

**The Insider's Guide to
Customer Service Knowledge Bases**

Technical Guidelines and Best Practices for Answering Customers'
Questions Quickly, Accurately and Cost-Efficiently

© 2003 RightNow Technologies, Inc.

Contents

The High-ROI Knowledge Base	1
What is a Knowledge Base?	2
Critical Knowledge Base Attributes.....	3
Ease-of-use	3
Customer-driven content.....	4
Self-adjustment over time	4
Graphical content	4
“Pointers”.....	4
Knowledge “subscriptions”	4
The “Self-Learning” Knowledge Base.....	5
Step 1: Seeding	5
Step 2: Capture	5
Step 3: Organization.....	6
Step 4: Management`	7
Best Practices for Maximizing ROI	8
Applying one knowledge base across all customer communication channels.....	8
Optimized site design.....	8
Reporting and metrics	9
Educate and train	9
Conclusion: Getting Started	10
About RightNow Technologies	11

The High-ROI Knowledge Base

Great customer service is all about answering customers' questions quickly and accurately every time. An effective knowledge base can be extraordinarily useful in fulfilling this critical business requirement. In fact, many companies have achieved dramatic business results through the use of customer service knowledge bases. These results include:

- **Significant cost reductions.** Every time your customers find the answers to their questions on your web site, rather than calling or emailing you, you save money. Every email you avoid saves you around \$10. Every phone call you avoid saves you up to \$33.
- **Happier customers.** Online knowledge bases can give customers instant answers to their questions 24/7. Plus, by reducing call center and email workloads, the knowledge base improves phone and email response times. The combined effect is faster service and increased customer satisfaction.
- **Consistent answers across all channels.** Customers should always get the same information, regardless of how they contact your company. This is particularly true when legal liability or regulatory compliance is involved. Use of a single knowledge base across all channels ensures that consistent answers are given on the web, by email, on the phone and in chat.
- **Super-scalable capacity.** Companies that rely exclusively on call centers to answer customers' questions have to quickly add phone lines and customer service representatives to handle short-term "spikes" in demand—whether those spikes occur because of an unforeseen problem or are simply part of a normal seasonal business cycle. A web knowledge base, on the other hand, can scale to answer tens or hundreds of thousands of queries quickly and inexpensively.
- **Greater staff productivity.** Knowledge bases put huge amounts of information right at the fingertips of customer service representatives (CSRs). So even new trainees can be empowered to quickly answer a huge range of customer questions.
- **Increased revenue.** If you can answer customers' pre-sales questions quickly, you'll book more orders. This is true whether you're taking orders on the phone or doing business on the web. So having a knowledge base to deliver those answers has been proven to directly lead to higher revenues.

Happier customers, more streamlined customer service operations and greater revenue. These are the benefits that thousands of companies around the world are reaping today through the use of effective customer service knowledge bases. But the right technology and best practices are key to achieving these business benefits.

What is a Knowledge Base?

A knowledge base is an organized collection of information. It's useful to think of it as a library. The books in the library all contain useful information. Those information "units" are organized to make it easy to find the one you want. You can locate the book you need using the card catalog, which lists each book alphabetically by title, author, and subject. Someone searching for the book "Moby Dick" by Herman Melville will find the same entry if they look under the title "Moby Dick," under the author "Melville," or under a subject like "whaling."

Instead of books, the information units in a customer service knowledge base are answers to common questions. At many web sites, only the crudest sort of organization is applied to these answers—often referred to as "FAQs" (frequently asked questions). These FAQs may be grouped by category, but usually aren't indexed in any meaningful way. So searching for the specific answer you need is like going to a library to find "Moby Dick" and discovering that all the great American novels have been thrown together in one big bin. You'd have to search through the entire bin to find what you want.

A better solution is obviously to make it easy for users to find answers based on any conceptual "handle" they might want to use. This would be like sitting at a PC in the library and typing, "Where do I find that book about the big whale?" or "What's the name of the story about crazy Captain Ahab?"

The library analogy can be applied another way. If you walk into your local public library at the end of March or early April, you'll probably see all their books on tax preparation set out in a special seasonal display. That's because they know that lots of people are looking for those books at that particular time. To help them find those popular books immediately, they put them somewhere where they don't even have to initiate a search.

Customer service knowledge bases can do the same thing—putting answers to customers' "Top Ten" questions front and center. This speeds customer service and gets people into the habit of using the knowledge base.

Thus, a knowledge base can be much more than just a collection of information. It can be a well-organized, dynamic information resource that makes it easy for even the most novice user to quickly pinpoint whatever information he or she needs at the moment.

