Enterprise 2.0 and the Context of Work

Carmen Medina
Director, Center for the Study of Intelligence, CIA
Traction Software User Group
Keynote TUG 2009 Providence RI

Notes are keyed to slide numbers

1 ----------

Show the first 4-5 minutes of the presentation

Make point that even individual who are trying to think about a technological future get things fundamentally wrong. Point to the large dials and levers. They totally did not anticipate miniaturization and digital systems.

Mention the "coal-powered" spacecrafts. How difficult it is to tell which one is really burning

2 ----------

So it with some trepidation that I set out to talk about how Web 2.0 (and 3.0 and 4.0 and 5.0)--the capabilities these terms represent--will alter work--no, that's not correct--will lead to a fundamental reconceptualization of work itself. I'm certainly not unique in making this point--many others such as John Seely Brown and Marina Gorbis of the Institute for the Future are also thinking about this. As an aside, I might note that just since I started working on this talk earlier this summer, the term Web 2.0 has begun to sound silly to me--like the word banana. The term already sounds antiquated.

But no matter how many minds cogitate on how things can be different, we, just like the creators of Flash Gordon, will get many, essential things wrong. We'll get fundamental structure and architecture aspects just plain wrong. (which is actually a point about collaboration, but I'll return to that later.) (DELAY SLIDE)

3 ----------

But before I go any further on how Web N.0 will change the context of work, I want to tell you a story, a true story. A story some of you may know a little or even a moderate amount about. But hopefully no one in the audience knows a lot about it because then I'll be BUSTED--you'll catch a lot of my errors, cuz you see I am in no way an expert on electricity. (REVEAL SLIDE)

If you busted me, I would look foolish, a condition of which I generally try to steer clear--which reminds me that recently I've been trying to live every day so that I make no mental errors. Like in baseball, if you r supposed to throw home but instead, for an inexplicable reason, throw to first instead, you've made a mental error.

Why do we make mental errors? I think that's an interesting question, you know. Doesn't it feel sometimes like you make a mental error because your RAM or your internal cache is full--or maybe I'm just obsessing about this because I'm getting older.

But this mental error issue--this also reminds me about collaboration, but I'll get to that later.

The story I'm going to tell you is the story of the 2003 electricity blackout on the East Coast. I remember at least the start of the blackout quite clearly. I was flying to Ireland that evening--14 August 2003--and arranged to fly out of Philadelphia, rather than the more usual New York or Boston--and Philadelphia's airport was the northernmost airport to escape the blackout.

What I would like to do is describe the events, processes, and yes mental errors that eventually resulted in the massive blackout. Almost all of this is drawn from the official report of the US-Canadian commission charged with ascertaining the cause or causes.

4) ----------

What I'm going to do is detail the chronology of events that led to the blackout. I'll start with the first event the commission identifies in the causal chain and keep going until, as the commission judges, the cascade of events that led to the blackout became inevitable.

But first we need to level set with a few facts about electricity.

As I think we all know, it is very difficult, in fact impossible, to store large amounts of electricity. As a result, electric utilities must constantly monitor the health of their entire system, the power stations, the substations, the transmission lines. The physics of electricity are fascinating--there are different types of power and they have to be kept in balance, the flow can change direction--I learned enough to realize I couldn't explain it other than to say that electrical systems can only be understood in terms of real-time flow. That's another point I want to build on later when I talk about collaboration--so that's three points about collaboration that I've parked.

Electric utilities estimate how much power they need to meet typical demand on a given day, factor in a whole bunch of extra power to handle contingencies, and then come up with a threshold number above which they need to stay. They typically share power among different electric utility systems

5 ----------

August 14, 2003 was a pretty normal day in the Mideast--in the Ohio/Pennsylvania region. It was warm, but not that warm--in most areas it never broke 90 degrees. It was humid, but not that humid. The area was not in heat wave conditions. The usual amount of power generation was down for some kind of repair or another.

At 1215 the state estimator operated by the Ohio-based Midwest Independent Transmission System Operator fails. (From now we'll call them MISO.) A state estimator mathematically processes raw telemetry-type data coming from the different parts of the electrical system. It then calibrates this data against a model they have of a well functioning electrical system to determine if anything is out of whack.

