

Revolutionizing customer support through MT

Lori Thicke

You just took some great pictures with your new camera that you want to share online, but you can't get the transfer function to work so that you can get them onto your computer. Later, your new productivity software is turning into a real time-waster as you try for the third time to install it.

We've all been there. For most of us, rather than call customer support, we'd like to find answers to our problems online. We may consult company-generated content such as knowledge bases and frequently asked questions (FAQs), but we may be just as happy looking for answers from other users, in forums, tweets and blog posts.

Luckily for us, we speak English. But imagine you bought your camera in Finland. Or your productivity software in Korea. Chances are slim that you'll find the answer to your problem if you are not very comfortable in English, or in one of the main commercial languages such as French, Spanish or Japanese.

Customer support has not caught up to the globalization of customers. According to Greg Oxton, executive director of the Consortium for Service Innovation, a nonprofit industry alliance of high-tech customer support organizations, "Growth markets for most companies today are not in their home country. Their markets are global." The implications of this for the translation and localization industry are clear: "The majority of customer support interactions are with content, not people, and that is driving demand for fast, economical localization capabilities."

Today, however, relatively few companies offer online technical support in the languages of their customers. Even when global customers account for more sales than domestic customers, support content is often available only in the domestic language, or in a limited panel of mainstream languages. This leaves out a great, and growing, number of users who also purchased products or services but who, unlike their domestic counterparts, are not fully supported.

In the past, supporting only one or a few countries was understandable. The amount of support material contained in companies' knowledge bases is massive. According to Jaap van der Meer, founder of the innovation think tank TAUS, what is currently being translated is only the tip of the pyramid. The base of

the pyramid consists of vast stores of knowledge-based content around a product or service.

Customer support content represents the lion's share of what a company needs to translate today. A typical company might have two to four million words of customer support – or much more. Microsoft has over 120 million words in English knowledge base articles. Multiply that by the languages of Microsoft's users and you see the problem. The cost of human translation is prohibitive.

But it gets worse. Support content is ever-changing as new products and services come out. Around 2,000 of Microsoft's articles are updated every week. Remember, this is just company-generated content. In fact, user-generated content is where the revolution is really happening, and that content is expanding exponentially each and every day. Yet it's not just this massive volume of content that makes human translation out of the question for most if not all languages. It's the speed with which we need to be able to access that content for it to be useful to us. No one wants to wait a week to have the fix to a problem. As the world goes social, mobile and global via a variety of devices (even in the emerging markets), customers are connected by social media and looking for answers now. Massive quantities of content needing real-time translations means that machine translation (MT) is the *sine qua non* for customer support.

Supporting international customers, especially in countries with relatively few customers, has been seen as a "nice to have." But Common Sense Advisory has strong arguments that translation is a "must have" for companies that want to grow more. In its 2006 "Can't Read, Won't Buy" report, we see that customers are up to six times more likely to buy when addressed in their own language. Furthermore, many emerging markets are English intolerant and can't or won't use English websites at all.

Support is also a differentiator, according to Oxton. "In a world where business is global, social and mobile, service excellence equals maximizing customer value and minimizing value erosion." Talk to your customers in their language and they will like your brand a whole lot more. Customers who can access



Lori Thicke is cofounder and general manager of LexWorks, cofounder of Translators without Borders and a member of the MultiLingual editorial board.

information in their own language report greater satisfaction and increased brand loyalty.

Here's one more reason to offer solid support in the local languages of your international customers: it'll help you stave off homegrown competitors. Consumers in countries such as China and India have been quick to vote with their renminbi and their rupees for companies that offer them products and services – and support – in their own languages.

MT for customer support

Global is expensive. Vast online knowledge bases contain a wealth of problem-solving material, but on that quantity, traditional translation could run into millions of dollars. On the other hand, to paraphrase Maurice Chevalier, consider the alternative. If translating millions of words of technical support content is expensive, the alternative of maintaining international help desks ain't cheap either. The cost of providing telephone or e-mail customer support can run \$100 or more per incident.

Companies need to solve their customers' problems before they turn into expensive calls to the help desk. The good news is that with better knowledge management, customer self-service and peer-to-peer communities, MT can drive customer engagement.

Where the costs of a human translation are prohibitive, in our experience MT works almost as well, and for a fraction of the cost. Even imperfect (unedited) MT generates higher satisfaction rates when the alternative is no translation at all. This is related to Oxton's concept of "sufficient to solve." Support content has different quality requirements than sales or marketing content. Users are more willing to accept language imperfections in exchange for useful information.

For example, a study by the Consortium for Service Innovation revealed that when customers go onto a support site, what's most important to them is "Information that is technically accurate and relevant." This may not accord with what many language professionals believe, but according to this same study, what is least important to customers is punctuation, grammar, complete sentences and correct spelling.

