Cover Sheet

Responding to RA number: DARPA-RA-11-52

Title: **Participatory Sensemaking about Real-Life Cyber-Security Stories with Rakontu 2.0**

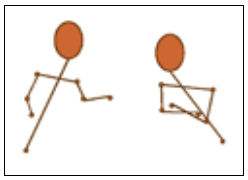

Organization name: Cynthia F. Kurtz (dba Kurtz-Fernhout Software)

Type of business: Other small business (woman-owned sole-proprietorship)

Principal investigator: Cynthia F. Kurtz, CEO

Other team members: Paul D. Fernhout, CTO, Kurtz-Fernhout Software

Technical contact: Cynthia F. Kurtz
Address: [redacted] Edinburg, NY 12134-5211
Telephone: [redacted]
Email: kurtz@kurtz-fernhout.com
Fax: [redacted]

Administrative contact: Paul D. Fernhout
Address: [redacted] Edinburg, NY 12134-5211
Telephone: [redacted]
Email: pdfernhout@kurtz-fernhout.com
Fax: [redacted]

Place of performance: [redacted] Edinburg, NY 12134-5211
Period of performance: Eight months, such as October 1, 2011 to June 1, 2012

Sub-proposer information: N/A

Proposal validity period: 122 days
Proposal valid from: September 14, 2011

Table of Contents

Executive summary

Cyber-security professionals need *examples* of security successes and failures to make inferences about threats and opportunities. Examples support the emergence of broad strategies and collective values that support critical decision making in times of need. Collective structured dialogue and analysis tools applied to thousands of real-life non-fiction cyber-security example stories (such as in RISKS Digest) could help professionals make sense of complex security issues. We would like to create essentially **an advanced distributed RISKS Digest-inspired system**. To do this we would like to improve our open-source social media platform (Rakontu) to more fully support Participatory Narrative Inquiry, a sensemaking approach based on stories of real-life experience. We would like to do this in an agile way with the feedback of cyber-security professionals (both novice and experienced) who gain real benefits from using Rakontu for collective sensemaking.

The design of Rakontu is based on work by Cynthia Kurtz (including at IBM Research and Cognitive Edge) on systems like DARPA's Genoa I and II projects and Singapore's RAHS project, as well as on narrative projects for dozens of government and corporate clients. In a nutshell, Rakontu helps groups of people collect stories *and relevant metadata about stories* in a way that mimics the flow of storytelling in natural conversation. By adopting a conversational approach, Rakontu both engages people in reflecting on their experiences and reduces tedious after-the-fact hand-coding of meta-data by third parties to the story. Rakontu then helps a community make use of their collected stories to discover valuable insights through group sensemaking. The fact that the same people *tell* the stories and *make use* of them makes Rakontu participatory. In this environment the value of tacit knowledge held by experienced professionals can be more easily recognized.

We have learned much from building the first version of Rakontu. We are looking for application areas to which we can apply Rakontu and gather feedback as we further develop it. And to be frank, we are looking for funding to continue work on it. It seems to us that Rakontu has much to offer the field of cyber-security. Though many potentially valuable stories exist in collections (such as the the otherwise excellent and informative RISKS Digest), they tend to be stored in barely organized “heaps." This limits their utility in supporting cyber-security professionals as they look for potentially subtle patterns across their combined experience.

To our knowledge, no one has done exactly what we propose here. We are not cyber-security specialists, so we may be unaware of all efforts in this area. However, we have seen many organizations gain traction on intractable problems through the use of participatory narrative methods, and we are confident that Rakontu can help cyber-security professionals achieve worthwhile benefits.

In the proposed project we would improve Rakontu by adding a real-time sensemaking workshop capability, integrating our data analysis software (NarraCat), and making other improvements to the database architecture and implementation. We would seed a Rakontu workspace with cyber-security stories available on the internet (like on comp.risks) and create a working, expandable proof of the concept that applying narrative methods to security issues can provide value to cyber-security professionals. We expect the project to take eight months and cost US$199,840. Our aim would be to start a snowball rolling to the point where cyber-security professionals begin to adopt the software for sensemaking about their own stories.

Conceptual screenshot mockups and related figures

The stories in my projects for clients

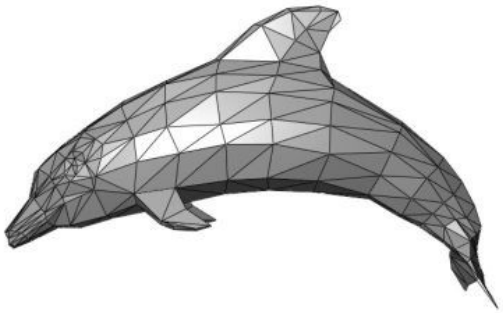

are connected in **webs.**
They **jump** into patterns and reveal insights.

The stories I see people sharing on-line

are piled in **heaps.**
They lie still and **obscure** insights.

Rakontu screens: patterns

Look at	*with*	*filtered by*	*question*	*answers*
Choices **Scales** Texts	Choices **Scales** Texts	**Choices** Scales Texts	Storyteller: Age Storyteller: Profession **Storyteller: Position on dam**	for undecided **against**

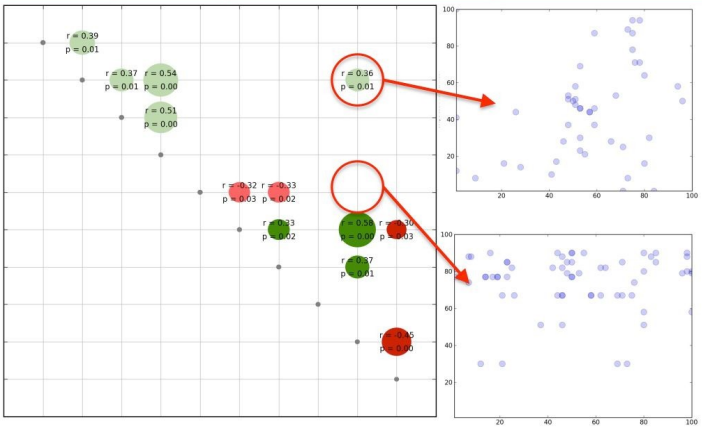

People can discover trends in their answers to questions about stories.

Notes

Do you see how there is a correlation between memorability and irritation, but only among those against the dam project.

Rakontu screens: café

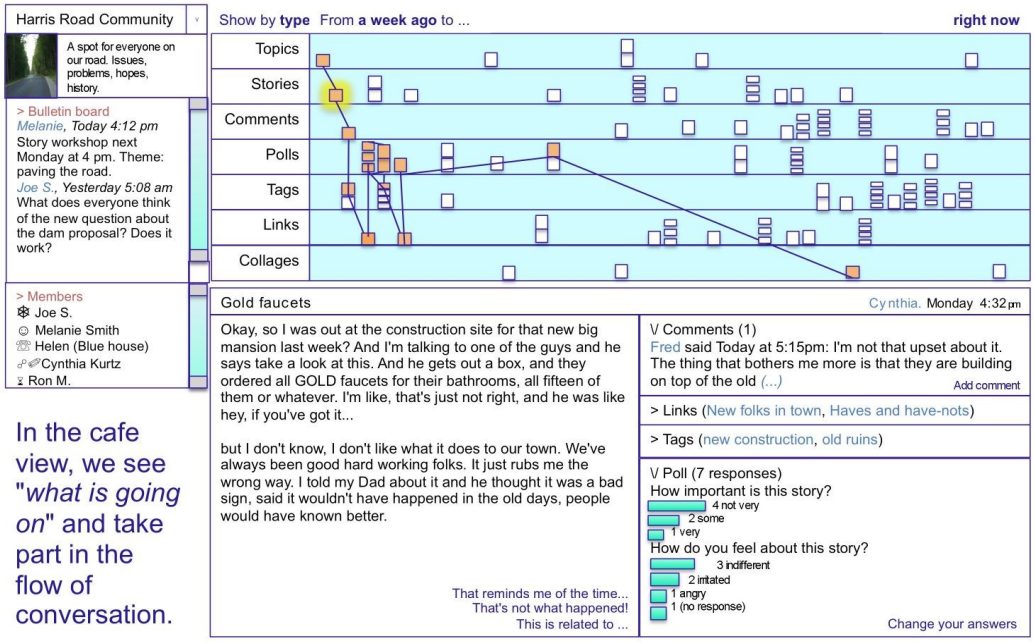

In the cafe view, we see "*what is going on*" and take part in the flow of conversation.