Essentially what happened is that at 1215 the State Estimator produced a solution--a current state of the system--outside the bounds of system dynamics, i.e. it was a nutty solution that indicated not a real problem, but some kind of technology fritz. As you can imagine, the state estimator is only as good as its data, so the operators troubleshot the state estimator to see where the problem was. They found the problem, they fixed it and got a good reading. But to troubleshoot the state estimator, you have to turn the system off, AND THE OPERATOR FORGOT TO TURN THE SYSTEM BACK ON. CLICK

Because this software was not working, MISO could not perform COMPLETE contingency analysis of generation losses within its area of responsibility for the next 2+ hours.

A little over an hour later at 1331, the Eastlake 5 Ohio generating plant tripped and shut itself down. CLICK The reasons why it shut down have to do with some arcane fact of electricity generation that I couldn't understand well enough to explain concisely. The key point about the unexpected loss of this plant is that First Energy's electrical system was now much closer to its threshold line than it would like to have been. First Energy itself did not conduct a contingency analysis after the loss of the Eastlake plant, and you'll recall that the guys at MISO had not turned the State Estimator back on so they were unable to do the analysis that would have discovered that, without the Eastlake plant, the First Energy system would be in trouble if transmission lines started to go down. Hmm... I wonder what will happen next.

Two minutes after two, the first overhead transmission line, the Stuart-Atlanta line, fails. CLICK This line actually is not important in terms of power transmission, It only really matters because MISO never knew it was down and so at critical times later in the afternoon it kept trying to understand system status assuming this line was up, which led to all sorts of miscalculation.

Twelve minutes later, the alarm and logging software at First Energy's control room failed. CLICK This was the software that when it received certain kinds of data, would issue visual and audible alarms so that individuals in the control room would know something was amiss. The data kept coming in but without the alarms the engineers needed to manually scan their systems, some of which were several screens or steps down, to know there was a problem. They did not realize they needed to do that.

At 1420, several remote consoles at substations fail. It's not clear whether these failures were a result of the large failure in the First Energy control room. As each console failed, First Energy's information technology staff was paged. CLICK

Seven minutes later, the second transmission line--the Star/South Canton line, fails. CLICK Lines usually fail because of tree contact that occurs after a line sags. The lines sag as they become hotter, because of the ambient temperature or because they overheat from carrying excessive power. This is why it is important that trees around power lines be trimmed regularly. This first line did not fail because it was carrying excessive power--it was at less than half its recommended load. It failed because the trees had not been properly trimmed.

Five minutes later, a neighboring power plant calls the FE control room to report that the STAR/SOUTH CANTON line must have tripped because they are getting whacky readings. The FE engineers say: "Nah, there's no problem. There's no....ALARM."

Now we get to some of my more favorite parts. At 1441 the primary server at First Energy fails. This is the server that hosts the Energy Management System, which you will recall isn't working correctly because the alarm software has failed. CLICK The commission determined that the primary server probably failed as a result of the queuing of non-alarmed data and the failure of the remote consoles. Now given that the EMS system is so important, there is a backup server always in hot standby. So the alarm application is moved intact--failure and all--to the backup server. 13 minutes later the backup server--afflicted by the same unresolved data queues--fails. CLICK During this time, by the way, the screen refresh rate for the engineers has gone from one to 59 seconds. When the backup server fails, the IT staff is, of course, paged again.

At 1505, the third line, the Harding/Chamberlain line fails. CLICK Needless to say, the First Energy engineers do not notice because they are in Dante's ninth level of software hell. Three minutes later, the IT staff completes a warm reboot of the primary server. The IT staff thinks everything is normal but in fact the alarm application is still frozen. And the IT staff, and truly it pains me to say I myself have experienced this kind of miscommunication once or twice in the past, well anyway the IT staff does not confirm with the control room operators that the alarm system is in fact operating properly before they declare victory.

And finally, after the blackout, we learn the EMS system, which had been purchased in 1995, was not running the latest version of the software. And why was that? The software had not been updated because First Energy had decided to replace the EMS system with one from another vendor.

The area served by First Energy receives its first indication the system is becoming unstable when at 1517, voltage dips temporarily on the Ohio portion of the grid. Afterwards, major users of electricity on the grid, beginning to notice the voltage fluctuation, start taking temporizing steps. Fifteen minutes later the Fourth major transmission line, the Hanna/Juniper line, fails as it sags into a tree because of the extra power it is carrying. CLICK It should be noted that even this line was still only at 88% of its peak carrying capacity.