Keeping in mind what customers really want – understandable content – the right question is not "Is this a good translation?" but rather "Did this translation answer your question?" MT may not be as "good" as a human translation, but for solving customers' problems, it comes darn close.

Just how close has been measured by Microsoft, and presented to the AMTA 2012 conference by Chris Wendt, group

program manager of MT for Microsoft. Martine Smets from Microsoft customer support looked at a range of languages such as Chinese, Korean, Dutch, German, Italian, Turkish, Portuguese, French, Spanish, Russian, Polish, Japanese and Arabic and concluded that MT was only slightly less effective than a human translation for solving customer problems.

Intriguingly, the average knowledge base resolve rate (the percentage of customer problems resolved online via human translation versus MT) for these languages is only slightly lower – just 8% – than a fully human translation. That is, MT solved customer problems 54% of the time, while a fully human translation solved their problems 62% of the time. MT is capable of helping almost as many customers, and for a fraction of the cost.

Customers overwhelmingly support MT when the alternative is no translation. A leading software publisher recently found that, depending on the language, from 72% to 95% of customers reported higher satisfaction when they were provided with MT support information. Intel found that visits to its Spanish customer support site since the introduction of MT went from 20,000 to 90,000 in just the first four months.

With regard to professional customers, Microsoft found that its Japanese MVP (Most Valued Professional) customers overwhelmingly considered MT to be worthwhile. Fully 83% found MT helpful, and 92% wanted more MT content.

ROI on customer support

Will Burgett, product manager of translation innovation and services at Intel Corporation, reported as far back as 2008 that by using MT for customer support, Intel has seen the following benefits:

- Empowering customers to find the help they need
- Delivering content to customers faster
- Improving customer loyalty and satisfaction
- Reducing the customer-support translation cycle from ten days to one
- Decreasing localization costs by 95%
- Growing international sales, especially in emerging markets

But perhaps nowhere is the return on investment (ROI) so easy to measure as by the reduced burden on help desk/call

work in
perfect harmony

XTM cloud

XTM Cloud is an award winning on-line translation tool. Whether you work independently or collaborate with colleagues, XTM Cloud will provide you with the latest cutting edge technology at a competitive price. XTM Cloud is a complete solution incorporating an advanced translation workbench with translation memory, machine translation, terminology and workflow management. Register with XTM Cloud today to join the XTM community and win more business.

Sign up for a free 30-day trial at xtm-intl.com/trial

T: +44 (0)1753 480479
e: sales@xtm-intl.com
www.xtm-intl.com

center staff. With the high cost of each and every customer incident, maintaining international help desks is labor-intensive, and therefore cost-intensive. Every call that is avoided reduces costs. We've seen MT result in call deflection rates from as low as 10% to as high as 90%, but even at a conservative average of 25%, this makes MT worth at least considering.

MT, however, is not attractive solely on budget. In fact, we have noticed that when we help customers save on one or more languages, they tend to shift those funds to adding new languages. MT helps deflect significant costs while at the same time increasing international sales, customer satisfaction, brand engagement, customer loyalty and even employee satisfaction.

Employee satisfaction is perhaps one of the more surprising benefits of MT for customer support, but far from resenting MT, many help desk employees appreciate not having to field repetitive calls on the same problem. They can leave MT

to solve the simple issues while they are on hand to solve the complex and more interesting ones.

Best-of-breed solutions

MT for customer support is not one size fits all. The needs of support organizations vary, and so does the engine that can meet those needs. Being comfortable with a variety of engines, we believe in using the best of breed for the individual context.

Some companies want their content translated in advance and ready to be accessed. Others want a widget to sit on top of forums and translate that information in real time with 24/7 availability. Others want a mixture of both, with perhaps post-editing on the top 20% of articles that are consulted every day. Still others may wish to engage their community in crowdsourcing – or crowd-improving – their translations. These different use case scenarios mean different MT approaches: rules-based (RBMT), statistical (SMT) or

hybrid. The important thing is to use the best-of-breed solution for whatever solution is best for your company.

People often ask us what engines we recommend. This is somewhat like asking what car you should buy. It all depends on what you need it for. If you need a car for running around the city, we might make a different recommendation than if you said you needed to haul two-by-fours.

It's the same with MT engines. You want to look at language pairs, content type, file formats and so on, because all these have a significant impact on MT quality. But for customer support, this question is somewhat simpler to answer, depending on whether the support needs to be real-time widget-based or in a searchable database.

The one invariable rule is that training the engine is essential. There is no off-the-shelf solution that we would recommend without training, whether it's RBMT, SMT or hybrid. Using a generic online engine without customization is

YOUR COMPLETELY INDEPENDENT TRANSLATION MANAGEMENT SOLUTION

"Now our translator can focus on localization and not recreating the document. The efficiencies we've gained have been as dramatic as moving from a typewriter to using Microsoft® Word."

