

Does Your Knowledge Need a Makeover? Part 1

By Livia Wilson

To date, most knowledge environments haven't produced the results we hoped they would. Companies have bought well-developed software; implemented good business processes enabled by the software, but a year down the road, most still haven't realized the great ROI they thought they signed up for. After reviewing a few dozen knowledge-centered environments, we *consistently* found knowledge usage benefits degrade after the first nine months of the initial knowledge base (KB) deployment. In these studied environments, the top three trends were:

- Volume of content objects numbered in the thousands.
- People developed coping mechanisms to get around the problem of not being able to find relevant content.
- The problem resolution process stopped showing a reduction in time to resolution.

The problem is the ROI is not sustained after the "honeymoon period". In fact, some groups showed a growing reduction in productivity after the first nine months — even as they used the same knowledge.

For example, VeriSign and Legato have chosen to re-evaluate and re-configure their knowledge environments to re-gain and sustain knowledge ROI. These organizations, like others we worked with, experienced many of the symptoms listed in the next section. These symptoms grow from content being developed, reviewed, reused and published without effectively managing the "context" of the knowledge. It was believed if the knowledge contained the customers' wording of their need (i.e. the customers' context), then it would be findable and usable. However because the customer is seeking to resolve a need they are not able to articulate, there is a gap between what the customer searches for and what he/she in fact needs. That gap must be bridged with expert context.

What occurs in technical support organizations is varied levels of context and varied approaches to solving a problem. Those variations have to be mapped at a context level to guide the user to an effective conclusion.

When the context is mapped according to an expert perspective, such that it causes a logical resolution approach to emerge from the set of content, it cuts the resolution time to a fraction of what it is when people are just using their own reasoning to search content. And, when people search and don't find, they create more and perpetuate the problem.

When we started to work with companies on this issue, we found they had a high “hit” or reuse rate – but the majority of the content was not used more than once or twice. People searched and found lots of content which was irrelevant and they spent more time searching than analyzing the problem. Triageing or isolating the problem to be solved took almost 80% of the resolution time. Consequently, when we mapped the content according to the triage logic, we saved between 53% and 82% of the resolution time.

If your organization has a knowledgebase see if this could be your situation by checking all that apply to you and your organization.

[] Search but can't find

People know the KB content exists, but can't find what they're looking for. Users think they have to get the wording just right. Most also believe that the content should be better categorized or that filters should sort content about unrelated products out of the search results. They think searching is a skill.

[] No relevance ratings

Many of the relevance ratings on the returned search results are the same. The search results should show relevance ratings with one to three results with very high relevance and the others with significantly low relevance. In this case, the right content objects don't float up to the top with a higher rating, so people want to see more content selections.

[] 1:1 case to knowledge ratio

A high hit or reuse rate gives the impression that there's useful content. When you compare the number of content objects used with the feedback ratings, you can't tell whether the majority of the content actually solved the problems or if they were just caught in the search results. Comparing the number of objects hit on or linked, you find there are nearly as many used (i.e. linked or hit) content objects as there are cases/trouble tickets. For example, if 500 cases/trouble tickets are processed and 400 content objects used — then it's a problem. In a healthy support environment, the relationship between cases and the number of content objects for a give performance period should be such that 40% or less of content objects resolve 60% or more of the cases.

[] Ratio of reuse is flat, no 80/20 rule

The high hit/reuse rate isn't distributed across the content objects used. For example, 20% of the content objects used 80% of the time demonstrates a more natural rate of redundancy. The number of total reuses/number of content objects reused should correlate to the natural level of redundancy in the problem space. For example, if 8 out of 10 calls are about how to upgrade to the new release and I have 8 reuses, it should be 8 reuses for only one content object. If we have 8 reuses across 6 different content objects then customers are being advised on how to upgrade in 6 different ways. There may in fact be 6 different ways to upgrade, but then the

customer would need to augment their expression of need to distinguish the basis for knowing which of the 6 ways should apply to them (in which case there is no 1:1 redundancy between their situations).

[] Time to use

If you look at the time stamp for when a content object was used working cases, you see that the use/link/hit time is not within the timeframe of when the problem triage/isolation occurred. If the content is referenced after the fact, it's more likely to cause the resolution time to go up rather than down. It becomes a documentation step instead of a resolution step. The time stamp may indicate the content is not working.

