and opportunities to generate new business. These opportunities are not immediately obvious to business leaders that might be set in doing business a certain way and are not familiar with the benefits that machine translation can offer for their industry.

Business consultants that have industry experience and shared industry knowledge how the new technology can be applied can help businesses to make the right decisions on machine translation adoption to stay ahead of the competition.

3.3 Industry Adoption – Adoption Models

3.3.1 Different models

The development principles described in chapter 2 fostered the adoption of the Moses package for use in governments and industry. The crucial factor for government and industry adoption however, is the licensing under an open-source license allowing commercial use – in this case the LGPL⁵ (Lesser or Library GNU Public License), which allows use for any purpose. The LGPL requires that changes to the code are shared with the users of the code once the code is distributed (e.g. installed at a customer site). It does allow the use of the software as a library and the code of the application using the library does not need to be published.

Together these factors created favourable conditions for non-academic use, making Moses the most popular and widely used machine translation package. TAUS observed the development of Moses in the non-academic space closely, we conducted user surveys from 2011 to 2014 and we actively promote the use of Moses through outreach events and dissemination activities.

⁵ Details about the LGPL license can be found here: <https://www.gnu.org/licenses/lgpl.html>

We observed the following adoption patterns among the users of Moses:

Model 1	Model 2	Model 3	Model 4
Moses out-of-the-box	Moses with open-source/free add-ons	Commercialized Moses	Hybrid MT
Download and build Moses and associated components yourself	Pre-built installation to make getting started easier Additional scripts to support language industry needs	Web-hosted or self-hosted UI to build MT engines and translate documents Various add-ons for data cleaning, post-editing, language support etc.	Moses combined with other MT technologies (e.g. RBMT engines) to provide high-quality MT output
Numerous translation buyers, research institutes, start-ups, software developers, language service providers and a technically savvy translator or two Internal Users Logrus Adobe European Commission WIPO SAP eBay CrossLang EMC and many more (many not publicly known)	DoMY Moses for Mere Mortals Moses for Localization Capita TI MosesCore packages Casmacat Matecat	Solutions: Tauyou <text> Capita TI Smartmate Asia Online Language Studio™ Pangea MT Tildes Let's MT Simple Shift myMT ⁶ Sovee SmartEngine Precision Translation Tools DoMT KantanMT.com Iconic Translation Machines IPTranslator Lingua Custodia Alpha Engines Morphologic Globalese	Systran Enterprise Server ProMT Translation Server EMTGlobal

Table 2: Moses Adoption Patterns

⁶ Open source under the AGPL v3.0 license at <http://olanto.org>

By conducting events and publishing outreach TAUS had the opportunity to gather a large repository of use cases for all Moses adoption patterns, types of players and offerings. This is a great resource for any organization thinking of adopting Moses for its MT needs. References to the use cases are included in the appendix in the sections “MT Showcase Use Cases” and “MosesCore Newsletter Use Cases”.

3.3.2 Model 1: Moses out-of-the-box

Since the inception of Moses compiling the software from the source code has been the most commonly used method to obtain the software for expert users. The source code can be downloaded from GitHub for many different branches and compilation instructions are provided. This method requires knowledge how to build software from source code on UNIX systems. Necessary additional components like libraries, word alignment and language modelling software need to be downloaded and built from separate source code repositories. This method is still the only method to get the latest innovations and bug fixes.

Since the start of the MosesCore project, the Moses team has also branched release versions of the code at the start of each year, stabilized the code and released it in source code form and binaries for the Linux, Cygwin and OSX platforms. These release versions are schedule-driven rather than feature-driven, so that any features present in the main branch at the time of the release fork get included in the release.

The team has resolved side-by-side installation issues, so that different versions of the software can be installed on the same machine – which is often necessary as data formats still evolve. The binaries include necessary 3rd party components like a language modelling toolkit and a word alignment toolkit. However, installation packages are not provided by the team and extensive UNIX system administration knowledge is still required for installation and use.

Once the Moses binaries along with the word aligner, language modelling toolkit and support scripts are installed, Moses out-of-the-box provides a complete suite of command line tools to train MT systems for a core set of European languages and use the trained systems to translate plain text. Corpus data to train the systems and tools for languages outside of the core set have to be provided by the user.