So again, bad tree trimming is partly to blame. MISO, because if you recall all throughout this period its state estimator is not working because it was left turned off, does not discover the Hanna/Juniper line tripped until after the blackout occurred. The inadequate awareness of lines that were failing led models and human operators to not understand or project current and future levels of distress. It was in fact when the Star/South Canton and Harding/Chamberlain lines failed, that the demands on the rest of the system began to exceed contingency levels.

From now on we truly are moments from disaster but not yet beyond the Carole King line. (it's too late, baby, now it's too late) Less than ten minutes after the Hanna/Juniper line trips, a line co-owned by First Energy and American Electric Power, the Star/South Canton line, fails, leading to the tripping of a circuit breaker connecting the two grids and another cascade of line failures.

CLICK Two more major lines are slated to fail before the blackout, according to the US-Canadian commission, becomes inevitable. First the Tidd/Canton Central line trips 5 minutes later as energy keeps trying to flow north to meet demand in Cleveland. CLICK Twenty minutes after that, the Sammis/Star line fails due to undervoltage and overcurrent, whatever that means.

CLICK The Commission believes that during those 20 minutes the First Energy operators needed to realize that cutting off the Cleveland /Akron area from the electrical grid would have prevented the catastrophic blackout that ensued. It is unclear whether FE operators ever fully realized the trouble they were in, although the control room supervisor did inform his manager at 1545 that it looked like they were losing the system. But did they realize they were about to lose most of the East Coast as well? I don't think so.

6 ----------

So I'm hoping that by now most of you in the audience have some sense of why I've told this long and complicated story as part of a talk on Web N.0 and the context of work. Dare I say it, the lightbulbs are going on. But to drive the point home, let's turn to the conclusions of the Commission report.

Four groups of Causes were identified for the blackout.

First, at a very strategic level, the major electric utilities did not really understand how vulnerable the system had become because they did not really understand voltage stability conditions and needs.

The second fundamental group of causes, one that concerns us and to which I will return in a moment, was inadequate situational awareness and understanding of the problem.

The third was the failure of all involved to adequately monitor tree growth and follow trimming policies.

And the fourth, which again is relevant to our discussions, was the failure of other ancillary organizations to provide real time diagnostic support.

If I may take a moment to point this out, these four causes are a useful template for diagnosing big goofups in any situation: first you probably have had a strategic failure to clearly understand the broad parameters of the situation and particularly how it has changed/evolved over a period of time, second, you have some fundamental knowledge awareness and processing problems, third, you have some tactical issue--in our case tree trimming--that was messed up due to lack of attention, maybe cost saving measures whatever, and finally there is a lack of effective collaboration from others that could have helped. Take one from each column, stir, and you have a disaster.

When you read through the discussion of information system and knowledge processing issues, you realize there is also a problem with simple collaboration and teamwork--except of course that teamwork is never simple. As the

commission notes, transmission system operators must have visibility not only over their own transmission facilities but they must also recognize the impact on their own systems of events and facilities in neighboring states.

The commission report also notes that the most critical factor delaying adequate assessment of what was happening to the grid was lack of information sharing.

The commission notes that rarely were any of the critical clues shared with fellow operators. They cite four factors for this:

1. physical separation of operators

2. the lack of a shared electronic log

3. lack of systemic briefing during shift changes and

4. infrequent training of operators in emergency scenarios.

Finally, the commission made several recommendations pertaining to information technology and knowledge management including the improvement of IT forensic and diagnostic capabilities, the development of IT management procedures, security governance and strategies, and the development of capabilities to monitor network and system health and better manage incidents.

7 ----------

In fact, electric utilities are now pursuing new IT, knowledge awareness, knowledge management, and collaboration solutions that seek to imbed their work within the framework of networked applications applying principles of social networking. I saw some of these initiatives when I visited Pacific Northwest National Laboratories earlier this year.

