



How AI Moves Knowledge Workers Up the Value Chain

For most professional services organizations, especially those who are in IT Consulting, Management Consulting, or IT Staffing, responding to Requests for Proposals (RFPs) are the lifeblood of your business. Price alone is never enough to win a bid. A critical element of any submission involves your ability to convey your capabilities and your differentiation. You must be perceived as the only logical and viable vendor capable of delivering on the stated requirements contained within the Request for Proposal (RFP), and the unstated requirements hidden throughout the RFP. You need to convincingly present your qualifications by articulating, in a highly tailored manner:

- The work you've done
- The people who did the work, and
- The way the work was done

This presents you with a challenge. More specifically, the knowledge workers within your organization that are tasked with pulling this content together have quite the challenge. They must present the information in a creative, compelling, and highly tailored manner. Your knowledge workers are constrained by the many versions and locations of content, and the various formats of the source material needed to write your submission.

Let's review the facts.

The work you've done — These are projects your company has completed over its lifetime. The details of these projects are contained in project records that identify the client, timeframe, scope, value and resources. If you've been in business for a while, this represents a lot of records in multiple formats that are distributed across numerous storage mediums. There's no standard for managing this content, which is very difficult to quickly and intelligently maintain, discover and repurpose. Your knowledge workers are expected to find and identify the right records, quickly, on their own.

The people who did the work — Let's just agree that this content is best reflected in individual resumes which are project-based. Each resume has been tailored numerous times for specific project opportunities. For one resource, can you accurately identify the master version of their resume, or which versions were used for specific projects? Do you know which version is best to use for a specific bid? Your knowledge workers are expected to know which resume is the original, which are the alternate versions, and how to choose the best one.

The way the work was done — These are the methodologies, frameworks, case studies and marketing collateral that describe your company's capabilities. This is your differentiation. This sets you apart. You expect your knowledge workers to know this stuff, and to be actively authoring any content to fill the gaps.



Here is the sad reality. You've got content all over the place. In order for your knowledge workers to find your stuff, you have to tag your stuff - but nobody tags their stuff. Those who attempt to tag their stuff frequently get it wrong because a) they didn't author it and b) they don't understand the context where it was used.

Your knowledge workers need help if you want to tame this beast. Understanding how they work will identify the tools they need to succeed tomorrow.

Understanding the Knowledge Worker

Peter Drucker invented the term "knowledge worker" in 1964. According to a McKinsey article published in 2013, the number of individuals employed to produce or analyze ideas or information has grown to over 200 million people worldwide. They span all industries but have at least two things in common, which directly relate to the preparation and submission of a winning RFP response:

1. Knowledge workers are goal oriented and not process-determinative
2. Knowledge workers make decisions under demanding time pressures

Who are the typical knowledge workers we're talking about here? When it comes to creating an RFP response, we are talking about your proposal writing professionals, your company's executives, your consultants, and your business development and marketing staff. In that context, the above two points have never been more true.

Think about the process of preparing an RFP response. You have a deadline. You have competing tasks and obligations that dramatically impact the time you can allocate to the authorship of your submission. Decisions need to be made quickly. This is the stress your knowledge workers are faced with. Compounding the issue, they are suffering from information overload. What they need is the right information to do their job: the resumes, projects, references, processes, frameworks, case studies, and collateral to make good decisions and produce the best submission, in the shortest time possible.

You have a deadline. You have competing tasks and obligations that dramatically impact the time you can allocate to the authorship of your submission. Decisions need to be made quickly.

These people play a critical role in your success. You don't have a choice. You must provide them with the core requirements they need to perform their work. This means a flexible, information-centric, technology-enabled work environment so that knowledge workers can focus on their end goals.

Knowledge Worker Technology Empowerment

In every organization, regardless of size, there is a collection of technology components that comprise the toolset of the knowledge worker. From simple office productivity applications, to online resources and databases, to in-house content repositories, the knowledge worker has many tools in their tool-bag. The challenge with this is multi-faceted:

- What works for one knowledge worker does not necessarily work for another
- Not every knowledge worker is familiar with the tools at their disposal, which often results in the least capable tools being utilized as they are the most understood by the broadest group of contributors
- Tools and processes become very personal and connected to individual knowledge workers
- Every knowledge worker develops tribal knowledge over time and they eventually learn what content to use, when to use it, and where to find it. Typically, this knowledge is not often formally shared
- When individual knowledge workers leave the company, the organization loses their technology and tool familiarity and is dramatically impacted

These issues are not new. The cost and impact of these issues is real and quantifiable.

Enterprise Content Management

Over the last few decades, Enterprise Content Management (ECM) systems were adopted by many organizations to support the knowledge worker. They provided a storage platform for the collection of documents and data used by knowledge workers. The issue is that ECM systems cannot structure content without labor-intensive manual efforts. Further, because ECM systems are document-centric, they largely lack any process or workflow capabilities beyond providing document lifecycle management.