3.3.3 Model 2: Moses with open-source/free add-ons

Developers in the broader open-source community attempt to address needs surrounding out-of-the-box Moses for non-expert users and users in the translation industry. These efforts revolve around three main issues – the ease of installation of the software, the ease of use of the software and the integration of the Moses system into existing translation industry systems, e.g. via APIs or translation industry standard formats.

For the released Moses versions the MosesCore partner Capita TI provided installation packages for Redhat and Debian based Linux distributions as well as for Windows. The Windows version provides a basic user interface to translate text and plain text files. The MT systems for use with the Windows version have to be trained on Linux.

An early effort to make Moses more user-friendly was started by some individuals in the European Commission appropriately named Moses-for-mere-mortals. It was an ambitious effort to provide a

user-friendly Windows user interface to train, evaluate and use a Moses system on a UNIX server. Unfortunately the project has been only infrequently updated.

Another open-source effort is the DoMY project lead by Precision Translation Tools which focuses on installation and data processing. It does not provide a user interface and recently Precision Translation Tools refocused its development efforts on their proprietary DoMT offering.

The Moses for Localization project consists of two parts: a component to translate different file formats with the open-source localization package Okapi and the Adobe Moses Tools, allowing corpus processing, MT system training and evaluation. Development on this effort has also slowed down.

Casmacat and Matecat are two web-based open-source translation environments funded by EU projects that use Moses as the MT engine backend. They allow tracking translator interaction and edits and dynamically adjust the MT engine based on the edits.

Each of these efforts addresses only part of the needs of non-expert users of Moses and neither has picked up a large user and/or contributor base. An easily installable, usable Moses package for users in the translation industry is currently not on the horizon.

3.3.4 Model 3: Commercialized Moses

The innovative Moses open-source code base built over a number of years, the missing user interface and translation industry tools/standards integration and the licensing under the LGPL license provided an excellent opportunity for MT system integrators. A growing number of MT system integrators have taken advantage of this opportunity and built solutions with web-based user interfaces, integration components and evaluation tools on top of the Moses code base. The integrators sell these solutions into the growing market for MT solutions, particularly for the translation industry. Profiles of these companies can be found at the end of the report.

Most of the commercialized Moses solutions focus on MT system customization using user-provided training data in order to differentiate themselves from the large online MT providers Microsoft and Google which are often seen as benchmarks. The assumption is that translation providers have privacy concerns in uploading translation memories to the online services. Some providers offer professional customization services, while others want to achieve scale and market penetration with largely automated customization. Some providers offer pre-built systems for popular vertical industries like travel, automotive and patent translations. For users with increased privacy requirements some MT providers offer on-premise installation, but most focus on offering SaaS solutions.

Other than differentiation on the technical level, which frankly often only consists of a rather thin user interface layer on top of Moses, there is also differentiation around pricing – this can be a one-time fee, subscriptions or word-based pricing. Suppliers in this expanding and rapidly changing market of MT solutions try to gain market share or find the right niche providing a sustainable economic basis for their offering.

3.3.5 Model 4 – Hybrid MT

Model 4, hybrid MT, is similar to Model 3 but differs in that additional proprietary MT technology is integrated into the solution in addition to Moses – a technical “secret sauce” that can significantly increase MT quality for certain use cases or language pairs. (Side note: “hybrid” here refers to the

combination of open-source and proprietary MT technologies, rather than the hybridization of different MT approaches)

Particularly traditional RbMT providers, with the emergence of SMT systems in the marketplace, could not justify the high customization costs of RbMT systems anymore and were forced to fuse SMT technology into their RbMT systems to make customization easier and/or increase MT output quality.

3.3.6 Summary

Through our monitoring of the market, including user interviews, we observed that Moses is adopted by suppliers of Model 3 and Model 4 solutions in the marketplace as well as large internal users like government organizations and large multinational companies, which see the full control of their MT destiny as a strategic goal. Only these organizations have the resources to employ an expert team to build well-performing MT engines and adapt Moses to their users needs.

We see little or no direct Moses adoption among small and mid-size companies and the companies of this size that adopted Moses early report diminishing use with the availability of customization in online MT services (Microsoft Translator Hub) and the increasing availability of competitive MT offerings in the market (many of them Moses-based). Individuals or small companies that do adopt Moses admit in our interviews that they do not know if the required significant time investment will ever be matched by sufficient returns on investment, i.e. they regard their use of Moses as a way to educate themselves about SMT and experiment with it.