Rakontu screens: library

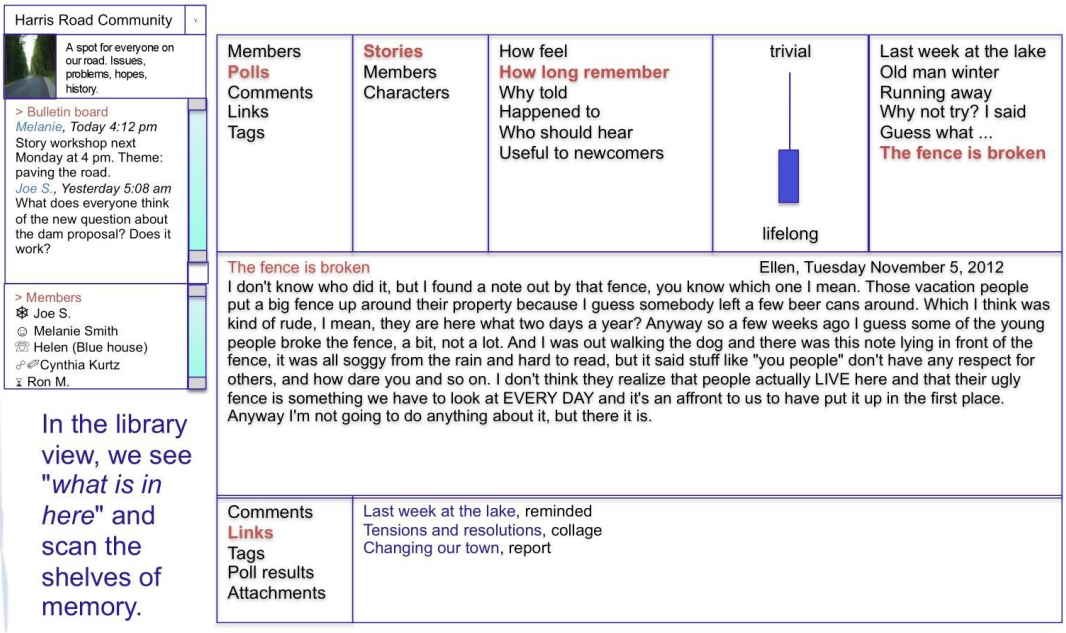

In the library view, we see "*what is in here*" and scan the shelves of memory.

3. Rakontu is about ...

many motivations to contribute!

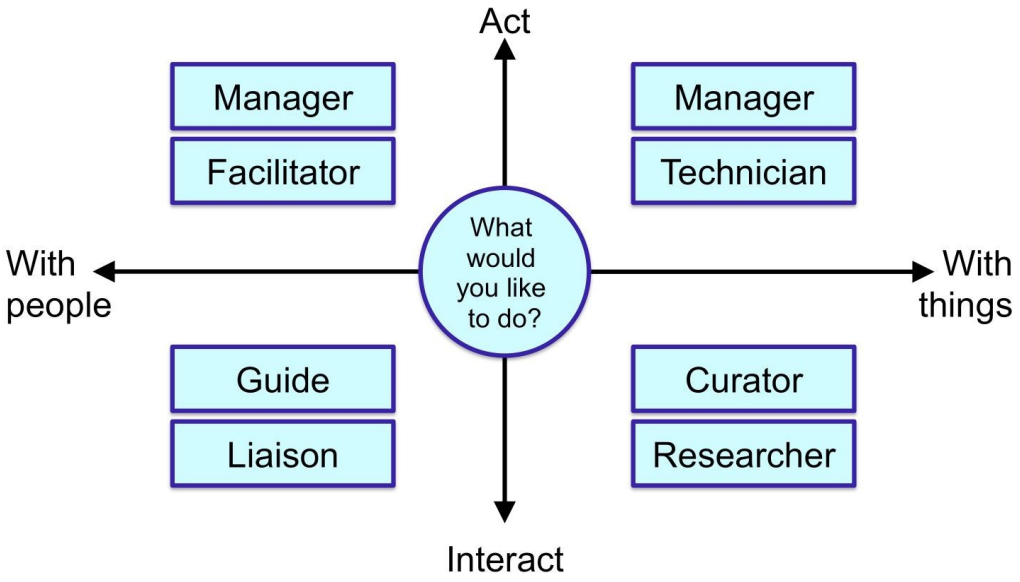

(this is based on Bartle 1996)

Rakontu screens: roles

Guide

Messages Requests **Topics**	Subjects Plans **Memory joggers** Reports Ideas	✓ Quotes from the old folks SillyPerson 1/2/2012 12 responses ★ **Photos of the old cemetery** **Joan M.** **12/2/2011** **0 responses** ✓ What the bridge used to look like Helen2 11/5/2011 1 response ✓ Remember the one-room schoolhouse? Wanda 6/9/2011 33 responses ✓ Report from August workshop erik 6/8/2011 9 responses

Curator

Show	*with*		
Topics **Stories** Collages Entries Workshops Members Characters	**no** few many duplicate conflicting varying unvarying	**tags** comments links poll results texts versions	Last week at the lake Old man winter **Running away** Why not try? I said Guess what ... The fence is broken

Each role has its own "place to go."

Manager

Members Appearance **Settings** Flags Messages	Invitations **Roles** Attachments	**Guide** Curator Liaison Facilitator Researcher Manager	The **guide** role is: agreed Text to be read when member takes on role: Guides answer questions, encourage people to tell stories, and maintain the vitality of the member community. Guides agree to receive messages so they can answer questions.

Rakontu screens: workshop

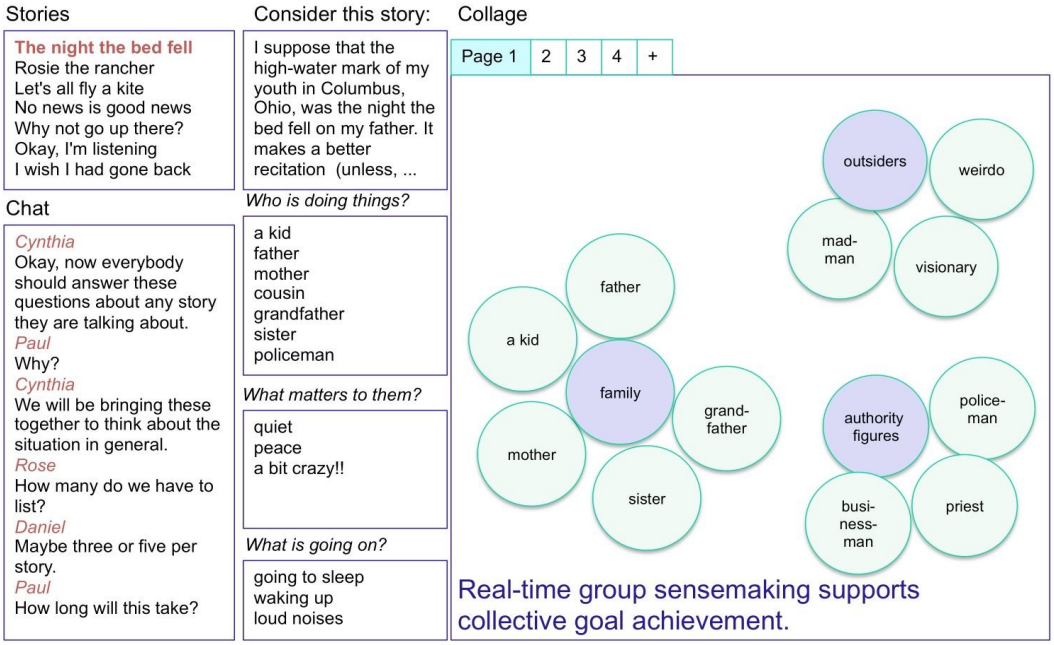

Rakontu modularity

We envision building the infrastructure of Rakontu in such a way that it can support not just narrative sensemaking, but *many* group activities based on various approaches, such as:

- Agent-based Exploratory Modeling and Simulation Games
- Appreciative Inquiry
- Comprehensive Anticipatory Design Science
- Critical Decision Method
- Delphi Method
- Enterprise Canvas
- Medicine Wheel
- Most Significant Change
- Scenario Planning and Horizon Scanning
- Structured Argumentation
- Structured Dialogic Design
- SWOT Analysis (Strengths, Weaknesses, Opportunities, Threats)
- System Dynamics
- Systems Thinking
- Tipu Ake ki te Ora Leadership Model

... and more!