Business Process Management

Knowledge workers need predictable, repeatable processes to achieve business objectives, namely, sourcing the information necessary to prepare a winning RFP response. As a result, this led to the creation and adoption of Business Process Management (BPM) platforms as a layer above ECM. This allows you to measure and control operational processes as well



as apply IT solutions directly to those processes. You gain visibility in the ongoing work via scheduling, tracking and monitoring. For this to succeed you need in-depth documentation and planning. This requires a large degree of predictability. As a result, the majority of business processes that are modeled in BPM systems are linear decision trees used to direct the actions and the decisions of the knowledge workers.

In the world of preparing proposal submissions, we can all agree that the only predictable thing is the objective of creating a winning bid. The processes that are used in accomplishing that goal are anything but predictable. BPM systems lack the flexibility needed by goal-oriented knowledge workers.

Case Management Frameworks

To understand knowledge workers' needs, we need to recognize how they use, generate and manage information. Typically, they must find, identify, weigh and review information before they can reach necessary decision points, communicate those decisions and the supporting information to others, and then generate the required content deliverables. This isn't easy when all of your source material is distributed across locational silos, in a variety of formats, and require an assortment of applications to engage it. Before the advent of technology enablement, this information was stored in a centralized location in folders, or files, or cases. Today, we've virtualized this construct in the form of Case Management.

Case Management Frameworks (CMF) sit on top of ECM and BPM in the knowledge worker technology stack. CMFs are used to assist in the creation of custom, non-linear processes to match dynamic business objectives.

Smart Process Applications

The highest level in the knowledge worker technology stack are Smart Process Applications (SPAs). They're the glue in the CMF process because they connect the industry-specific business activities to cases. Extremely relevant in the bid creation, SPAs support processes that are people-intensive, highly variable, loosely structured and subject to frequent change. SPAs bring it all together, including:

- Content capture, output and management
- BPM tools to execute the required tasks, providing visibility and accountability
- Structuring and searching tools to allow knowledge workers to educate themselves on the cases
- Analytical tools for critical decision making
- Collaboration features for the actual content creation process

Looking at the stack you'd think knowledge workers have what they need. ECM systems provide access to content. BPM platforms provide repeatable workflows. CMF enables non-linear knowledge work, and SPAs provide the ability to organize tasks into a fluid workflow. Yet, despite appearances, the stack remains incomplete:

- Information is still overwhelming
- Time demands are still intense
- It's still too hard and too manually-intensive to analyze documents and uncover the answers you need

If you're going to empower your knowledge workers to create amazing RFP responses, you need to provide them the intelligence and insight that is necessary to find the content they require, in an easy and time-effective manner, so that they can prepare a winning proposal.



Machine Learning and Artificial Intelligence Changes the Game

As mentioned at the start of this paper, there are three rules to content discovery:

1. In order to find your content, you have to tag your content
2. Nobody tags their content
3. Those who attempt to tag the content frequently get it wrong because they didn't create the content and don't understand the context in which it was created

These truths make it hard for your knowledge worker to find the content that is necessary to assemble a winning submission. It gets worse.

Advocates of an ever-evolving technology stack will suggest you need to incorporate automation capabilities into your proposal response process. The issue is that automation drives you down a linear, binary path. However, no two RFPs are the same and your business is in a constant state of transformation. In part, that's why SPAs were introduced to the stack. What you need to focus on is acceleration. It's driven by user interactions, adapting minute-by-minute to ever-changing conditions.

What if we could use Artificial Intelligence (AI) to interpret both the intent of the RFP criteria as well as the inventory of content and resources we already have on file? Our knowledge workers could then quickly find the right content, with the right nuances, required by the RFP. AI would mean the end of tagging and classification to make content discoverable and repurposable. With AI, we would have:

- No more manual classification and tagging of content
- No more libraries of questions and answers
- No taxonomies with terms and synonyms to maintain

This gives you instant access to content that is accurate, relevant and dynamically linked to the thread that defines its original and complete context.

reDock for Bids' AI creates a semantic network that defines the context of content and the meaning behind the content with every user interaction. reDock for Bids breaks the content down into contextual, usable segments (all linked back to their source), then it goes further. reDock for Bids extracts semantic and contextual meaning from requirements (criteria) and finds matching content (Projects, Resources, Methodologies, etc.). reDock for Bids learns through interaction, recognizing patterns of use and relevance, as well as formats. It also learns as processes evolve, adjusting to changing conditions and influences to suggest highly relevant actions.

Looking Under the Hood

How well a system copes with search and text analysis largely depends on how well the system understands the text. Understanding requires linguistic analysis, which enables computers to work with meaning rather than with character strings. Even more importantly, true understanding requires knowledge of why these words were written in the first place and what concept or object they refer to: was it in response to a requirement, as a corporate reference, or as a description of the services your company can offer? Once you understand the contextual meaning, you can deliver relevant search results which is exactly what the knowledge worker needs. reDock for Bids parses to detect the meanings of the words (semantics) and establish how they relate to one another (syntax) in light of where they were found (context).