Moses, for non-academic purposes, is therefore an open-source project that is similar to Android or the Linux Kernel. Android provides an operating system to build a mobile device, but applications, user interface modifications and the hardware need to be added. The Linux Kernel is a battle tested operating system kernel, but libraries, a compiler and a user interface to deliver a functioning operating system need to be added. End users do not build mobile devices with Android or assemble operating systems using the Linux Kernel – system integrators do. In the case of Moses suppliers of Model 3 & 4 solutions and large internal users are the system integrators.

Moses in its current form does not offer a complete, easy to install and use solution like the open-source solutions Open Office, MySQL and several Linux distributions do. Attempts have been made with Model 2 projects, but up to now have failed to catch on. We think this is because machine translation is a niche market in comparison to the market for these more general purpose open-source solutions. In addition, the in-depth background knowledge required to train and operate Moses further reduces the number of potential adopters. One also has to keep in mind that even the mentioned easy-to-install-and-use open-source solutions can be regarded loss leaders for the large companies (Oracle and Redhat) that now own them. It remains to be seen if a comprehensive Model 2 solution will ever emerge.

We do think that it is instructive that Moses can be compared to Android or the Linux kernel to chart a sustainable path forward for Moses development for industry and government. Google is the corporate backer of Android who is interested in moving its development forward, but Google also wants to exert some control on the direction of the Android, e.g. via the Open Handset Alliance⁷ and Android One. A similar corporate backer has not emerged for Moses, as the large online providers Google and Microsoft have their own implementations of statistical machine translation. A potential

⁷ <http://www.openhandsetalliance.com/>

other stewardship model can be seen in the example of the non-profit Linux Foundation⁸ – a broad coalition of Linux Kernel users which got together to ensure the sustained development of this essential resource used in devices ranging from airplane in-seat entertainment systems to compute clusters in the financial industry. The Linux Foundation supports the development of the Linux Kernel, helps companies navigate the ever-changing Linux landscape and promotes the use of Linux in new scenarios. The options for sustainable future development of Moses are outlined in more detail in the sustainability report D5.5.

In summary, the availability of the open-source Moses has significantly changed the dynamics in the MT supplier market, though maybe not to the same degree as the online translation services from Microsoft, Google, Yandex and Baidu.

4 The Moses MT Market as Part of the Larger MT Market

4.1 Market Volume and Sizing

There is a wide variety of uses for machine translation including gisting (assimilation of content), search and discovery, sentiment analysis of foreign language content, ecommerce, post-editing and speech translation. MT is truly an enabling technology for many new innovations in the language sector. For this report we define the MT market purely as the market of MT solutions, not of any of the solutions/services that use MT. For example if Bloomberg offers a monitoring service for foreign language use using MT, it would be very hard for us to determine which share of the revenue generated by the monitoring service would be attributable to MT.

Post-editing revenues, which make up a large share of the business use of MT, can per accounting rules not be counted as technology revenue showing that the major source of the revenue generation is considered to be the work of the human post-editor, not the MT component pre-translating the source text.

TAUS in early 2014 surveyed the entire MT solutions market and we estimated the size of the MT market in 2014 at \$250 million⁹. Included in this estimate are revenues from licensing MT software (both server and desktop licenses), professional services (customization, consultancy) and hosted MT service offerings (SaaS and volume-based charges for the use of an MT API). Some companies quoted their annual revenues, but for most of them we have made our estimates based on circumstantial information and data. The most uncertain data are the estimates for the use of the APIs from Google and Microsoft. We are also certain that we have missed quite a few language service providers in our assessment that do offer a brand new hosted MT service or customization service. Of the 64 identified MT operators in TAUS' MT solutions list about twenty were established in the last five years or launched their MT offering in the last five years. However, revenues from these brand new MT providers and operators, including many Moses-based solutions, are still very low, most of them well below \$1 million.

For this Moses MT market report we identified 21 of the 64 MT operators as Moses-based (the appendix "MT Suppliers" lists the suppliers) and we estimate the market share of these operators to be about \$45 million or about 20% of the entire MT solutions market. The annual revenue growth rate the surveyed Moses-based MT supplier report averages 15% which matches very well with the

⁸ <http://www.linuxfoundation.org/>

⁹ <https://www.taus.net/think-tank/reports/translate-reports/mt-market-report-2014>

estimated growth rate of 16.9% for the MT market overall – the Moses-based MT suppliers don't grow faster than the market. We expect their market share therefore to stay roughly the same.