Critical Knowledge Base Attributes

Customer service—especially customer service over the web—places heavy demands on a knowledge base. Following is an overview of the key attributes that make a knowledge base effective for customer service.

Ease-of-use

The easier the knowledge base is to use, the more frequently customers will opt to get answers for themselves via the web rather than making a call or sending an email. Several attributes can make customer service knowledge bases especially easy to use:

Keyword search

Customers (as well as CSRs, if the knowledge base is being used internally) should be able to simply type in a single word or phrase to find the answers they seek. Keyword search results should be tuned so that they don't return too many results or too few results. Too many results slow down the process. Too few may mean that potentially useful items have been excluded.

Natural language query

Natural language searching takes into account the contextual elements of written language. Natural language searching can distinguish between queries about how to operate a given product and where to purchase it. This helps ensure the speed and accuracy of searches.

Related items

If an initial query is unsuccessful or returns too many knowledge items, the customer should not have to start from scratch. Instead, the next query should choose from answers that are closely related to the first. This way, each step in the search can be progressive.

"Top Ten" lists

Just as a library puts its tax books front and center in the weeks before April 15, the knowledge base should put certain knowledge items front and center whenever those items are in peak demand. This may be because of predictable seasonal issues (such as return policies in the weeks following Christmas) or unanticipated crises (such as recall of a defective product). In this way, the greatest possible number of customers find the answers they seek with a single mouse-click.

Easy escalation

Finally, it should be easy for customers who want to abandon their search and contact a CSR directly to do so—whether by email, phone or chat. Otherwise, they will be discouraged from attempting self-service again in the future. As will be explained below, these escalated direct contacts actually create a perfect opportunity to build the knowledge base so that customers with similar questions in the future can find what they need without assistance.

Customer-driven content

Some companies try to build their knowledge base by having subject-matter experts devote a lot of time and effort to authoring knowledge items. There are several problems with this conventional approach to content creation:

- 1) There is no guarantee that the answers created by internal authors will actually match up well with the questions that are most frequently on customers' minds.
- 2) Lots of upfront effort is required, as the authoring team tries to formulate answers to every conceivable customer question. This often bogs down knowledge base projects and delays time-to-benefit.

A powerful alternative is to allow interactions with customers to drive content. That is, each time a customer asks a new question, an associated knowledge item is added to the knowledge base. The "self-learning" technology that automates this process is described in the next section of this paper.

Self-adjustment over time

Customers' needs can change due to seasonal fluctuations, new product releases, or unforeseen events. To be effective, the knowledge base must respond to these changes over time. Knowledge bases that depend entirely on manual adjustment will be slow to respond. Often, the necessary changes won't be made at all. And the cost of managing the knowledge base will remain high. By contrast, a well-automated knowledge base can greatly ease ongoing maintenance. This is also explained in the next section.

Graphical content

A picture is worth a thousand words. This can be particularly true when it comes to answering certain types of questions—such as installation or repair instructions. That's why it's often a good idea to use diagrams, photos and other graphical content as knowledge items. A good customer service knowledge base must therefore support content other than pure text.

"Pointers"

Some questions may best be answered by content outside the knowledge base itself. Graphical content, for example, may already exist in another section of the web site. Or the solution to a technical issue might best be addressed by content on a third-party site. In such situations, the knowledge base may contain a hyperlink or other "pointer" rather than a complete answer itself.

Knowledge "subscriptions"

Some answers may be so dynamic and/or critical to customers that they will want to know when that answer changes. So it's good to enable customers to "subscribe" to a notification of any change in that answer. This proactive approach to customer service can head off problems before they occur.

The "Self-Learning" Knowledge Base

As noted in the preceding section, effective high-ROI knowledge bases are:

- 1) driven by customers' needs, and
- 2) require as little manual maintenance as possible.

This alignment with customer needs and automated fine-tuning over time is achieved through self-learning technology. Self-learning knowledge bases are characterized by a four-phase process: seed, capture, organize and manage. These phases place the customer at the center of the process. They also provide the automation necessary to save time and money.



Step 1: Seeding

The first step in developing a self-learning knowledge base is to create a knowledge "seed." A knowledge seed typically consists of 15-50 initial knowledge items of commonly asked questions and answers. These initial seed items must contain accurate and relevant information, but they do not have to be comprehensive. Most organizations can compile this knowledge seed from existing documents within a few hours. This knowledge seed is the starting point for a self-learning knowledge base.