Daniel Ng | **GPRO Technologies**

- Side-by-Side Translation Editor
- Translation Memory Suggestions
- Termbase Management
- Advanced Project Reporting and Statistics
- Unicode Support for Western and Eastern European, Double-byte Asian Languages
- Import/Export Multiple File Types, Including:
 - Adobe® FrameMaker® and InDesign®
 - Microsoft® Word and PowerPoint®
 - HTM/HTML, XML, XLIFF and More

Sales@MadCapSoftware.com | +1 858.320.0387 | MadCapSoftware.com

Copyright © 2012, MadCap Software, Inc., and its licensors. All rights reserved. MadCap Software, the MadCap Software logo, and MadCap Lingo are trademarks or registered trademarks of MadCap Software, Inc., in the United States and/or other countries. Other marks are the properties of their respective owners.

not advisable for customer-facing implementation: it's risky, potentially brand-eroding behavior.

On the other hand, only an online engine can give you 24/7 availability without the need for a huge infrastructure to support it. If this is your goal, Microsoft Translator Hub (also known as Bing) is a great solution because it can be trained on your data. In Lexcelera's internal studies we found that training improved the performance of the Microsoft Translator Hub by a full 20 BLEU points.

For documentation and online help, our studies have shown that a trained hybrid system such as SYSTRAN's performs best: quality is higher and terminology is respected, which means superior post-editing productivity when the goal is to reach fully-human quality.

To translate support content that is the most consulted – for example, the 10% to 20% of knowledge base articles that are consulted the most – we would recommend using a trained RBMT-based hybrid engine and post-editing that text lightly or fully, depending on quality needed and how often that content is accessed. If you're creating separate support sites in different languages, this high-quality documentation is a good starting point. If parts of the support site are meant to be dynamic, with a real-time translation widget like Microsoft's,

this high-quality post-edited content will provide valuable input as training material.

User-generated content

After training, the next most critical consideration in choosing an engine is the provenance of the content. One of the reasons a rules-based hybrid engine works so well with documentation and help files is that they tend to be correct grammatically. But you can't always say that for user-generated content.

Spelling mistakes and bad grammar aren't as much of a hindrance for Microsoft Translator Hub because it is an SMT engine (and doesn't need to "understand" the text as does a rules-based engine) and because it has already been trained on the contents of the whole web, flaws and all.

When the source text is user-generated, it's essential for an engine to have the ability to handle colloquial content, with the occasional spelling or grammatical error. RBMT engines don't have it. They need correct source language to make sense of a sentence. In a study conducted in our office by Fabrice Chabot, we found that while a trained SYSTRAN hybrid engine performed better on documentation than Microsoft Translator Hub, the opposite was true when the content was FAQs and forum content: two sentences out of three were "fully understandable"

when translated by Microsoft Translator Hub, while fewer than one out of three met the same criteria when the SYSTRAN hybrid was faced with error-rich content. The fact that the Microsoft engine already has a rich base training in user content is a huge advantage.

In the previously mentioned study, the Microsoft Translator Hub had been trained in two ways, and this distinction is crucial to understanding why it performed well. Unlike a generic Moses engine, for example, the Microsoft engine was already rich with more user-generated content than a company would typically have access to. The second training was conducted by Chabot, who augmented the base engine with enough customer-specific data such as translation memories to make sure that the engine preserved important terminology.

One of the users of the Microsoft Translator Hub who shares this dual training approach is Autodesk. After training, a widget was embedded into its customer support site. Today this widget offers real-time translation in 36 languages. They also use Microsoft's collaborative "improve this translation" solution; a dashboard behind it allows them to evaluate user contributions and approve them before they go live. The existing base training as well as the training on Autodesk data allows the company to coax higher performances from this engine.

According to Mirko Plitt, senior manager of language technologies at Autodesk, "To go into additional languages that are being left out, most companies can't turn to Moses because they don't have enough data in those languages to train an engine. It's not a technical question but a data question."

Although Autodesk has one of the most sophisticated Moses installations in the industry, the need for real-time translations also prompted the choice of Microsoft Translator Hub. Plitt says: "Microsoft understood a market need. Another reason we chose Microsoft is that it is available all the time. The alternative would have been to create an infrastructure available 99% of the time and we are just not set up for 24/7 support."

All of this points to the fact that customer support will be the next big wave in the translation and localization industry – though perhaps a tsunami might be a more apt description considering the masses of content needing to be translated. **M**

**audio
localizations**
in country, on time

London, Paris, Frankfurt, Milan, Rome, Madrid, Lisbon, Amsterdam,
Prague, Budapest, Warsaw, Zagreb, Istanbul, Athens, Stockholm,
Oslo, Copenhagen, Helsinki, Moscow, Cairo, Tel Aviv, Mexico City,
Buenos Aires, San Paulo, Beijing, Shanghai, Taipei, Seoul, Tokyo.

 www.binarisonori.com 