Figure 2: Revenue per offering type Moses-based solutions

In our survey of Moses-based MT suppliers we also asked for a breakdown of revenue by product or service offering type and region. We had to exclude the numbers reported by CSLi/Systran so that they would not totally dominate the estimate – the revenue of this provider is still roughly an order of magnitude higher than of the mostly new Moses-based MT market offerings. We wanted the charts to reflect the general market for these new solutions rather than just one provider.

The revenue breakdown by offering type for Moses-based solutions (Figure 2) is very similar to the breakdown for the MT market overall (Figure 3) with the exception of the word/volume-based offering. The reason for this is that no large online MT service is powered by Moses (yet).




Figure 3: Revenue percentage per offering type MT Market overall

We also asked our Moses-based respondents where geographically their revenue is coming from (Figure 4) and again contrasted this with TAUS estimates for the MT market overall (Figure 5). Moses-based solutions show strength in Europe which can be attributed to the healthy eco system for Moses-based start-ups. This eco system is also due to public funding for the project including talent development initiatives. Moses-based solutions also show strength in emerging markets which is typical for many open-source solutions – providers in these markets can use the low cost of entry for the base open-source solution and low-cost labour to build solutions addressing local market needs without having to rely on larger suppliers potentially not even particularly interested in these markets.




Figure 4: Revenue percentages per geography Moses-based solutions

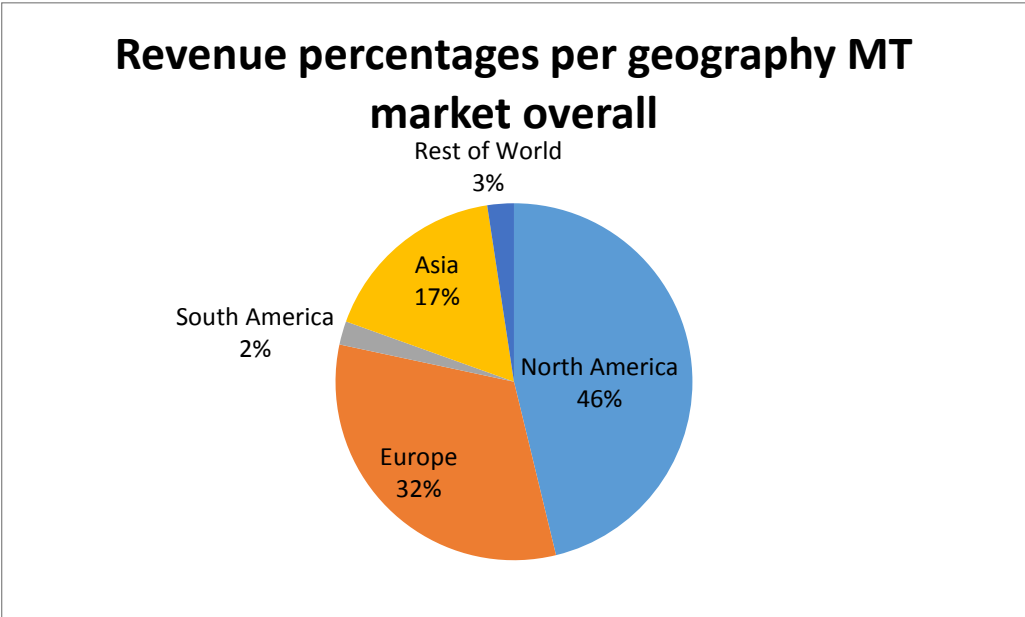


Figure 5: Revenue percentages per geography MT market overall

4.2 Moses Impact on the MT Market beyond Solutions

Moses impact on the MT Market goes way beyond the mere participation of Moses-based solutions in the MT solutions market. Like many other open-source projects it has an impact that isn't always quantifiable in absolute Euros and Dollars and often is not even trackable as the use of the project doesn't require registration.

Here are some of the impacts that TAUS observed.

4.2.1 Cost Savings

Users of Moses including Value-added resellers of machine translation, in-house users and also MT suppliers benefit from the years of collaborative development effort that went into the code base. Estimating that the code would cost in the seven-figure Euro range to develop from scratch and the number of users can likely be counted in the hundreds, we estimate the overall cost savings for industry and government users in the hundreds of millions of Euros.