Step 2: Capture

Once the knowledge base is seeded, the "capture" stage of a self-learning knowledge base takes place. During the capture stage, users search the knowledge base. If they can't find an appropriate answer, they submit their question via a web form, email, chat session or phone call. The question is then routed to an appropriate human respondent—such as a customer service representative or an appropriate subject-matter expert.

The human respondent then answers the question. In the process of doing so, the answer is captured for potential inclusion in the knowledge base. In most

companies, that answer is quickly reviewed by an editor to ensure accuracy and clarity of language.

This capture stage achieves two important goals:

- 1) **It creates new knowledge items based on user needs.** New items are automatically created whenever a user encounters a “gap,” which means that authors don’t have to guess where those gaps may be.
- 2) **It automates knowledge authoring.** Rather than requiring a separate laborious process, authoring is done as part of an existing business function.

This capture stage quickly fills any holes in a knowledge seed. With active use, a complete knowledge base is generally compiled within just a few weeks. The capture stage keeps working over time, filling any gaps that may emerge in even the largest, most mature knowledge bases as a result of the inevitable and continuous changes that occur in products, services, companies and markets.

Step 3: Organization

Seeding and capture alone would only result in a huge knowledge base with no intelligent organization. We have all experienced such knowledge bases, where a simple search produces hundreds or thousands of matches that are of little use for our immediate needs. That's why knowledge base organization is so vital.

Organization ensures that the information that has been historically most useful to users becomes the most accessible information for subsequent users. In this way, the most users find the most answers with the greatest ease.

Self-learning knowledge bases use a variety of techniques to categorize knowledge items and rank their usefulness. These techniques include:

Usage-based ranking

A good knowledge base will automatically re-rank knowledge items based on frequency of use—adding them or removing them from the “Top Ten” as appropriate. The knowledge base can also alert administrators about items that have not been accessed for a long time and/or that are no longer effectively answering customers’ questions. It’s often best to remove such items so they don’t clutter up the knowledge base and slow customers’ searches.

Implicit feedback

By tracking the path that users take through the knowledge base, a knowledge base can determine the effectiveness of its answers. For example, if viewing a particular knowledge item almost always results in the customer sending an email, it’s likely that the answer is particularly unclear or not useful. Or, if

after entering a given keyword, customers consistently complete their searches after viewing one specific answer rather than all the other answers returned by that keyword, then that answer should become the primary answer for that keyword. These types of adjustments help increase the total percentage of customers who quickly and successfully find the information they need online.

Explicit feedback

Another great way to monitor the usefulness of the knowledge base to your customers over time is to ask them. Short, "one-click" surveys can be presented to customers as soon as they complete a knowledge base session, asking them about their experience on a numerical scale. Based on their responses a variety of adjustments can be made to the knowledge base as a whole or individual knowledge items.

Although these techniques are technically complex, they yield simple and compelling benefits. In fact, self-service rates of 75%-95% are quite common with self-learning knowledge bases. That's because self-learning knowledge bases respond to every site visit—becoming even more effective every hour of the day.

Step 4: Management

The final self-learning step automatically ensures the timeliness of the knowledge base. This step forces information that does not continue to be useful off the top of the knowledge base so that it becomes less accessible and doesn't interfere with access to more important items.

Relevancy is achieved by automatically aging the data collected in step 3. The result is that more recent usefulness-ranking data (i.e. that which has been collected from more recent site visitors) is given greater weight. This aging process can be thought of as the "half-life" of the data. That is, data about a knowledge item ages in the same way that the potency of radioactive material decreases over time.

Information items that continue to be useful stay at the top of the knowledge base, while items that diminish in usefulness descend from the top and become less visible.

For example, an e-tailer's information about gift-wrapping would probably be very useful between Thanksgiving and Christmas, whereas information on how to return merchandise would be more useful during January and February. The management step performs this re-ranking automatically.

Best Practices for Maximizing ROI

While advanced knowledge base technology itself can be extremely useful in optimizing the quality and cost-efficiency of customer service, there are also a variety of best practices that have been proven to further enhance overall ROI for such deployments. These best practices include:

Applying one knowledge base across all customer communication channels

Customer service organizations often focus on knowledge bases as a means of delivering superior customer service on the web. However, for maximum ROI, it's best to apply a common knowledge base across all customer communication channels. That is, CSRs should be able to use the same knowledge base that customers use as they respond to incoming phone calls, emails and chat sessions.