Other than the serendipity that brought me to PNNL in August, there is another reason why I chose to illustrate my talk with the story of the 2003 Northeast blackout. Electric utilities fall in the category of High Risk, High Reliability organizations--and the CIA, by the way, is one of these as well. HR-squared organizations do activities that are inherently risky but have to do them with a high level of reliability/performance. High-performing High Risk, High Reliability Organizations have several noteworthy characteristics often cited in the literature:

1. They are determined to avoid failure, so they think about failure a lot. They study it carefully when it happens.

2. They don't simplify easily or prematurely. They're OK with messy. They worry more about jumping to conclusions.

3. They are designed for resilience--multiple eyes, 360 degrees awareness.

4. They are operations and process-oriented. They care about the means--the ends can never justify the means because you are always striving to identify predictable, reliable causal relationships.

5. They defer to knowledge and expertise, i.e. not just to organizational relationships and hierarchy.

One of the most successful high risk--high reliability industries is aviation. The safety record of commercial aviation is testament to modern engineering and organizational methods. But aviation benefits from a very interesting attribute--it routinely engages in the riskiest parts of its enterprise. Every day there are tens of thousands of takeoffs and landings--the two most stressful aspects of flight. Takeoff places the greatest strain on the equipment; landing places the greatest strain on the humans. In that way aviation differs from many other HR-squared organizations, such as electric utilities. I will return to this point in just a little bit.

The conditions of work that high reliability/high risk organizations must achieve are exactly the conditions achieved by redesigning work to be done completely within the web environment. (If you're going to tweet the bottom line—that's it!!) I'm not talking about posting the results of work on the web or even about creating a coproduction space that connects you to your customers or taking some discrete segment of your process and webbifying it. I'm talking about actually performing your tasks in web applications that can theoretically be connected to an infinite number of other nodes--in the case of our story, neighboring electric utilities, other related processes, other smart machines, and other humans. This requires designing your work explicitly to be performed, as much as possible, collaboratively, transparently, in real time, and in multidimensional space.

As an old college debater, this is an extremely attractive rhetorical position for me. Rather than present the capabilities of Web N.0 in diaphanous, new-age terms that only invite the snickers of Real Women, the case I am making is that high reliability, high risk organizations need to embrace these concepts to further increase their excellence ratio. So it's true you see that Real Men do Collaborate.

8 ----------

So what are the specific advantages of redesigning work so that it is performed in parallel networks, through dynamic interactions, with maximum transparency and

collaboration, and, by design, minimum control. I say minimum control by design because even in systems where we think we have a lot of control, as was the case with the electrical system in 2003, well you see, we really don't.

The value propositions for inherently web-based work fall into three categories, I think. The first set of advantages support how organizations deal with the normal. NEXT SLIDE

I think some of these are obvious so I won't elaborate much. Dealing with complexity is one, however, that deserves some discussion. The Age of Complexity is a cliché, but like all clichés, it is a cliché because it contains truth. Individuals are not able to understand the causality and possible interactions of complex systems because an individual human just cannot scale to that level.

Karlene Roberts from the Cal Berkeley School of Business notes that as a world we keep developing more and more complicated technologies that require more and more complicated organizations. On top of this complexity, I would argue, we have layered a temporal dimension.

The average time to translate a basic technical discovery into a commercial product in the early part of the 20th century was 30 years.

Today our new technologies get to market in 2-3 years and may be obsolete in five.

John Seely Brown talks about how the pace of change today is too fast to allow for institutional development. How then will managers be able to design the optimum system?

There is some hope, however, that an interconnected system is better able to learn, understand, or perhaps it would be better to say, sense and adjust to its own complexity. By definition, an organism has to scale to its level of complexity or....it falters and eventually dies.

Our present organizational models are built on the failing assumption that individual control nodes can understand the complexity of a large system. Future organizational models will assume that only a system with the property of frictionless interaction and sense making can hope to understand itself. And you know what, there will be some real unexpected, and perhaps in a truly existential way, unpreventable catastrophes. Just last week, T-Mobile and a Microsoft subsidiary had to announce that Sidekick data was permanently lost in the cloud.

Reducing the need for redundancy is also an interesting point. Reducing the need for redundancy is also an interesting point. Redundancy is often cited as an

important technique for high reliability/high risk organizations wishing to avoid accidents. Redundancy is expensive, however, and as our blackout story illustrated, you can never have too much of it. Awareness, shared roles and responsibilities, and collaboration can serve some of the same purposes as redundancy for much less cost. Although I like this point because it is inelegant, I worry about it in fact because it is elegant. Fundamentally, I don't think we can be precise enough in our understanding to ever eliminate redundancy.