Technical Description

Conceptual Background on Narrative Sensemaking

Sharing stories of real-life experience is essential to human learning and the collective refinement of best practices for specific task domains. One such story might be “I installed Windows XP on a new PC and it was compromised through the internet on a default open port before it could download the first service pack.” This proposal is about making the process of sharing stories easier in the domain of cyber-security through the use of our system for real-life experiential story exchange called Rakontu ("tell a story" in Esperanto).

Experiences, whether direct or indirect, are the basis on which humans inductively come to appreciate certain task-relevant values and the wisdom of broad strategies and habits of practice that support those values. Science and engineering relies on *semantic* reasoning, which codifies these experiences as explicit semantic information. Semantic information is generally expressed as facts, concepts, and links; e.g., “ A software port is a virtual/logical data connection that can be used by programs to exchange data directly; software ports can be a vector of intrusion.”

By contrast, *narrative* reasoning expresses experiences as narrative information, or stories that involve sequences of actions and their experienced consequences related to goals in context. The human brain has faculties for working with both semantic and narrative information and uses them in complementary ways. Narrative information often functions to organize and prioritize sometimes overwhelming semantic information (e.g. “Open ports in a default install? Why should I care when I need to get my work done right now?”).

Those who study military conflicts have long studied real-life stories of battles, campaigns, and wars to understand deeper strategic issues about security, even as they generally also know a lot about specific tactics and the operation and limits of specific military equipment. Nothing can replace real-life experience, but we can make the most of the experiences we have or know about. Learning by reading about other people's mistakes and successes, or learning by reflecting more deeply on our own experiences, is cheaper and faster than learning the hard way through new mistakes.

As a relatively new field, cyber-security could benefit from more of such learning. A cyber-security professional may master a variety of tools and concepts (semantic information), but he or she may still lack a bigger picture in relation to an organization's goals and priorities. This lack of contextual awareness makes it harder to know which tool (from a vast conceptual toolbox) to use and in what order to minimize an attack surface, to respond to an active intrusion, or to take advantage of a broad opportunity to improve security. Stories provide an experiential framework on which to build a sound contextual awareness.

What do we mean by real-life stories? When we speak of creating software to support story exchange, people often believe we intend to help people create polished Hollywood fiction, corporate advertising, or political propaganda. The narrative literature itself can be confusing because “storytelling” can be used to mean anything from recounting what actually happened (non-fiction), to making up a story to inform, entertain, or persuade (fiction), to performing a presentation of a fictional or non-fictional story (performance). Polished, purposeful,

performed stories are not what Rakontu is about. Rakontu is about helping people share their real-life, non-fiction, raw, true experiences so they can increase their collective ability to make sense of conditions they encounter and make the best decisions they can in context.

We also want to head off another misunderstanding we often encounter: that ordinary people are not storytellers. Many of the people we talk to define stories as only the polished fictional items that professional Hollywood storytellers craft and present, or they compare themselves to the best storytellers they know. They do not define stories as the rough and ready, often halting, things they say to their acquaintances about their own lives (e.g. "Norton quarantined my whole mailbox because of a virus in an attachment and it took two hours to figure out what happened and how to fix it; sorry I did not get the report to you sooner"). This is analogous to people believing they are not musical because they have not spent ten thousand hours learning to play classical music on the violin at a virtuoso level. But those same people might hum to themselves, drum on desktops, compose Filk songs with friends, play the harmonica, and sing in church on Sundays. As with music, stories are an essential part of the human experience. Everyone has the essential equipment installed to participate.

To help small groups make the most of their narrative capacity, we created a working prototype called Rakontu 1.0. Rakontu is free and open source software that small groups of people can use together to share and work with their stories. It's for people in neighborhoods, families, interest groups, support groups, work groups: any group of people with stories to share. Rakontu members build shared "story museums" that they can draw upon to achieve common goals. Rakontu is about small groups sharing stories for a reason. Rakontu communities are private spaces where people share personal experiences about something they all care about and in the process build something they can all use. Usually people in a Rakontu community will have something they want to do together, some common goal, and they will be interested in collecting and working with their stories as a means of getting there.

Disclaimer: *"Imitation is the sincerest form of flattery."* We have no affiliation with Peter G. Neumann, the moderator of RISKS Digest, or SRI (although Cynthia Kurtz has worked with other people at SRI on a joint project in the past related to the Genoa project). We have not consulted him about this proposal and it is not in any way endorsed by RISKS Digest or the ACM. However, based on his experience with RISKS Digest, he might be an excellent person to review this proposal or to provide other feedback on it, and you have our permission to distribute a copy to him. He might have valuable contributions and suggestions about such a project based on his years of experience thinking about computer risks and how people discuss them. We would be very open to using any such suggestions to improve the project.

What are you trying to do? Articulate your objectives technically and succinctly.

Sensemaking is the process by which people give meaning to experience. We want to help security professionals share real-life experiences and make sense of them collectively (primarily in small groups of generally less than 200 users) using a software platform called Rakontu designed specifically to support narrative exchange and sensemaking. Using Rakontu to accumulate cyber-security stories about both successes and failures and do sensemaking with them will lead to a better understanding of security-related best *and worst* practices, not in the abstract but linked to real-life stories in which the *consequences* of various practices are evident. It will improve the training of new cyber-security professionals by fostering a better

understanding of security strategies and values. It will improve the security of software systems through technical and social improvements. It will foster enhancements to intrinsic security (via hardening) and mutual security (via cooperation) across organizations.

Here are some titles from a recent issue of RISKS Digest:
http://catless.ncl.ac.uk/Risks/26.54.html

- Air France 447: Smart planes still vulnerable to human error
- Man unable to open car from the inside and dies of dehydration
- Bitcoin + Cloud Computing = Approx. USD$231K Up In Smoke
- Chinese newscast apparently reveals their cyber warfare games
- 4G and CDMA reportedly hacked at DEFCON
- Why Governments Are Terrified of Social Media

An excerpt from one of those stories:

> "Slip-Up in Chinese Military TV Show Reveals More Than Intended; Piece shows cyber warfare against US entities (*Epoch Times*)
>
> A standard, even boring, piece of Chinese military propaganda screened in mid-July included what must have been an unintended but nevertheless damaging revelation: shots from a computer screen showing a Chinese military university is engaged in cyberwarfare against entities in the United States. The documentary itself was otherwise meant as praise to the wisdom and judgment of Chinese military strategists, and a typical condemnation of the United States as an implacable aggressor in the cyber-realm. But the fleeting shots of an apparent China-based cyber-attack somehow made their way into the final cut. …"

There are thousands of such stories in RISKS Digest alone. People naturally find telling and reading such stories engaging and enlightening. Presumably they are learning a lot from this process already. Could we do even better than an ad-hoc approach to working with these stories? Could we improve the overall social process by more complete metadata and better tools for structured discussion and analysis?

Rakontu collects metadata in a conversational way. For example, say a cyber-security professional is contributing (or reading) the above story about the Chinese military propaganda slip-up. He or she might be asked questions such as:

- How long do you think you will remember this story?
- How high do you place the risk level apparent in this story?
- Can you think of a few words that summarize any opportunities you see in this story?
- From your point of view, would you say the story ended well or badly?
- To what degree do you think answers to the previous question by a variety of readers around the world would agree or disagree?
- Did you experience any of the events in this story? If so, would you like to recount your version of events?
- Does this story remind you of another one? Would you like to find or tell it now?