Downloaders and users of the Moses code base are not tracked and many of them stay anonymous for a variety of reasons. This approach is typical for open-source projects and used for Moses since its inception. The approach was intentionally not changed during the funding periods Euromatrix, Euromatrix+ and MosesCore. Funders have to be aware that this implies that meaningful ROI estimates in terms of the overall market are hard/impossible to calculate.

Of course one could argue that in the absence of an open-source solution not all of these users would have develop an MT system by themselves and instead could buy a solution from a proprietary MT solution supplier. With license fees running into the thousands or even tens of thousands of Euros this was and is still not a financially viable solution for many of these users.

With the open-source system Moses available, users can invest some of the cost savings into specific adaptations they need for their processes. Value added resellers and solution suppliers can invest in features their customers most desire.

So Moses, along with the large online MT services from Microsoft, Google, Yandex and Baidu has put downward pressure on prices in the MT market. With the availability of Moses statistical machine translation (SMT) has effectively been commoditized and value is moving to adjacent services using MT. Despite the commoditization of the software we still see the training data essential for the training for SMT systems retain its value.

4.2.2 Increased Diversity of MT Offerings

With the majority of available MT solutions now based on Moses, Moses has not only led to increased price competition, but also feature competition.

Today we see in the MT market diverse offerings with different feature sets, different pricing models for low resource languages (e.g. the Baltic languages). Users are able to pick exactly the solution that fits their needs even if choosing from many solutions might be harder. This diversity in offerings would not be possible without the availability of Moses as an open-source solution.

5 Conclusions

In the market of MT solutions, a market that is small overall, but playing an outsize role to provide machine translation as an enabling technology, Moses-based solutions have captured a significant

market share of about 20%. Moses has contributed significantly to the diversity of solutions in the market place and with the low barrier to entrance has put pricing and innovation pressure on all providers in the market, thus benefiting the MT end user.

Moses' influence goes beyond the easily observable and quantifiable influence in the solutions market – it is one of the most flexible, complete and mature MT platform that large organizations and start-ups can base their machine translation technology on, enabling innovation and cost-savings.

Though a complete easy to use and integrate open-source version of Moses providing machine translations to non-technical end users like translators still is not available we see many efforts in this direction and hope that such a solution can emerge from the increasing use of MT in the language industry. Other open-source solutions like Okapi and OmegaT in this industry have also taken years to get traction, likely due to the smaller user base compared to more mainstream open-source applications like Firefox or Open Office.

Because community participants on all the described levels of adoption have a significant stake and interest in the future of Moses, TAUS is quite positive about the vibrancy and importance of the Moses MT system in the future, independent of any particular source of funding. We hope to be part of its success in years to come.

Appendix

Profiles of Moses-based MT Solutions

Product Name	Company	Country	Solution Type	Solution release	Moses Adoption Model	URL
<text>	Tauyou	Spain	SaaS, Onsite	2012	3	http://www.tauyou.com/newindex.html#text
Smartmate	Capita TI	United Kingdom	SaaS, Onsite	2011	3	https://www.smartmate.co/
Language Studio™	Asia Online	Thailand	SaaS	2009	3	http://www.asiaonline.net/EN/Default.aspx
Pangea MT	Pangea MT	Spain	SaaS	2011	3	http://pangeamt.com/en
Let's MT	Tilde	Latvia	SaaS, Onsite	2012	3	http://www.tilde.com/mt/letsmt
myMT	Simple Shift	Switzerland	Onsite	2013	3	http://simple-shift.com/myMT.html

SmartEngine	Sovee	United States	SaaS	2014	3	http://www.ovee.com/translation-technology/
DoMT	Precision Translation Tools	Thailand	Onsite	2013	3	http://www.precisiontranslationtools.com/products/
KantanMT.com	KantanMT	Ireland	SaaS	2013	3	http://www.kantanmt.com/features.php
IPTranslator	Iconic Translation Machines	Ireland	SaaS	2013	3	http://iconictranslation.com/
Alpha Engines	Lingua Custodia	France	SaaS, Onsite	2014	3	http://www.linguacustodia.fr/en/our-services
Systran Enterprise Server	Systran	Korea	SaaS, Onsite	2009	4	http://www.systransoft.com/translation-products/
	ProMT	Russia	SaaS, Onsite	2010	4	http://www.promt.com/company/technology/smt/
EMTGlobal	Safaba	United States	SaaS, Onsite	2013	4	http://www.safaba.com/machine-translation-solutions
Globalese	Morphologic	Hungary	SaaS, Onsite	2014	3	http://www.globalese-mt.com/