When you webbify your work, you also gain the potential to build real maps of your work processes. You know who contributed when, who lagged when, and can track process patterns that precede success or failure. This is what I mean by modeling success and failure. I believe this could lead to interesting new value propositions and business models that consultants will be able to use to make lots of money.

No matter the organization, we all want to make better use of talent. But most of us probably suffer from horribly archaic job definitions and roles that don't fit any specific talent particularly well. By freeing people to contribute to projects across organizational boundaries--actually by eliminating these organizational boundaries--we will be better able to apply our talents where they will contribute to the most good.

Another related quality enabled by web platforms is the virtualization of work. One employee, in ten minutes, may be able to contribute a solution to a problem vexing another section. Our current modes of organizing work do not support these micro-contributions, but web work platforms supported by the right protocols will make such kinds of micro-contributions easy. These types of micro-activities are already happening on the Internet in other contexts--think microfinancing and of course Twitter.

Finally, networked organizations should find that their workers will have a greater sense of ownership of the entire mission space. Isolated roles equate to uninvolved workers. "It's your network."

9 ----------

The second category of advantages lies in helping organizations deal with the unusual, with the emergencies.

First, networks--particularly ones rich in diversity--are notoriously good at dealing with weak signals. First Energy engineers didn't notice some of its lines were coming down, but other energy companies with greater dependency on a particular line, did notice. The fact of the matter is that what are weak signals to a large organization are always a strong bop on the head to someone else in the

system. A networked organization allows these sensitive pickups to more easily flow through the whole.

One of my favorite lines these days is that we knowledge workers have an irrational attachment to rationality. We like to assume order and rationality, when in fact irrational or unexpected behavior is, dare I say it, much more common. This is really just a sub-category of the weak signal point, but worth calling out.

Finally, if you recall when I was talking about aviation's safety record, I noted that aviation practices every day with every flight, the most dangerous parts of its process--landings and takeoffs. Most organizations, however, do not have that happy circumstance.

In fact the way they are organized --in stovepipes designed to deal with the normal--prevents them from practicing the really challenging parts of their work on a regular basis. I think that NetWorked organizations will find it much easier to bring people together to simulate an emergency scenario--because that's how they will do their work every day--together.

Furthermore, by operating and collaborating on a regular basis, the various parts of large organizations will, by bantering and chatting, create trust and connective tissue that would otherwise not be there and which are necessary to handle stress and strain.

You know I've used the aviation example a lot, and Greg Lloyd, when he was reacting to my remarks, noted that if my points are apt for a social-mechanical system, think how much more relevant they are for social-social systems.

I told him I would use that but it leads me to think of a recent Aviation disaster that clearly was not practiced--The Air France accident in the South Atlantic, and that represents a clear breakdown in social mechanical systems. If you recall the very modern aircraft was sending signals to the home office about its operational conditions that are now being pieced together to try to solve for the accident's cause. What I ask myself is why these readings were not available to the cockpit crew? It was obviously a question of design. It's an interesting question, isn't it?

What it makes me think is that as smart machines are integrated into work, we humans are going to have to get with the program. The digital machine doesn’t care about steps or protocol. It will be spitting out the data all the time. Can we organizationally step up to that challenge?

A point I haven't explicitly called out, because it runs in the background of many of these dynamics, is the ability of realtime sense making and collaboration to reduce or even eliminate LAG.

10 ----------

As I mentioned earlier, I've recently become fascinated with the question of mental errors. What are they really? And why do they happen?

I don't have the answer but would note that in a networked environment designed with transparency from the get go, mental errors should be easier and faster to identify and correct. Like our engineer who went to lunch and didn't turn the state estimator back on. A sentient network with the ability to determine whether key components are operating normally would have "noticed" the system was sleeping and probably been able to turn it on automatically. (This, by the way, will probably always be easier to do with machines than with people.)

All too often today in organizations, knowledge is assumed to exist in the hierarchy--otherwise why does it reserve the right to make certain decisions. Or knowledge exists in some database or manual. Knowledge management or capture solutions often don't work very well because the employees, just when they thought their job was done, have to now take several extra steps to record their learnings. This generally only leads to resentment and mumbling. But in Web 2.0 enabled work, the knowledge can actually exist explicitly in the relationships--the recorded interactions-- that enable outcomes, and these interactions can be captured as a matter of course in the system's history.