The reader might also have previously been asked demographic questions like these:

- How many years of experience do you have in your line of work?
- How large is your employer?
- What is your professional specialty?

When people tell each other stories in face-to-face conversation, both storytellers and audiences routinely annotate stories with evaluative questions and comments that serve to negotiate meaning in context. For example, a person might preface a story with a claim to utility like "Here's something I learned about trust," which is an answer to the implicit question of why the audience should give the teller the floor to tell a story. Similarly, an audience might place a story in context with an evaluative comment like "I never heard of such a thing!" which is the answer to the question "Do you find this story surprising?" Or they might ask a question like "Did that really happen?"

In other words, when people tell each other stories in real conversation, they don't put their stories into isolated boxes. They *aggregate* their stories into negotiated webs of meaning. Rakontu supports such aggregations by asking people contextually-appropriate questions about their stories and those of others, questions people might ask and answer if they were exchanging stories in person. Assigning tags to stories may be a burdensome task, but answering questions about stories is an activity socialized adults find not only familiar but obligatory and even sometimes enjoyable.

These sorts of conversations about stories can interact and accumulate to form webs of meaning that support decision making in contexts not always perceivable in advance. The benefits of such a system to a profession whose performance strongly depends on learning from collective experience could be significant. When such webs of meaning support structured real-time sensemaking, their benefits increase still further.

The charter of comp.risks from 1986 reads as follows:
http://catless.ncl.ac.uk/Risks/1.1.html#subj1

> "This is the first issue of a new on-line forum. Its intent is to address issues involving risks to the public in the use of computers. As such, it is necessarily concerned with whether/how critical requirements for human safety, reliability, fault tolerance, security, privacy, integrity, and guaranteed service (among others) can be met (in some cases all at the same time), and how the attempted fulfillment or ignorance of those requirements may imply risks to the public. We will presumably explore both deficiencies in existing systems and techniques for developing better computer systems -- as well as the implications of using computer systems in highly critical environments."

Fast-forward 25 years later, and here is a web resource related to comp.risks, written by the moderator, which says:
http://www.csl.sri.com/users/neumann/illustrative.html

> "NOTE: Many recent RISKS cases are not yet included. Maintaining this file has become increasingly labor intensive. However, the Election Problems section is

now up-to-date as of 14 March 2011. For other recent items, try the search engine at http://www.risks.org.
Copyright 2008, Peter G. Neumann, SRI International EL243, Menlo Park CA 94025-3493 …"

This illustrates the problem of having a third party add metadata (even just for top level categorization) to many stories after they have been told. If one person is trying to reorganize a quarter century of security risk stories, no matter how heroic and tireless the effort, we can assume that a lot of information is going to have to be ignored due to time constraints. The comment also implies copyright issues in trying to use such data comprehensively. Imprecise or missing copyright licensing statements potentially render 25 years of RISK contributions inaccessible as far as reliably transforming them into new resources in new contexts. An ideal system would make it easy to add metadata (including about free licensing or digital signatures) when an author creates a story or when others read the story or discuss it, so that at least going forward, systems could be created with clearer copyright status like Wikipedia.

Our hope is that once such an improved story-oriented system is available, security researchers will add stories in local repositories, with some stories being shared in public repositories. People could add stories from RISKS Digest and similar sources, although in practice copyright issues may limit re-use. The big win will be if security professionals add their own stories and related metadata. Essentially, this system has the potential to become a much-improved RISKS Digest-like system for the 21st century, built on narrative methods proven to work in other contexts.

While Rakontu 1.0 is a working prototype, it could be improved. The proposed work includes:

- porting existing code for story exchange (Rakontu) to a more flexible and capable and locally-runnable platform, moving from primarily HTML to a primarily Java Swing-based desktop application (with limited web support where the desktop application can publish content to the web in a variety of ways, or can use the web for simple data collection forms but with most interaction occurring through the desktop software);
- moving from Google App Engine to a different locally-runnable database with a triple-store transactional paradigm like refined in the Pointrel system and elsewhere;
- refining in-house code for mixed-methods analysis (NarraCat) into end-user code;
- developing new code for features of the system designed but not yet built (like real-time sensemaking workshops, story-gathering workshops, and graphical timelines);
- creating a sample database with stories about security issues drawn from RISK Digest and other public sources; and
- encouraging and supporting a cyber-risk-focused user community as they use the software and database.

The move back to the desktop from the web may seem like a step backwards, but based on research, there can actually be a benefit in small communities by increasing the barrier to entry somewhat by requiring a specific application to join a community. We have also found JavaScript toolkits to often have limited accessibility support, even the best in that area like Dojo, may not provide a consistent user experience across platforms and may require more related engineering time and produce worse results when doing sophisticated things like interactive data visualization. Digital signatures and end-to-end encryption are also harder to support on web applications than desktop ones with local keystores using locally auditable

code. Still, part of this move is also drive by our own limitations and expertise as well as the current state-of-the-art in JavaScript libraries; no doubt there may be fancy variants of the Rakontu system that are web-based in years to come. But, we need to have a well working and secure system first to build from, which Java provides for us. The success of "Minecraft", a Java desktop application for collaborative virtual worlds that has gained millions of users over the past two years with little advertising, shows what is possible with a cross-platform desktop application that can be run either stand-alone or using a service on the network. As an analogy, Rakontu 2.0 will be to stories what Minecraft is to one meter cubes, supporting groups of people who want to build structures from them that make sense to themselves.

Each of these elements poses risks in the form of technical as well as social challenges. These range from getting the requirements wrong to typographical mistakes in implementing well-defined algorithms to not being able to get the attention of security professionals (the last being, frankly, the biggest single risk to the project).

To minimize these risks, as two experienced computing professionals, we have already done a lot of background work. We have a limited first version of Rakontu working on Google App Engine. We have learned from that project and have created a Rakontu 2.0 design document based on user experiences from Rakontu 1.0 (though it was not used in the security domain). We also have done some preliminary prototyping of part of the real-time sensemaking aspect we want to add. We have thought about the design based on our previous experiences creating end-user software. We will learn from another project (RKB Explorer) that was able to use content from RISK Digest in an alternative way. We plan to use test-driven development for creating new code with minimal defects that meets requirements and yet can be refined further in an agile way. As an open source project, we will benefit from continual user feedback and engagement from the cyber-security community as the project progresses.

Keeping the software free and open source will lower the barrier for adoption of this software by security professionals, but even with that, we will need to publicize it through posts to mailing lists, Twitter, and other social media. After the concept is sufficiently proven, we plan to encourage computer security professionals to use and further refine the open source software in-house. Eventually public databases might be created and maintained, and a larger community might organize existing security stories (like those from RISK Digest) as well as add new ones.

We propose an ambitious eight month schedule for creating Rakontu 2.0 and getting the security community to start using it, split into four two-month phases. It is possible that some phases will take longer than two months (each comprising eight work weeks) if we encounter significant issues or get significant user feedback on essential features. Nonetheless, we are committed to seeing this project through once we start, so we will do whatever it takes to ensure these deliverables are completed even if they take longer than estimated.

As the software will be free and open source, we do not plan on gaining additional revenues from selling it. However, we do hope that (as with our previous work) we will be able to sell related services to some small percent of the software's user base. Those services might include training, improving the software in specific requested ways, or helping with data analysis, whether for cyber-security or other content domains. We hope to sell a printed user guide that is otherwise a free download. In these ways we expect to support and expand

Rakontu without additional grant funding from DARPA. However, because we are creating a modular system, there would always be opportunities for creating new modules in the future to support any specific government organization's policies or processes. Such modules would be outside the scope of this proposal.