Among the substantial business benefits that accrue from such cross-channel knowledge base implementation are:

- 1) Consistent answers across every channel. Customers should get the same answers to their questions regardless of how they contact a company. Using a common knowledge base ensures that this is the case.
- 2) Improved CSR productivity. CSRs can't always know everything about everything. But, with access to the knowledge base, they too can quickly find the current, most accurate answer to the customer's question. In fact, armed with a knowledge base, even new CSR hires can answer many questions like an experienced pro.
- 3) Efficient knowledge creation. Companies that don't use a common knowledge base across all channels have to author content separately for each type of use. This is wasteful and inefficient.
- 4) Risk and compliance. In some cases, the answers that CSRs give customers must be quite precise in order to avoid potential legal exposure and/or ensure regulatory compliance. By formulating these answers in advance and making them available verbatim to CSRs (as well as online), companies can ensure that they stay within the guidelines prescribed by their legal departments.

Of course, sometimes CSRs have to be given access to knowledge items that would be inappropriate to expose to customers. A good knowledge base should therefore provide for different portions of content to be made visible to different audiences.

Optimized site design

The business benefits gained through the online use of a customer service knowledge base are directly related to how well customers are guided to that knowledge base. The most successful companies place a link to their

knowledge base in a prominent spot on their home page. Some even go as far as to place their Top Ten questions right on the home page. This ensures that customers coming to the site with questions immediately begin interacting with the knowledge base.

Many companies have also found it useful to put links to the knowledge base at appropriate places throughout the site. This is especially important in the case of online retailers and other companies that sell via the web, since it allows customers to get answers to their questions without abandoning the site or the purchase process.

Reporting and metrics

Even with all the automation that today's knowledge base technology offers, it's still wise for customer service managers to continually review the reports that their knowledge base generates and take appropriate action. These reports include keyword use, item utilization and survey responses. By monitoring these reports, managers can continuously improve knowledge base performance—thereby steadily increasing both customer satisfaction and operational savings.

A particularly important statistic to track is sessions vs. escalations. This is one of the best measures of knowledge base effectiveness, since it gives a clear indication of how many customers did or did not find the answers they sought. By tracking this statistic, customer service managers can determine how much impact various fine-tuning measures are having on the customer service experience. This statistic can also be used to calculate ROI, since it provides a basis for quantifying the total number of phone calls and/or emails eliminated by online self-service.

Educate and train

While it shouldn't be necessary to train customers *how* to use the knowledge base, it's still a good idea to train them to *use* it. This can be done in a variety of ways. For example, the "hold" message in the contact center can suggest that customers use the knowledge base while they wait. CSRs may also politely point out when handling calls that the answers they provided were readily available on the company web site. They may even take the time to point out exactly where on the site similar answers could be found in the future.

Customers can also be led to relevant knowledge items when they attempt to send an email query from the site. This can be done by automatically "scanning" the text of their emails (in a process similar to a natural language query) and presenting them with existing knowledge base items. This promotes both the use of the knowledge base and reduces the total number of emails that CSRs have to respond to manually.

Conclusion: Getting Started

With the right technology and smart implementation, customer service knowledge bases can offer significant business benefits to companies who want to delight customers while keeping a lid on operational costs. Effective customer service knowledge bases deliver accurate, up-to-date answers to customers 24/7. They help CSRs become more efficient at answering phone calls and emails. They remove obstacles to online purchases.

Because today's knowledge bases can also grow quickly in response to live customer queries, they also make it easy to get started. Rather than committing to a huge content-authoring effort, companies can simply start with whatever content they already have on hand and build from there. The knowledge base itself will then take over to rapidly refine and expand content as necessary to fulfill customers' needs.

In fact, with such substantial benefits and such low pain-to-initiate, there's no reason why every company should not launch some sort of customer service knowledge base program. Customer service is, after all, a critical competitive differentiator in the marketplace. And the quality of that service is largely contingent upon the ease and speed with which customers can be supplied with the knowledge they want.

About RightNow Technologies

RightNow Technologies is a leading provider of on-demand software solutions designed to optimize customer service operations for businesses of all sizes. RightNow's comprehensive customer service solution features a self-learning knowledgebase that is designed to seamlessly support multiple communications channels including web, interactive voice, email, chat, telephone and proactive outbound email communications. RightNow offers its solutions through a multi-tenant, hosted on-demand model to reduce the cost and risk associated with deploying traditional enterprise customer relationship management, or CRM, software. RightNow also provides business process optimization and product tune-ups throughout the lifecycle of its client relationships. Founded in 1997, RightNow is headquartered in Bozeman, Montana, with offices in Europe and Asia. For further information visit www.rightnow.com.