

KNOWLEDGE MANAGEMENT STRATEGIES: EXPLORING THE LONG TAIL



In 2007, Chris Anderson, then Editor-in-Chief of *Wired*, wrote a seminal article entitled *The Long Tail*, in which he described a unique advantage held by big online retailers, such as Amazon, over traditional physical media outlets.

For bricks-and-mortar stores, it has never been economical to stock anything beyond a relatively thin slice of titles. Even large bookshops carry only a tiny fraction of the books in print. An average record store needs to sell two copies of a CD each year to pay the rent on the shelf space the CD occupies (and, as a result, even large retailers like Walmart might only carry a range of music corresponding to the top 40,000 tracks on the market).

In contrast, Anderson cited the usage statistics of Rhapsody, a streaming service which — at the time — carried 750,000 tracks. Regardless of whose numbers were studied, the retailers or the streaming service, there was a huge demand for a limited number of the most popular tracks, and very low demand for most others.

For Rhapsody, however, shelf space was not a consideration. A digital track's only footprint is the storage required to hold it, and storage is cheap. Interestingly, Rhapsody discovered there was a monthly audience not just for the top 40,000 tracks, but the top 400,000. Each month, more than half the tracks in its inventory found *somebody* wanting to listen to them.

Anderson named this effect the "Long Tail," after the distinctive shape of the power curve obtained by plotting interactions for each item.

In the digital workplace, one area where the Long Tail increasingly manifests itself is [Knowledge Management](#). The growth of SaaS has pushed up the average number of cloud services consumed by enterprises to over 900. The Internet of Things will bring tens of billions of new smart devices, in many previously unconnected areas of the business. That's going to be a lot of new stuff to support and maintain, and a lot of demand on a broad range of knowledge.

If usage of Knowledge Articles in a large and mature content library is charted, we typically see a similar power curve to the one described by Anderson. A small number of articles will sit on the left: accounting for the vast majority of views, citations, and re-uses. The vast bulk of the articles find themselves on the Long Tail of the power curve, but each is hardly touched. So is it worth bothering with them at all?

To answer that, we can look at the three rules defined in Anderson's article, for providing value from a Long Tail. Applying these three principles consistently and over time ensures that a solid set of long-tail knowledge is built up over time, and, importantly, that it provides value. As the digital workplace evolves, this broad content may be vital in helping your support organization to meet the challenge.

1) Make everything available

The Long Tail's key differentiator is its near-unlimited breadth of coverage. Anderson's article contrasted bricks-and-mortar DVD retailers with the then-embryonic Netflix. Having "broken the tyranny of physical space," it did not matter to Netflix how infrequent or widely-spread the viewers of a particular TV episode were, "only that some number of them exist, anywhere." With incidental costs cut to negligible levels, almost any level of demand, however low, makes it worthwhile hosting an obscure film or episode.

Likewise, for our Knowledge Repository to produce the benefits of a Long Tail effect, we need to ensure that content is available for the large breadth of obscure, rare, and quirky issues which we might need to support. It's important that it is there, however rarely each instance might be visited. It might be straightforward to manually identify the top 10 issues reported by users, but the next 1000 are an unscalable challenge. You can't be selective. As Anderson puts it: "In a Long Tail economy, it's more expensive to evaluate than to release. Just do it!" However, this requires some sharp focus on costs, something zeroed-in on in Anderson's second rule.

2) Cut costs

Anderson argued for differential charging at the consumer level — to pull customers "down the tail" with cheaper prices. For Knowledge Management, however, consumption costs are not as relevant as production costs, so this is where we must apply the rule.

A typical benchmark might put the cost of a first-line resolution at \$25, versus \$75 for the second line. If popular, a frequently used knowledge article gets used at the front-line a thousand times, and prevents an escalation each time, it is generating a theoretical savings — by itself — of \$50,000. It's worth putting significant effort into these articles. Embellish them with video. Promote them on the front page of the service portal.

Meanwhile, one thousand knowledge articles on the Long Tail, each used once, might deliver the same savings, but if we spend more than \$50 producing each of those articles, the savings is wiped

out (and that's before we account for the cost of any articles which are never re-used at all.)

An effective Long Tail strategy for [Knowledge Management](#), therefore, requires very lean costs in knowledge production. It's a high-volume, low-margin process. This is where frameworks like Knowledge Centered Support (KCS) are attractive: KCS looks to derive knowledge content from day-to-day support processes, rather than from a separate, unaligned task of knowledge creation.

3) Help me find it

Anderson's article highlighted the story of *Touching the Void*, a book about a disastrous mountain climbing expedition in South America, written in 1988 by Joe Simpson. After modest success, it was nearly forgotten and almost out of print when, in 1997, Jon Krakauer published *Into Thin Air*, an account of a tragic series of events on Mount Everest. Krakauer's book was a bestseller, but something strange started to happen. Online sales of *Touching the Void* began to soar. The publisher rushed out a new edition to meet the sudden demand.

The reason for this was, primarily, Amazon recommendations. The online retailer's algorithms were linking purchasers of Krakauer's book to Simpson's *Touching the Void*, having lurked for years as just another title on Amazon's Long Tail, became hugely relevant due to an algorithm.

Users of Knowledge Management are typically under time pressure. One service desk agent, in 2014, describes the "gamble" that he faced on every customer interaction. Knowledge can help drive better first-time fix performance, but that was only one of his key targets. The other was average handling time. If agents feel it will take them a long time to find knowledge, or even worse, that they might search in vain, they will worry more about the time target.

To benefit from that Long Tail, therefore, we need fast, preferably automated identification of the right Knowledge Article from the wider bulk of content. If we can't find the right content, the Knowledge Base risks becoming a WORN (Write Once, Read Never) technology.

This rapid identification has been a big focus of our innovation with Remedy: the Smart IT UX finds knowledge automatically, in context. The agent no longer needs to gamble. Learn about the new [Remedy with Smart IT](#).

Read the rest of our series on Knowledge Management:

[Knowledge Centered Support](#)

[Importance of Knowledge Management for the Digital Enterprise](#)