While there remains much work to do, we feel confident that in a period of about eight months of 1.5 FTE work, we could port the relevant core of Rakontu software, enhance it in various ways outlined in our Rakontu 2.0 design documents, and create an initial workspace stocked with at least 200 stories pulled from RISKS Digest or similar sources, along with related added metadata, to prove the concept and act as conversation starters. Because the system itself will prove its value when it is actively used by security professionals within their own organizations and with their own stories, part of the project would be to inform security professionals about this new tool and encourage them to use it in-house.

The people who will care most about this project will be the security practitioners themselves, as well as, by extension, those who manage or depend on such professionals. The impact of this project will be a reduced time for new security professionals to come up to speed in the field, a lower loss of tacit knowledge with employee turnover, and potentially new deep insights into best practices in cyber-security. Potentially, this project could become the new platform on which something like RISKS Digest is eventually based, with all previous items submitted to RISK Digest maybe even someday eventually encoded into the system by the community (subject to any third-party licensing issues for that content versus “fair use”).

The participatory aspect will in general encourage security professionals to make their knowledge more widely available to other professionals. The general architecture will support several levels of scale, so that individuals, organizations, and the public can easily each have their own private or public databases, and stories can be moved between them as appropriate.

Explain what is new in the approach and why will it succeed.

In the distant past, the storage and organization of stories in human society was supported by people in caretaking roles such as griots, shanachies, bards and other master storytellers (who listened as much as they told). Such story caretakers could intelligently supply stories appropriate to a current situation and adapted to a listener's current practical understanding and cultural background. Today the role of story caretaker has diminished or disappeared in most organizations and communities, though each community of professionals has its elders who still perform these roles to a limited extent.

Our approach relies on the participation of security professionals, especially the elders of the profession, in semi-structured conversations that enable the accumulation of contextual and interpretive metadata about stories of experience. For example, a professional might describe a problem solved, then describe which tools, methods, strategies, or values were critical to success. These stories and metadata will create webs of meaning that help security professionals find the stories they need to solve new problems. As more and more stories and metadata form ever larger webs, patterns will appear which reveal critical insights, such as where particular tools are best used and where they provide little help, or even just on how to get organizations to cooperate better. These will inform the entire practice of cyber-security through the thoughtful accumulation of experience.

One problem common in the computer field (and so by extension the cyber-security field) is that with rapidly changing technology, there has been an attrition of older and more experienced technologists for a variety of reasons. One reason for this is that the importance of elders in the field as repositories of wise strategies and values is less apparent than the importance of someone with the latest technological skill. As a result, as shown on RISKS Digest, the same old security problems seem to be re-implemented over and over again in new languages and in new libraries. Any tool that helps make tacit knowledge held by elders more explicit and recognizable will help the community both in holding onto that knowledge and in seeing the value of people with broader perspective and experience. All successful communities need a mix of people whose strengths complement each other's weaknesses in different contexts. Sharing stories through a formal system is one more way for technology professionals with broader experience to contribute in a recognizable way. Otherwise such contributions often are made informally and go unrecognized except in their absence (when systems tend to fail because essential values and strategies have been forgotten, even to the point where the community has forgotten that it even used to know of a certain strategy). Stories can also help communities in identifying the people to “go to” for certain issues.

Some expert systems and knowledge management archives collect stories from real professionals but with context and interpretation stripped away. Most of these systems have third parties append metadata without consulting those who told them. For example, when a builder of an expert system adds metadata on, say, the utility of a story told by a human expert without asking the expert, an opportunity to place the story in context is lost. Some people, building beyond Roger Schank's work on newspaper reading software in the 1970s, have worked towards an AI-based solution to such issues, hoping that eventually some smart machine will do much of the thinking for us and bring together facts and context to reason about them. Or they hope tools similar to Google will help in finding the right story at the right time. Many builders of expert systems consider the annotation of experiences with relevant metadata to be clerical work and look for ways to minimize it. We have found instead that the opportunity to reflect on one's own experiences and those of others can be enabling, not burdensome -- if it is embedded in natural conversation in context. We have worked towards a Doug Engelbart-style “Augmentation” or “Collective IQ” approach to help a human community work together and co-evolve with its tools, datasets, and processes.

A variety of bug databases like Bugzilla collect reports about specific problems users have encountered with specific software and track their status towards resolution. Metadata collected in these systems is generally factual and of limited use in narrative sensemaking for decision support. For example, a user recounting an experience with a software bug might be be asked to rate the bug's severity. But in a narrative system they might also be asked questions of context like why they use the software in the first place, how long they have been using it, what they were trying to do when they encountered the bug, how they felt about the bug (angry, confused, amused), whether they intend to stop using the software because of it, and so on. In general metadata focused on narrative recountings of events is less about the facts recounted than on their context and relevance to storyteller and audience. Even such things as superstitions about software and beliefs about differing motivations among its builders and users ("they don't care what we need") can contribute to the discovery of insights that improve software in use. Still, such experience-based bug tracking systems are indirectly related to what we want to do, and in theory a future version of Rakontu might be somehow integrated more directly with a bug tracking system.

RKB Explorer, by Hugh Glaser and Ian C. Millard (to support "Resilience for Survivability in IST"), displays semantic content related to RISKS Digest as well as other databases. But RKB Explorer focuses on organizing stories using semantic tagging, and while that is useful, it does not explore narrative-specific annotations such as we describe above. Our project could be seen as an attempt to go further than RKB Explorer in a new and complementary direction.

Advanced tools for detecting patterns in collected stories do exist. These include Cognitive Edge's SenseMaker software, which we researched, built and maintained until a few years ago. That work was based in part on research we conducted in DARPA's Genoa I and Genoa II projects and Singapore's RAHS project. Software that supports qualitative and quantitative data analysis, such as NVivo and SAS, is also related and useful in this area. But these tools are geared toward supporting better decisions by top leadership based on information from a wide variety of sources (often anonymized to protect privacy). They are not designed to be used in a participatory way by those telling and reflecting on their own real-life experiences.

Narrative techniques are also used in ad-hoc ways in online discussions in newsgroups and chat sessions as well as in structured face-to-face workshops by experienced practitioners. Cynthia Kurtz's book *Working with Stories* helps people support such work in their organizations and communities. By embodying these techniques in software we hope to make these techniques easier to use and easier to learn by structured supported practice. SRI's SEAS and Angler projects (Cynthia participated in Angler design discussions) have overlap with Rakontu, including a participatory aspects, but they are not open source or primarily narrative-focused or, to our knowledge, used on cyber-security stories. Knowledge management tools are in use at places like IBM Global Services, and provide a way to create searchable knowledge bases, but they are generally not narrative-focused or participatory.

What is innovative about this project is that this will be the first time (that we know of) that an open source tool specifically designed to promote participatory real-life story sharing and narrative sensemaking will be applied to cyber-security stories. It is the gathering of stories of experience together with interpretations of those stories *by those most qualified to speak about them* that makes our approach uniquely valuable. The approach has been validated by years of research and practical experience in other domains. The project will succeed because it brings together an approach that works, a need that demands attention, and the combined energy of many professionals devoted to improving their collective ability to act in the best interests of those they protect.

Describe how the proposed effort increases or shrinks existing cyber attack surfaces and what new surface areas are created.

The size of cyber attack surfaces presumably depends on the training, experience, effort, and perspective of the people developing and supporting live systems and related standards. We believe that improving the ability of security professionals to share not just experiences but webs of meaning built from them has the potential to make computer systems more secure. Of course, the proposed system could itself be attacked in a variety of ways (detailed below). The tools could also be used to share stories about successful exploits against systems and promote an attacker culture. So, there are risks associated with this proposal. Yet improving the culture of security by sharing stories about security issues is something many already have

a long-standing faith in, such as is illustrated by RISKS Digest itself.

Is the solution buying tactical breathing space for existing problems or driving toward convergence by changing the playing field?

This proposal is intended to alter the playing field of how security professionals learn and share their craft with each other. This in turn will minimize the landscape of cyber-attack surfaces by increasing collective capability throughout the security community.

How does this solution incentivize the adversary? If the solution were deployed, how might the adversary use the solution to further their own goals? What are potential unintended consequences of the proposed solution?