MT Showcase Use Cases

The MT Showcases, organized by TAUS and funded by MosesCore, presented a broad range of use cases for all Moses Adoption Models described in section 3.3.1. This provides a great repository of knowledge for future adopters of Moses. Links to all use case presentations can be found on <http://www.statmt.org/mosescore/index.php?n=Main.Videos>.

Organization	Location	Month	Topic	User Type	Moses Adoption Model
Trusted Translations	Monaco	March 2012	Moses from the point of View of an LSP: The Trusted Translations experience	Value-added reseller	1
AVB Translations	Monaco	March 2012	A Moses engine for legal translation	Value-added reseller	3

Logrus	Monaco	March 2012	"Moses, Moses: Let my people go." Moses MT engine feasibility study	Value-added reseller	1
Applied Language Solutions	Monaco	March 2012	High quality self-serve MT in SmartMATE	MT supplier/Value-added reseller	1
Tilde	Monaco	March 2012	Moses on the Cloud. Do-it-yourself Machine Translation	MT supplier	1
tauyou	Monaco	March 2012	Friendly Machine Translation	MT supplier	1
Adobe	Beijing	April 2012	Moses tool set. A set of tools based on Adobe technology to simplify your usage of Moses	In-house user	1
Applied Language Solutions	Beijing	April 2012	High quality self-serve MT in SmartMATE	MT supplier/Value-added reseller	1
Institute of Automation, Chinese Academy of Sciences	Beijing	April 2012	How we use Moses to develop our multi-lingual Machine Translation systems	In-house user	1
Pangeanic	Paris	June 2012	Pangea MT putting open standards to work	MT supplier	1
Hunnect	Paris	June 2012	The ups and downs of implementing an MT environment for English - Hungarian	Value-added reseller	1
Trusted Translations	Paris	June 2012	Moses: The Trusted Translations Experience	Value-added reseller	1
CrossLang	Paris	June 2012	Bologna Translation Service: Making education accessible across Europe	MT supplier	1
Sybase	Paris	June 2012	An MT journey: MT in use at Sybase, a SAP company	In-house user	1,3
Symantec	Paris	June 2012	Moses inside Symantec	In-house user	1,2
TAUS	Seattle	October 2012	TAUS Dynamic Quality Framework	In-house user	1,2
Adobe	Seattle	October 2012	The Simple Install – Streamlining Moses Setup for Industry Scale Users	In-house user	1,2
tauyou	Seattle	October 2012	Language Processing Techniques for Statistical Machine Translation	MT supplier	3
Safaba	Seattle	October	Full Service Enterprise-	MT supplier	4

Translation Solutions		2012	Specific MT for Global Enterprises		
AVB Translations	Seattle	October 2012	Two Practical Use Cases at AVB Translations	Value-added reseller	1,3
TAUS labs	Seattle	October 2012	Moses Tutorial and Other Open Resources	In-house user	1,2
TAUS	Seattle	October 2012	The Landscape	In-house user	1,2
TAUS	Singapore	April 2013	TAUS DQF	In-house user	1,2
Precision Translation Tools	Singapore	April 2013	A Small LSP's Guide to Commercialized Open Source SMT	MT supplier	2,3
Google	Singapore	April 2013	Google Translator Toolkit	MT supplier/Value-added reseller	n/a
Hunnect	Singapore	April 2013	Hunnect's Use Case	Value-added reseller	1
Institute for Infocomm	Singapore	April 2013	MT for Southeast Asian Languages	In-house user	n/a
Asia Online	Singapore	April 2013	Strategies for Building Competitive Advantage and Revenue from Machine Translation	MT supplier	3
TAUS	Singapore	April 2013	Moses and Other Resources	In-house user	1,2
TAUS	Singapore	April 2013	Introduction and Overview	In-house user	1,2
TAUS	London	June 2013	The Dynamic Quality Framework Tools	In-house user	1,2
Pangeanic	London	June 2013	I Used to Be a Translator. Now I Run MT	MT supplier	3
LexWorks	London	June 2013	Moses in the Mix. A Technology Agnostic Approach to a Winning MT Strategy	Value-added reseller	1
Safaba Translation Solutions	London	June 2013	The True Value of MT to Global Business	MT supplier	4
TAUS Labs	London	June 2013	Moses and Other Open Resources	In-house user	1,2
University of Edinburgh	London	June 2013	Moses Past, Present and Future	In-house user	1
TAUS	London	June 2013	The Evolving MT Landscape	In-house user	1,2
Tilde	Santa Clara	October 2013	MT & Terminology: better together	MT supplier	1
Welocalize	Santa Clara	October 2013	The WeMT Program	Value-added reseller	1
Microsoft	Santa	October	Microsoft Translator	MT supplier	n/a

