



Knowledge Management Metrics

We all know that what gets measured gets done. Yet, many support centers collect no or only a smattering of knowledge management metrics. If you're serious about knowledge management, come along for a tour of knowledge management metrics.

Knowledge Creation

Measure productivity for the knowledge base just as you would measure case productivity. At a minimum, measure volumes of documents created, reviewed, and published. The next level up would be to also capture adherence to the

workflow targets you have for the knowledge base. So if you want to get documents reviewed within two days you should measure and report on the percentage of documents that are indeed reviewed within 2 days.

Here's a sample basic report (volumes only).

KB Productivity for Week Ending 27 May 2005

Name	New	Waiting for Review	Reviewed	Published/internal	Published/customer	Waiting for Fix
Anna	3	2	0	1	1	1
Bert	2	0	0	0	0	0
Carlos	0	0	0	X	0	0
Denise	X	X	X	X	X	X
Ellen	X	X	X	X	X	X
...						
Total	8	3	5	3	2	X

Knowledge Usage

Having lots of documents is wonderful but you also want to know whether they are helpful both for customers and staffers. Start with basic measurements: are documents hit (retrieved) in searches? Are they read? What are their ratings? How often are they linked to cases (assuming you have the ability to link cases and documents, which I highly recommend to implement in your next round of upgrades if you cannot do it today).

Note that document ratings are notoriously unreliable since only a minute proportion of customers bothers to fill out the ratings at all. A better strategy is to measure escalations from self-service to personal support. This requires some work on the web site so you can capture when customers who had been using the knowledge base then decide to log a case (best practice is to capture case logging within 24 hours of a self-service session – and of course this is a coarse measure since a customer may search on one topic, find the answer, and proceed to log a case on a different topic).

Here's a sample basic report for knowledge base usage by customers. I like to look at the most widely used documents (here, the most widely read) because they give you a good feel for popular (problematic) topics. You can of course look at the whole of the knowledge base, and least-popular documents are great too!

Top KB Documents for May 2002

Category	Number	Title	Hits	Read (%)	Usefulness Rating
Installation	125	Installation Troubleshooting Guide	320	300 (94%)	27 (9%)
Configuration	111	Adding a Frame	1004	280 (29%)	279 (100%)
Configuration	502	API Guide	X	X (%)	X (%)
...					

Knowledge Maintenance

Assuming your knowledge base is past infancy, the real issue is not creating documents or even getting customers to use the knowledge base, but rather to keep documents current. The metrics to support maintenance efforts can be very similar to the ones discussed under knowledge creation: volume of already-posted documents that were reviewed, volume of documents waiting to be reviewed under your obsolescence rules, and so on.

You can also use more exploratory metrics such as:

Documents with the lowest ratings (or the highest escalation rates to personal support), so they can be revised or deleted

So-called failed or unmatched searches: were customers looking for information that is missing? Did they use terminology that is

In this issue

- [Knowledge Management Metrics](#)
- [Integrate Problem Reporting with Knowledgebase Publishing](#)
- [Knowledge for Shaping the Future of Your Support Organization](#)
- [Can The PC King Excel In Services?](#)
- [Tech Marketing Is Back... and Brand Moves to Center Stage](#)

not the one used in the knowledge base document? Quickly scanning lists of failed searches is a surprisingly effective method to tackle this point.

And don't forget metrics basics

Treat knowledge metrics as seriously as you would any other support metrics. The rules are the same.

Garbage in, garbage out. If you're not sure your data is solid, think twice about using it for metrics. Be suspicious of any data that requires an extra step from the support staff to fill out: it may be filled out hastily and not accurately. Better do without than rely on incorrect information.

Don't be average. Averages may be easy to compute, but they obscure the extreme data points, which are often quite interesting. Measure achievements against targets rather than averages whenever possible.

Slice and dice. Unless your team is tiny you'll have to present different views to different folks. For instance, the owner of product X will want to see only the results for product X. A benefit of slice-and-dice metrics is that individual owners will be able to spot any errors quickly – and get them corrected. It's hard to spot errors in overall reports.

Create a dashboard. Select a few key metrics (10 is probably too many) and create a visual display with graphs and trend information. Within a few weeks you will find your eyes become trained to immediately spot unusual results.

Make metrics available immediately and automatically. Support people are busy. If the metrics don't get delivered automatically in their mailbox, chances are they will not take the time to go get them.

About the Author

Francoise Tourniaire is the founder and principal of FT Works, a consulting firm that helps technology companies create and grow their support operations. She's the author of *Managing Support Strategically* and *The 10 Commandments of Support Pricing*, both practical guides for support managers and executives. For more information, visit www.ftworks.com or call 650 559 9826.