If this solution was deployed, an adversary might try to defeat it by attempting to seed it with misinformation, shut it down, or use it as an attack vector against security professionals. However, it is likely that there would be both public and private users of the software, so private use of the software within organizations could minimize some of the risk for this. Such private datasets could be leaked though, potentially compromising the identity of security professionals; however, many are already identifiable in the community through their publications or public posts to mailing lists. As in the case of the Chinese training video mentioned on RISKS Digest, once a covert act is known to have been discovered, attackers can try to cover up the evidence somehow. However, as people already share stories about risks (such as on RISK Digest), it is not clear how substantially increased the risk would be.

What is the asymmetry for this solution? Given an attacker and defender of the proposal, explain which role is more advantageous (effort, cost, time, etc.) and why? Explain briefly for short, medium and long-term strategies for both the attacker and defender point of view.

The proposed system could be used both by attackers and defenders to share and organize experiences. Both attackers and defenders need to learn their craft somehow, so both are motivated to read documents and to share information. However, the advantage is to the defender overall in the long term. Defenses tend to become stronger the more they are discussed while attacks tend to become weaker the more they are discussed (since known attack modes can then be countered, as “forewarned is forearmed”).

Attackers can be considered to be in two broad groups. The first is small opportunistic groups (or individuals like “script kiddies”). The other is attackers in large well-organized military-level organizations or potentially seriously organized large-scale financial crime groups. This proposal provides more advantages against the small scale attackers. Since individual domestically-based attackers can be detected and prosecuted based on one mistake, and security professionals have many years to learn their trade, learning from their mistakes and sharing what they have learned with each other about dealing with both domestic and foreign-based attackers, it is our feeling that overall the system will help reduce asymmetries in attacker/defender balance related to small group attackers. Every time an defender shares information they cement their position in a community (as an expert), whereas every time an individual attacker (outside a large attack organization) shares information, they risk their position through exposing their identity (“pride goeth before a fall”). In general, it is not

currently as much of a risk for a defender's identity or location to be exposed (generally these are knowable already) as it is for an attacker's identity or location to be exposed (which are generally kept private in order to avoid legal or political consequences).

The risk/reward dynamics are somewhat different in an organization like a military division devoted to attacking systems, where within that zone shared information on attacks may pose no risk. But the recent alleged issue with a Chinese training video shows that attackers are much more at risk politically from their stories getting out than defenders, so even there, the asymmetries are in favor of this proposal in the long term. We can also hope that an experience-based system used by security professionals to help secure their systems will benefit in the long term from lots of good advice by security professionals about keeping the system itself secure, which may benefit against large scale attackers.

Since adversaries themselves have computer equipment to defend, and since experiential storytelling is such a natural part of the human experience, it is likely that even potential adversaries may be tempted to use the system to learn from it and contribute to it in the short-term. Also, some potential young ("script kiddy") adversaries may become exposed to a different and more mature culture of defensive values (including about honor and service) and may gain a different set of values that might encourage them in a new career direction.

If attackers do use the system to codify their knowledge, they are essentially describing their own best practices as well as potentially identifying themselves in some way. Both of these factors potentially weaken any asymmetric advantage if such datasets become available somehow. Defenders could study those datasets to find ways to respond to the outlined attacks or to understand the thinking process of attackers. Still, a very clever long-term-oriented adversary might create a false and misleading database to waste the defending community's time in analysis as a sort of meta-level denial-of-service attack.

Albert Einstein said: "The release of atom power has changed everything except our way of thinking...the solution to this problem lies in the heart of mankind. If only I had known, I should have become a watchmaker." Now that digital watches and phones running watch-like applications have more computer power than was used to design early atomic weapons, even watchmakers face tough practical and contextual issues about what they do from a security perspective. In the long term, there is a fundamental irony to using the tools of abundance (cheap computing) to attack other people in order to gain access to resources. A more constructive alternative is to use computing power to produce resources directly (like through robotics, better design, or advanced materials simulations). In the very long term, people sharing stories about defense and abundance may lead to a new vision of intrinsic security and mutual security that reduces the overall level of strife on the internet.

If you were to have to defeat your own effort, how would you go about it? (Note: it is perfectly acceptable to identify deficiencies in the effort. It is not acceptable to believe that there are none.)

These are some ways the system could be misused or subverted:

- An attacker could launch denial-of-service attacks against story sharing sites to prevent their use (including by spamming them to waste security professionals' time).
- An attacker could use the system as a vector to mislead (or attack) specific security

researchers or security installations.
- An attacker could use the system as a vector of hateful propaganda.
- An attacker could prevent the system from being developed in the first place or contribute in negative ways to ensure it was developed badly.
- An attacker could subvert the underlying infrastructure below the system.
- An attacker could use information learned from the system to discern potential flaws in security thinking or potential unpatched problems in some older systems.
- An attacker could monitor public services to see if people were aware of the attacker's activities and if so change their attack approach.
- An attacker could use the knowledge in the system to harden the attacker's own systems against intrusion.

Likely there are other possibilities as well.

Who could benefit from this technology? Be specific and cite an actual use case.

A specific use case might be within a group tasked with ensuring secure US government communications such as within the NSA. They could use such a system to improve training quality and reduce the impact of employee turnover. Experienced professionals could share experiences about how they evaluated and secured government systems running, say, Debian GNU/Linux. Because these experiences would be surrounded with conversational annotations, they would retain essential elements of meaning in context that would accumulate into webs of meaning. From these webs new employees would learn not just best practices but the values, assumptions, and priorities behind them. Experienced security professionals would continue to learn and improve their skills as they participated in collective narrative sensemaking. Even state-of-the-art practices might be improved further by the discovery of insights based on collective experience.

This system could also be used within a large organization such as General Electric. It could be used by small organizations as well, either creating databases in-house or just by accessing larger public repositories (in the same way some people learn about basic themes in computer security by reading public resources like RISKS Digest or Slashdot).

Capability/Technology Information

We have previously worked on related narrative technology for the US DOD's Genoa I and II projects and Singapore's RAHS project. Aspects of this narrative technology have been proposed and developed as part of those, but with a non-participatory emphasis. It is the participatory aspect of this project that goes beyond those projects, although they also had other features that are not within the scope of this project (most of which we did not work on).

On our own initiative we built a Rakontu 1.0 working prototype as free and open source software intended to be used in a variety of domains by a variety of different groups as a general purpose tool. In a sense this is "cost sharing" in preliminary development. Rakontu 1.0 has some aspects of what we would like to have in Rakontu 2.0, but we have learned from that system and want to make substantial improvements to it, including making a desktop version using Java. These changes include porting it from Google App Engine to another database (SQL-based and/or CouchDB) to make it more easily deployable locally; removing

some things that did not work out well (as documented in Cynthia Kurtz's critique of Rakontu 1.0 in a blog post called “Steal these ideas”); adding a real-time aspect for collaborative real-time workshops, and integrating narrative catalysis tools from NarraCat (now in Python).

We have also done some initial test coding on part of the new GUI in Java to look at issues related to Java-CouchDB interaction and Java-SQL interaction, and in part to evaluate accessibility and security issues. Based on that we have decided substantial parts of the system will be in desktop Java for accessibility reasons, for security reasons, and for reducing initial development and support time costs. We created the beginnings of a proof of concept for real-time workshops working on CouchDB using JavaScript and the Dojo library. But that demo still has cross-platform updating issues, accessibility issues, and security issues that need to be resolved, and it is only a partial test of the technology for real-time workshops. So there remains much work to be done on that part, including porting it to desktop Java. A live demo of that technology preview running on CouchDB using JavaScript can be viewed here: http://pdfernhout.couchone.com/twirlip2/_design/twirlip/conceptMap.html?diagram=SecurityStories

Interactions with the Ad-Hoc Cyber Research Community

Software built as part of Genoa II and RAHS is currently being used by national governments in relation to national security threats and opportunities. A proprietary spinoff of that called SenseMaker is in current use for non-participatory research related to understanding trends and perceptions in organizations and communities. We researched and built the early versions of SenseMaker when Cynthia Kurtz was Cognitive Edge's Director of Research, but stopped working with Cognitive Edge about two years ago and know little about the current version of SenseMaker or its uses. Cynthia worked on two projects that used SenseMaker in defense communities to better understand issues related to retention of security professionals considering leaving national service. Those applications of SenseMaker (outside of Singapore's RAHS) do not relate directly to cyber-security efforts. Our Rakontu 1.0 software (released about two years ago) was used to support a “Mistake Bank” of experiences in a small group of interested participants. That use was general and not specific to cyber-security or even security in general. We do anticipate substantial interactions with the ad-hoc cyber research community as a result of this proposal going forward.