	Clara	2013			
KantanMT	Santa Clara	October 2013	Creating Competitive Advantage with Rapid Customization & Deployment of Moses	MT supplier	3
TAUS	Santa Clara	October 2013	The Open Source MT System Moses and Its Use in the Industry	In-house user	1,2
TAUS	Santa Clara	October 2013	Is the Translation Industry Ready	In-house user	1,2
Tilde	Dublin	June 2014	MT applications in the EU public sector	MT supplier	3
Sovee	Dublin	June 2014	Sovee Smart Engine 2.0: A Leap Beyond Base Moses Technology	MT supplier	3
KantanMT	Dublin	June 2014	Enabling MT for everyone!	MT supplier	3
Iconic Translation Machines	Dublin	June 2014	Beyond Data: Delivering Machine Translation with Subject Matter Expertise	MT supplier	3
EC	Dublin	June 2014	MT@EC for European public administrations and online services	In-house user	1
Translated	Vancouver	October 2014	Introduction to Matecat, the Open-source CAT Tool for Post-editing	Value-added reseller	1
Unbabel	Vancouver	October 2014	Seamless Globalization with Distributed Crowd Post-editing	Value-added reseller	1
Precision Translation Tools	Vancouver	October 2014	The Simplified Guide to Getting Started in SMT	MT supplier	2,3
eBay	Vancouver	October 2014	Machine Translation at eBay	In-house user	1
TAUS	Vancouver	October 2014	TAUS Introduction and MT Market Overview	In-house user	1,2

MosesCore Newsletter Use Cases

For this report we interviewed about 20 users of Moses in the industry – based on the interviews we published detailed use case articles for the most representative use cases for users that permitted publishing this information. Links to the newsletters including the use cases can be found on <http://www.statmt.org/mosescore/index.php?n=Main.Newsletters>.

Use Case	User Type	Moses Adoption Model
Trusted Translations (USA)	Value-added reseller	1
LexWorks (France/Canada)	Value-added reseller	1
Moravia (Czech Republic)	Value-added reseller	1,2
bmmt (Germany)	MT supplier	1,2
Citrix (USA)	In-house user	1
KantanMT (Ireland)	MT supplier	1

MT Suppliers

List of MT suppliers in early 2014 (Moses-based suppliers in bold):

SYSTRAN

Language Engineering Company, LLC

Kawamura

LexWorks

Reverso Softissimo

CCID TransTech Co., Ltd

Moravia

PROMT

Raytheon BBN

SDL

SkyCode Ltd.

SyNTHEMA srl

Mutilizer

Smart Communications

BABYLON

LionBridge

Trident Software

Automatic Trans

E4NET

GrammarSoft ApS

Translated

Worldlingo

Lingo24

MorphoLogic Localisation

AppTek

CrossLang

Eastlinden

Hunnect

Capita Translation and Interpreting

CSLi

Eleka Ingeniaritza Linguistikoa, SL+C14

Pangeanic

Lingotek

Lucy Software

tauyou

AsiaOnline

Safaba Translation Solutions

SmartLing

Olanto International Foundation

Tilde

Baidu

KantanMT

Sovee

UTH International

Modulo Language Automation

BMMT

Datamundi

Iconic Translation Systems

Unbabel

Prompsit Language Engineering

Google

Microsoft

Yandex

ABBYY

CSOFT

IBM

K2E-PAT

Kodensha

Lingenio

Lingosail Technology Co., LTD.

Linguistic Systems

NICT

Precision Translation Tools

Yamagata