What we're doing is not entirely new. Semantic analysis techniques, such as Latent Semantic Analysis (LSA) and Support Vector Machines (SVM), have been available for years. The issue is how these existing techniques fall short for your requirements. These approaches can tell you what words are on a particular page, where they are located, and what other words often appear near them in context. However, they cannot understand what those words mean, why they were written, and what they refer to. If the proposal process was a predictable and repeatable process, such as the ECM and BPM layers in the technology stack, then LSA and SVM would work. When you want to action information within unpredictable processes, such as those that occur at the CMF and SPA layers of the stack, the mathematical modeling of language shows its limitations.

Said another way, knowledge workers need context. It's that simple. They understand information by mapping it to concepts and objects, not through statistical correlations. reDock for Bids' Artificial Intelligence (AI) analyzes and uncovers patterns in the reference to those concepts and objects and builds them into stories to find intelligence and insights to make decisions. This is all done through language.

Knowledge workers need context.

To optimize the knowledge worker's technology stack you need intellectual property that can map language to concepts and objects. reDock for Bids' AI does this. Our comprehensive text analytics accurately identifies contextual segments and automatically extracts the underlying business objects. As you consider the vast variety of documents your knowledge workers have access to, along with the varying document formats and styles of writing, you can start to appreciate the power of reDock for Bids' AI. This includes co-reference to the same business objects that span segments and documents. You can extract the meaning needed to empower knowledge workers within domain-specific vocabularies and linguistic usage, without the need to construct complex rule sets or large statistical training sets. Our AI does all of this automatically. It analyzes incoming content in the background after it is captured to accelerate your knowledge workers' time-to-discovery and time-to-decision. In short, our AI provides you with the intelligence and insight necessary to do your job fast and effectively. Smart Process Applications need to provide knowledge workers with the right information, in the right tasks, at the right time — exactly, or they'll fail. The ability to do that can only come from a thorough understanding of the content so that it can be applied appropriately. Our AI provides the needed understanding for SPAs from the content itself. We structure the unstructured data so that your knowledge workers don't have to.



Looking to the Future

As discussed at the start of this paper, there were over 200 million knowledge workers globally as of 2013. That number has risen dramatically since then as the global economy continues its transition into a knowledge-based workforce. While we have focused on the utilization of our AI capabilities as part of our solution to the proposal challenge, it's clear it does not need to be limited to such a narrow application.

As we reviewed the knowledge worker technology stack, we focused on the real-time tracking of content, resources, and collateral that was related to the RFP response process. In most organizations, enterprise solutions such as Customer Relationship Management, Enterprise Resource Planning, Marketing Automation, Manufacturing Resource Planning, Sales Automation, Accounting, etc., are utilized daily to run the corporate operations. All of these software solutions are implemented and utilized by knowledge workers. The sheer

volume of structured and unstructured data flowing within and between these applications is staggering. Entire industries focused on middleware and integration have been established to ensure documents are exchanged to facilitate various information sharing requirements. Business Intelligence and Data Warehousing is a massive industry that has morphed into an even bigger market, now called Big Data, whose sole objective is to distill insights and correlations that would be otherwise hidden from us. It's all about making decisions, and recognizing new opportunities, with information contained within the enterprise but not easily discerned or discovered by your knowledge workers.

If all of the processes related to these solutions were static, predictable, linear and repeatable then current technology stacks could easily handle the exchange of data. However these processes, just like our RFP response processes, are anything but linear and static. They are dynamic. They are unpredictable. They are real life.

So, how does an Accountant and a Plant Manager and a VP of Marketing work together when their vocabulary, processes, forms, content, and reports (and potentially their spoken language) are so dramatically different? reDock for Bids is there to bridge the gap. reDock for Bids will discern the content, the context, and the intent as we continue to unite knowledge workers across entire enterprises, bringing silos of AI information together in one powerful solution.

reDock for Bids leads the way with massive productivity gains.

Most RFP response solution vendors cite a 40% productivity gain. At reDock, we consider 40% to be abject failure.

It's not magic, it's just math. Artificial Intelligence allows your organization to reimagine your proposal process today and your entire enterprise tomorrow.



Pierre-Olivier Charlebois
CEO & Founder
reDock Inc.

Graduating from McGill University in 2005 with an Honors Bachelor Degree in Electrical Engineering, Pierre-Olivier immersed himself in the science of Machine Learning and AI – first as a researcher in brain-computer interfaces at Thought Technologies to help children with ADD, and then as a Gameplay Programmer at international gaming author Ubisoft.

In his role as CEO of reDock, Pierre-Olivier is directly responsible for setting reDock's strategic vision and spearheading the Management Team while driving the Development organization.



reDock Inc. believes that truly great software should free people so they spend more time on the optimal use of their talents. Our software leverages Artificial Intelligence (AI) to extract meaningful corporate content into segments that can be repurposed to accelerate document creation and deliver winning business outcomes.

Contact reDock at 888-338-3448 x1 or email info@redock.com