Metrics

There are two broad ways of evaluating the success of this work. The first is simply in terms of whether the project succeeds in meeting its stated goals. That is, whether it substantially implements the Rakontu 2.0 specification defined by the figures included above and documents referenced in the statement of work; whether it substantially integrates the NarraCat tools into the overall system; whether it creates a sample dataset of 200 cyber-security risk-related stories; and whether it makes appropriate announcements to cyber-security professionals and constructively engages with them as evidenced on a development mailing list or similar venues like issue tracking systems.

The more important metric of success is whether cyber-security professionals begin to *use* such a system. This can be measured in part by monitoring such things as traffic on related user mailing lists and bug tracking systems, tweets mentioning #rakontu, downloads or forks

of the software, contributions of cyber-security stories made to any publicly maintained Rakontu servers, and contributed improvements to the software. Because this project aims to reach the point where professionals perceive it as a ready-to-use system for their own use, those quantitative metrics will not be available until after the system has been developed to a certain extent. Since we plan to announce the system after the first two-month task is complete, we hope there will be some involvement by cyber-security researchers to show by the end of the project (especially for the fourth task of user testing as we recruit beta testers).

Long term success could be measured quantitatively by the extent to which, as a result of use of the Rakontu software, organizations reduce their training costs and reduce the number of intrusions into their systems through enhanced capability. However, it may be difficult to attribute such improvements directly to the use of the Rakontu system. One qualitative long-term metric might be the extent to which the security community expands the Rakontu system and integrates it into its larger work ensuring cyber-security, including aspects such as semantic networks that complement the narrative approach.

Statement of Work (SOW)

The details of what we would like to do for the next version of Rakontu (2.0) are outlined at the following links and are hereby included by reference (although the figures at the beginning of this document are taken from these sources and define graphically the essential functionality we plan to implement):

http://www.rakontu.org/pitch.html
http://www.rakontu.org/Rakontu2_Architecture.pdf
http://www.rakontu.org/Rakontu2_VisualDesign.pdf
http://www.storycoloredglasses.com/p/narracat-tools-for-narrative-catalysis.html

The system will be primarily in desktop Java (but may also include a web component with JavaScript for ease of remote access) and may require some related freely-available dependencies like freely licensed open source libraries or toolkits to be installed. The project will be developed using free and open source software approach under the Affero GPL license or compatible licenses, with the code itself maintained on GitHub or similar places like SourceForge. The exact licensing of the sample content will depend on its sources, but our intent is to have at least 200 security related stories drawn from RISKS Digest and other public sources. However, those stories are mainly just to have something to work from, to test on, and to use as conversation starters. Ultimately the success of the project has to be measured by the number of security researchers and practitioners who decide to use it either in some public form or in-house and contribute their own stories and metadata. We anticipate potentially millions of cyber-security related stories stored in Rakontu systems taken as a whole (although that may be difficult to measure if many databases are private).

We propose four major milestones related to payments, each mostly depending on the previous stage being substantially completed. All deliverables can be verifiable when completion is announced by examining our progress with commits to a free and open source software hosting system such as SourceForge or GitHub as well as related public posts about the project. Software can be downloaded and tested for substantial compliance with the design documents. All documents will be produced in open formats such as Open Office document formats, HTML, or PDF files and made available on a project web site.

Each stage is anticipated to take two months (including time for talking with DARPA project officers and evaluators). Each task will be worked on by both Cynthia Kurtz and Paul Fernhout, with Cynthia working 0.5 FTE and Paul working 1.0 FTE. In common for all stages, we will be using open source development environments like Eclipse with source code control like with Git. For accessibility reasons, the GUI will be primarily written in Java (or other JVM languages) for use on the desktop or via Java webstart. There may also be limited functionality via web code that duplicates a part of what the desktop can do for support of publishing or story collection. We will use test-driven development to help assure quality of the code and an ability to subsequently use agile methods for any needed refactoring later.

Task 1: Core story exchange. The first two-month period involves getting the core of the existing Rakontu 1.0 system onto Java (and potentially other JVM languages) with a standard SQL backend and/or possibly CouchDB, which in general will entail rewriting much of the code of the entire system. This task will deliver working Java-based desktop code for supporting adding stories and topics, annotating stories based on answers to questions, and supplying basic personal profile information on the participant. It is intended that the system be useable by others at this completion of this task, even if it will be lacking in advanced features, under the mantra "release early, release often".

Task 2: Timeline and real-time sensemaking. The second two-month period will entail adding extra features to Rakontu (as outlined in the above referenced documents) to support real-time workshops and also to have an interactive timeline of what is going on in the system. During this period, we will also announce what we are working on and ask security professionals to begin to take a look at the system, to begin the process of getting user feedback as we continue to improve the system. This task will deliver working code for real-time sensemaking workshops and a an interactive timeline.

Task 3: Visualization and analysis. The third two-month period period will entail integrating data visualization and statistical analysis tools based on NarraCat. NarraCat is a set of Python scripts developed by Cynthia Kurtz for in-house data analysis and visualization in support of narrative projects. It was released in July 2011 under an open-source license. NarraCat is not yet designed for the end user, so adapting it will take additional development work to produce new code in Java, JavaScript, and/or Jython for the JVM. We will also continue enhancements to the system based on user feedback. This task will deliver code to support statistical analysis and visualization integrated with the system.

Task 4: Testing and refinement. The fourth two-month period will entail additional user testing with a small number of invited security professionals. We will ensure that there are at least 200 existing stories about security seeded into the database from RISKS Digest and other sources as conversation starters and for use in testing and evaluation. We will prepare a ten to twenty page final report, web presentation overview (probably as a ten minute video), and short introductory user guide on the project. This task will deliver a sample database of 200 stories, user guide, web presentation, and final report. The system will have substantially sufficient stability, security, and documentation so it can be easily set up in-house for private cyber-security story databases or with public databases for broader public participation in collecting public security-related stories in a participatory way (perhaps even eventually as a candidate for the ACM to consider to supplement the venerable RISKS Digest infrastructure).

Schedule/Milestones/Deliverables

Assuming an October 1st start date:

October and November 2011: **Task #1**
Deliverables: **Working code** publicly available and useable by others for collecting and discussing stories using a SQL backend with a front-end as some mix of desktop Java and other JVM language code.
December 2011 and January 2012: **Task #2**
Deliverables: **Working code** for real-time sensemaking workshops and an interactive timeline. **Announcement** of this project to security professionals.
February and March 2012: **Task #3**
Deliverables: **Working code** to support statistical analysis and data visualization integrated in with the system.
April and May 2012: **Task #4**
Deliverables: **Sample database** of 200 stories, **user guide**, **web presentation**, and a **final report** about a substantially useable system that government agencies and other organizations can begin to use in-house with their own private cyber-security stories.