[] Time to closure doesn't decrease with subsequent links/hits

When you compare the time to resolution between cases not using existing knowledge and those that did, you don't see a consistent decrease in resolution time. By learning from existing content, users should be able to resolve cases/issues faster and with more confidence — thereby promoting a more efficient closure. If the improvement trend by problem isn't clear, people aren't using the content the way it should be used.

[] Time to closure or relief is inconsistent

When comparing resolution time at each touch point and between people, exceptions occur; however, there should still be a consistent trend. The same problem should take relatively the same amount of time to solve, outside the exceptions. Before the KB, the variation in productivity between people in the same role is often seven to tenfold because of the difference in expertise and style. After the KB, the rate of productivity should be similar between people in the same role (i.e. people supporting the same technology, at the same level in the organization, relatively mature in their job).

[] Too much content

Knowledge bases can grow to thousands of content objects. When the content objects for a problem/product area are written so each has the same possibility of being found, (i.e. written "about" a topic without regard for how that topic relates to a customer experience) it makes searching during the resolution process (i.e. searching before we really know what we are looking for) ineffective. That many problems can't uniquely exist — only varied perceptions of a much fewer number of problems. When a product/problem area is managed in a CRM or ACD queue, the rate of content objects per queue should be under 400.

If you checked: 6-8 items, your need is EXTREME, 3-5 items your need is MODERATE, 0-2 items your need is SLIGHT. To learn how to address your need with Knowledge Normalization makeover, keep reading.

So what do I do now? Knowledge Normalization

Knowledge Normalization is the process of examining how knowledge works for the business and its customers at a systemic level. We want to understand how the content objects map to the problems we have to solve — and how they guide the user in the right direction to solve the problem. Through normalization, we create relevant troubleshooting logic within the content objects without having to script or hard wire content to drive users' ability to find and use the object. In this process, we work with subject matter experts to write troubleshooting guides within KB content objects to help analysts (or the customer him/herself) isolate the problem and solve it quickly. In most cases, the normalization process takes four to six weeks and creates relevance without having to go through each content object one at a time.

These estimates have come from projects conducted to normalize companies' KBs by eliminating the bloat of duplicate content and missing problem solving models built into searchable content. These companies experienced results within the first six weeks after normalizing their problem spaces. One of these leading companies was VeriSign. Madeline Brane, the lead program manager said,

“This is like an extreme makeover. It means people get the relevant resolution at the top of search results list and we don't have all that ugly content.”

At VeriSign, support professionals use problem-solving technology internally to solve customer issues and structure the resolution for reuse. Externally, customers use a problem solving engine where they can access resolution content through a guided search or an informational/unstructured search engine to accesses all types of content. Knowledge normalization enables all these methods of content access to work together intuitively and address customer demand by expertly guiding the customer to the most relevant content object to meet their need.

By uniquely structuring the content in the format of a troubleshooting guide, agents and customers alike are able to take a guided path to solving problems — where no such path existed within the content before. Most guidance resulted from another phone call or collaboration costing the customer and the service agent more time and money — and that path now exists in easily accessible content.

Not every company has a variety of search engines but most seem to share the challenges of creating relevant paths to the right content even with a single search engine. The knowledge normalization process paves the right path for customer demand so issues reach their conclusions directly and effectively without taking wrong turns.

Stay tuned for Part 2 of the series on Knowledge Normalization makeovers when it is published in the next issue. Now that you have recognized a potential need for Knowledge Normalization, the second article will define normalization further and discuss how to normalize content objects within your own KB.

Livia Wilson is a principal with OutSights, Inc. (OSI), a management consultation firm. OSI helps organizations implement new customer service programs and specializes in system design to deliver services more profitably, effectively, and leverages knowledge to benefit customer support organizations and their customers. The ability to translate complex service struggles into service solutions is an unusual and desired skill. This competitive edge allows OSI to empower clients with new knowledge, making them feel significantly more comfortable with the numerous service environment activities they manage.

For more information about Knowledge Normalization services and how to get started, please contact OutSights associate Farrell Dessert at 913-782-7336 or email fdessert@outsights.com.