I love to cook and whenever I embark on a new dish I search the web to compare the recipes. But the very best content/knowledge is contained within the comments of people who have tried the recipes and the banter back and forth when they discuss the pros and cons of different approaches. This is in fact how knowledge needs to be captured at work. And it is this process that will, quite literally, create exponential wisdom. So, to return to the Buck Rogers opening, collaboration at scale and with diversity does increase your range of understanding. Perhaps they wouldn't have had such huge levers.

Finally, I think the transition of work from steps to flow will allow us to discard metaphors which simply weigh us down with nonproductive imagery. Let me just pick on one here: weakest link. Weak links are particularly dangerous when the system lacks transparency. In fact a really bad scenario is when you have a connected system that is not transparent. (In some ways that describes the Northeast electrical grid in 2003.) Weak links can go unnoticed. Unfortunately, traditional management concepts may very well reduce transparency in a system just to preserve the concept of hierarchical control. But transparency and collaboration actually should dramatically reduce the consequences of weak links, In a connected, collaborative, and transparent system, you are in fact potentially as powerful as your strongest member.

11 ----------

All transitions are difficult. That is in fact why no one likes change. And a lot of learning and habit changing must be done, both within organizations and by individuals, to adjust to the demands of work that flows.

The transition will be particularly difficult for managers, I believe, who will have to reconceptualize much of what they do. An important component of most work design today is a periodic pause of some kind to allow for a management intervention, usually aimed at quality control. If we're right about the power of the network, much of the quality control will occur organically and dynamically.

Quality control will be a property of the process, not an add-on. I do not think it will be trivial for managers to learn to provide value in flow rather than in steps. You invariably hear about managers who are good subject matter experts vice managers who are good with people. Because work as it is now organized often allows or even facilitates such bifurcation, you end up with too many who are one-dimensional. I think the new designs of work will call for both skills continuously plus some additional, yet-to-be-defined expertise in network health.

I can't provide the answer here, but I think a lot of consulting horsepower will be devoted to this issue in the future.

Nor do I believe staff will get off easily. My experience is that, even when given the opportunity, employees have difficulty speaking up, dealing with each other honestly, and taking risk beyond their small areas of responsibility. The new dynamic of work will demand such behavior. For some, however, who are already practicing this behavior in the give and take of social networks, the change will come more easily.

But we cannot underestimate the stress this will cause. I was reading this weekend about the alarming instance of suicides among French telecom workers--many of them more senior--who apparently were affected by a significant transformation process underway. WE are, I think, talking about a secular change in how work is done that will likely be accomplished in less than a generation.

I think the way we define job roles and responsibilities today is archaic and nonproductive. As collaborative networks develop, new patterns of value contribution will emerge. We will need to find ways to adequately track and reward micro contributions. The kinds of behavior that make a network hum, at least in some cases, will be quite different from what we're used to.

Let me give you an example. I saw a presentation recently on how game design has been used by the Institute for the Future to crowdsource problem solving. Using the principles of game design, which is a whole other topic worthy of its own talk, the Institute designed attributes by which to measure successful contributors.

One of these attributes was "dark imagination"--yes, that pretty much means what you think it means. I don't know too many organizations today that reward employees for their dark imaginations but such new types of contribution will in fact emerge from the interaction of social networking technologies with work.

Eventually consultants will need to devise new concepts to evaluate the effectiveness of massively paralleled organizations. But my guess is this will prove quite difficult. Organic, dynamic work processes will work quite well in practice, but creating grand theories as to exactly what is happening, I predict, will be more elusive. Eventually, when you don't quite understand why something is happening, you find it difficult to solve significant disruptions. Lacking a strong theoretical framework, you are forced to solve problems through trail and error.

12 ----------

The other day someone was telling me the story of some new initiative she was advocating and how the manager she reported to wanted to see how it would all work.... She needed the whole thing laid out in front of her before she could approve it. And my friend said with some frustration.... that's impossible. I cannot design it in detail--the individuals involved will do that.

As managers and leaders, we need to abandon the hubris that we can design the perfect process or organization. The best we can do is facilitate a process from which can emerge goodness and, over time, excellence.