Cost

The labor cost of this effort is computed at 0.5 FTE for Cynthia Kurtz and 1.0 FTE for Paul Fernhout at rates of US$105 and US$75 per hour respectively. These are similar to previous rates we have received for government and corporate contract work. Additionally, we will factor in 10% for general overhead and 10% for profit. The cost for each deliverable milestone based on eight weeks of 1.5 FTE work during two calendar months will thus be US$48,960. The total **cost for labor** plus overhead and profit will be four times that, or **US$195,840**.
Reference: ((105 * 0.5 * 8 * 5 * 8) + (75 * 1.0 * 8 * 5 * 8)) * 1.2 = 48,960 per task.
This is an ambitious undertaking, and it is possible it will take us longer to reach each milestone, but the payment for each milestone can be considered a fixed cost. So, it is possible, even likely, that the effective labor rate will be lower.
We are based in New York State and so locations in the North East are drivable to us. We are anticipating that up to three road trips to make related presentations (such as to Washington, DC and to an East Coast military academy like West Point). Our past government reimbursed costs for a two-night stay to Washington DC has been approximately US$1500. We estimate a visit to some place closer like West Point costing about US$1000. So, we are including **travel expenses** of **US$4000**.
The **total cost** for this project for labor, overhead, profit, and travel is thus: US$**199,840**

Appendix A - Proposer Team Members and Other Required Information

Proposer team members at Kurtz-Fernhout Software:

Cynthia F. Kurtz, CEO, Principal Investigator
Paul D. Fernhout, CTO, Software Developer

Our CEO, Cynthia F. Kurtz, is one of the world leaders in using narrative methods involving real-life experiential stories to discover new insights about complex issues. See for example the widely cited: Kurtz, C. & Snowden, D. 2003. "The New Dynamics of Strategy: sense making in a complex and complicated world" in *IBM Systems Journal* 42(3): 462-483. She has a dozen years of experience helping large organizations and government agencies work with stories for decision support and organizational learning. She has previously worked on DARPA's Genoa I and II projects and Singapore's RAHS project, which made use of narrative methods to help analysts better discover weak signals and consider diverse perspectives in complex situations in the social, environmental, economic and security spheres.

Her husband, our CTO, Paul Fernhout, has been working on semantic software for over twenty-five years, implementing his semantic first triple-store in the early 1980s, and continuing those exploration with the free and open source Pointrel system.

The team has produced several large software applications together over the past two decades.

Subproposer team members: N/A

Non-US organizations involved on the development team: NONE(*)
(*) Because of the nature of open source software, it is possible that non-US citizens may make contributions to the code or content as it develops. For example, in the past we have had contributions to French localization of Rakotu 1.0 from a resident of France; such specific contributions can be determined by consulting the source code repository. All such contributions will be reviewed by the proposal team before integration in the mainline of development through a services such at GitHub.

Proposer or subproposer organizations belonging to a Government entity or FFRDC: NONE

Organizational Conflict of Interest Affirmations and Disclosure: NONE

Intellectual Property: All software developed by the proposers under this contract will be released to the public under the Affero GPL license or similar free and open source software licenses, and the government will retain Government Purpose Rights to that new software. However, the resulting product may include previously written software by us such as from Rakontu 1.0 licensed under the GPL or similar copyleft licenses that may restrict the licensing terms under which the government can use the software in terms of "copyleft" requirements for distributed changes (but which will not restrict free in-house use by the government for its own purposes). While substantial contributions by others are not expected in the first few months of work (small bug-fixes might be more typical), any contribution to the codebase under a copy-left license by others during the course of the project would further restrict the licensing options for the whole without re-engineering to remove that contribution. All related

published sample database content will be generally made available under as permissive a license as the original sources allow, such as a Creative Commons licenses that at a minimum would allow free local use, but to the extent there are restrictions on the use of original source materials such as in altering specific stories to create derived works, those restrictions would apply to government use of the sample datasets. This project intends to build on top of well-known widely-used free and open source software platforms such as SQLite, JavaScript, Java, Python, and related libraries, and those in turn have their own (generally well-known and manageable) issues related to copyrights and patents.

Human Use: NONE(*)
(*) Note that there is no use of human subjects directly in this work, but because it is intended to be used by security professionals to gather their stories, there are implications related to privacy. Those will be managed by using only publicly available stories in the sample dataset of 200 items (like those already posted to comp.risks or similar venues). However, there may be additional public datasets the project creates as part of its activities, but any new items contributed by people will only be accepted by people who are aware that their contribution is going to be integrated into such a public dataset and are in agreement with a specific license such as a Creative Commons license similar to Wikipedia's.

Publication Approval: Because this project will be developed as a free and open source software project in a public code repository like GitHub, to proceed it requires implicit publication approval by the US government under this contract for public release in advance. It is not intended that any secret or sensitive information be handled by the project staff during the course of this project or that any such information would be included in any sample public datasets the project creates for testing or “conversation starter” distribution under this contract. Any such information would be removed from a public dataset if we became aware of in relation to work performed under this contract (subject potentially to limitations of distributed source code control systems for expunging data even with best business efforts). Further, it will be specifically asked of any contributors that they do not contribute secret or sensitive information to the project's public datasets but instead create their own local private archives for such things or in other ways follow appropriate disclosure channels.

Appendix B - Additional Optional Information

Systems and research mentioned

RISKS Digest:
http://en.wikipedia.org/wiki/RISKS_Digest

Peter G. Neumann
http://www.csl.sri.com/users/neumann/neumann.html

RKB Explorer
http://www.rkbexplorer.com/about/
http://risks.rkbexplorer.com/
http://ckan.net/package/rkb-explorer-risks

Roger Schank's work on Newspaper Reading programs
http://www.rogerschank.com/biography.html

Doug Engelbart's vision on boosting "Collective IQ" via "Networked Improvement Communities":
http://www.dougengelbart.org/about/vision-highlights.html

DARPA's Genoa II program:
http://web.archive.org/web/20030621111224/http://www.darpa.mil/iao/GenoaII.htm

Singapore's Risk Assessment and Horizon Scanning program:
http://www.rahs.org.sg/

Cognitive Edge's SenseMaker
http://www.sensemaker-suite.com/

SRI's SEAS and Angler (Cynthia Kurtz was involved with some Angler design discussions):
http://www.ai.sri.com/~seas/
http://www.technologyreview.com/computing/13115/
http://www.ai.sri.com/~angler/

"Software is Hard" and Rosenberg's Law
http://gamearchitect.net/Articles/SoftwareIsHard.html
Rosenberg's Law: "Software is easy to make, except when you want it to do something new. The corollary is, The only software that's worth making is software that does something new."

Our work

Publications by Cynthia Kurtz:
http://www.workingwithstories.org
http://www.cfkurtz.com/publications.html
http://www.storycoloredglasses.com

Rakontu Version 1.0:
http://www.rakontu.org/
http://www.rakontu.org/tour.html
http://www.rakontu.org/screenshots_member.html
http://code.google.com/p/rakontu/

Rakontu 1.0 critique, "Steal these ideas"
http://www.storycoloredglasses.com/2010/08/steal-these-ideas.html

Rakontu Version 2.0 plans:
http://www.rakontu.org/Rakontu2_Architecture.pdf
http://www.rakontu.org/Rakontu2_VisualDesign.pdf
http://www.rakontu.org/pitch.html

NarraCat:
http://www.storycoloredglasses.com/p/narracat-tools-for-narrative-catalysis.html

Paul Fernhout's Pointrel Social Semantic Desktop experiments:
http://sourceforge.net/projects/pointrel/

Paul Fernhout on "Recognizing Irony is Key to Transcending Militarism"
http://www.pdfernhout.net/recognizing-irony-is-a-key-to-transcending-militarism.html

Paul Fernhout on "Some thoughts on strengthening security post-9/11"
http://slashdot.org/comments.pl?sid=2422714&cid=37371702

If this basic system becomes widely used by professionals in the security field, we hope it will help increase the amount of "cognitive diversity" available to security researchers and which can sometimes otherwise be limited by the security clearance process, as explained here:
http://www.phibetaiota.net/2011/09/paul-fernhout-how-security-clearance-process-harms-national-security-by-eradicating-cognitive-diversity/

We also have hopes to later expand the Rakontu platform eventually, moving towards a general "social semantic desktop" approach (including perhaps the sorts of additional modules mentioned as possible in the figure on modularity), but that is beyond the scope of this specific proposal. Here is information about those broader ideas and applications:
"Social Semantic Desktop for Sensemaking on Threats and Opportunities"
http://slashdot.org/comments.pl?sid=2368162&cid=37016386
http://www.phibetaiota.net/2011/09/paul-fernhout-open-letter-to-the-intelligence-advanced-programs-research-agency-iarpa/