Chapter Three

WHAT INFORMATION ARCHITECTURE FOR DEFENSE?

Martin C. Libicki

We live amidst an information revolution, which is to say, a revolution in the capabilities of information technologies and infrastructures. The quality and quantity of the information we receive have greatly increased, but for information to be truly useful, it must improve the quality of our decisions, which, in turn, are judged by the quality of the resultant actions.

Two-thirds of all personal computers and almost all networks and databases are used for business, not recreation. Plausibly, therefore, most decisions that information technology supposedly improves are those made in an organizational context. They result, one way or another, from interactions among people and their machines. As such, the quality of the decisions an organization makes is increasingly related to how it constructs its information systems. These architectures are often faithful reflections of the tacit assumptions about power and purpose held in these organizations. Architecture and organization are linked. An organization's tendencies shape its architecture; its architecture, in turn, helps shape its culture.

The theme of this chapter is how organizations respond to the opportunities and challenges that come from giving everyone access to rapidly increasing amounts of information and corresponding tools, such as analysis and communications. Should sensors be few and commanded from up high, or many and commanded from the trenches? Should information be pushed to the user based on what is deemed necessary, or should users be able to pull information and subscribe to data-flows they feel they need? Can information be arranged and rearranged by users to best fit their intuitive grasp of the

matter, or does one display fit all? Will bandwidth constraints limit the information and services users can access, or can a robust menu of alternatives be employed to expand or work around such obstacles? How much should be invested in providing the capability needed to discover information in realms ostensibly far from one's learned domains? Must security be enhanced by limiting what warfighters can know of the battlespace? Should interoperability be good enough to permit users to seamlessly skip from one domain to another, plucking what they need from what they can see? Will users come to feel they own their tools?

WHAT IS ARCHITECTURE?

A common-sense definition of architecture holds it to be the relationship of a system's components to each other. But since people are elements of almost all information architectures, architectures involve human participation. Like any policy structure, architecture reflects power relationships. Unlike many policies, though, it can alter the social underpinnings of power relationships.

The subject of information architecture has more of a future than a past. All organizations have an information architecture, but before there were information systems, the fundamental principles of an organization's information architecture directly reflected (1) command relationships and related functional responsibilities, (2) geographic distribution of personnel, (3) the distribution of clearance levels and related physical access privileges (e.g., keys, safes), and (4) personal relationships. In other words, information architectures were a direct reflection of management and institutional relationships and could be almost completely studied in that context.

The development of the telephone and telegraph altered this formulation by removing some of the geographical impediments to communications. Complex machinery with its sensors and controls possessed its own information architecture, but it was generally internal to the machine itself. Rarely were these machines instruments of an organizationwide architecture. Prior to World War II, the percentage of all workers with phones at their workstations was low, and the ability to control machinery remotely was negligible.

The entry of automated data processing systems in the 1950s created specific work flow paths for certain types of information—e.g., payroll processing and inventory management. Most people accessed computers through special-purpose terminals that limited their interactions to predetermined processes. General access to corporate information was rare except within top levels of management. Indeed, for the first few decades of their existence, computers tended to have a centralizing effect in that they permitted management to harvest much more information about its enterprise.

The development of recognizably modern information architectures began not much more than a quarter-century ago, when sophisticated measurement and control systems were coupled with workstations cheap and robust enough to proliferate outside computer rooms. Widespread personal computing in the workplace is no more than 20 years old; networking, no more than 15; and general Internet connectivity, no more than 10 (universities being the primary exception). The advent of pervasive information systems has so greatly changed the quantity of information available to people that it has changed the quality and hence the nature of the architectures.

Organizational information architectures are still evolving. Office workers in most modern Western organizations have workstations with intranet and Internet access. Most production and paperwork processes are monitored throughout their journey and have their minute-by-minute information logged. These data are widely, but not necessarily generally, available. Corporate America is in the midst of switching from phone and paper to the Internet as its primary link to suppliers, collaborators, and customers.

The next decade or two should bring further complexity and richness in organizational architectures, thanks to several trends:

- Access to both the Internet and organizational intranets through initially low-bandwidth and always-on small-screen mobile devices, cell phones, or cousins of the Palm Pilot VII™.
- Greater systemization of institutional knowledge bases to the point where such knowledge can be analyzed with data mining and other forms of logic processing.

- The growing ubiquity of small sensors and smart radio-frequency bar-coding both in process industries and in public,[1] and in quasi-public settings (e.g., hospitals, schools).
- Further increases in the extraction of usable knowledge about customers and the overall environment from external sources both public (the Web) and quasi-public.
- The granting to suppliers, collaborators, and customers of deeper access into one another's corporate knowledge bases.

The Global Information Grid

DoD's interest in information architecture is clear. Joint Vision 2010 and Joint Vision 2020 are paeans to the importance of information superiority to military superiority. Many of the technologies used to construct information systems were developed by or for DoD. For a long time to come, DoD's information architecture will revolve around mobile devices and sensors of the sort that corporate America is only starting to see. But the levels of automation within DoD vary greatly, partially because capital turnover is slower there; large parts of DoD are more digitized than is corporate America. It is a cliché to note that DoD has no information problems not shared by private enterprise, but there are real differences between the two.

DoD is beginning to contemplate an enterprise architecture for itself. It already has bits and pieces under way within service systems (e.g., the Army's nascent Force XXI initiative and the Navy's IT 21 project), joint internets (NIPRnet for unclassified work, SIPRnet for secret material), and its software suites, the Global Command and Control System (GCCS) and the Global Combat Support System (GCSS). The term *global information grid* (GIG) is often used to describe the anticipated agglomeration, variously known as the global grid (the 1992 JASON study), the system of systems (former Joint Chiefs of Staff [JCS] Vice Chairman Admiral William Owens, but also former

[1]The events of September 11, 2001, are likely to sharply increase the use of identification cards, such as proximity cards, for automatic logging of entries and exits.

Army Chief of Staff General Gordon Sullivan), and the battlespace infosphere (U.S. Air Force scientific advisory board).

How far along is DoD on its architecture? DoD typically divides architecture into *systems architecture* (what is "wired" to what), *technical architecture* (how interfaces are defined), and *operational architecture* (how data flow to carry out missions). Systems architecture exists, by definition, at every point, but whether it accords to its specifications is a different matter. In practical terms, the shortfall between how DoD's networks are configured and how they ought to be configured given equipment constraints is probably modest. Technical architecture also exists, but the persistence of legacy systems means that the actual state of interoperability falls far short of where it would be if all systems complied with the joint technical architecture, which, itself, covers only a small fraction of what it ultimately has to.

Operational architecture remains largely undocumented, although there are calls to make it more formal. Systems architects could determine where to make investments if they knew where the data were supposed to flow; technical architects could determine where interoperability was most critical if they understood where information had to be exchanged and understood. So what is the hang-up? Operational architecture sounds easy, but it is about such elemental issues as who talks to whom and which processes need what data—which is to say that it is about power, and power is extremely difficult to negotiate and codify in peacetime. Moreover, the nagging feeling persists that, to echo Von Moltke, no operational architecture will survive contact with the enemy. War generates surprise. Many systems fail in combat; others succeed brilliantly but are too few in number.

Yet all three forms of architecture somehow fail, even in combination, to capture the essence of architecture, the constructs under which people exchange information. Physical architects enter school believing that architecture deals with the arrangement of structural elements; they leave school understanding that its subject is really the arrangement of spaces within which people interact. And so it is for information architecture.

Need There Be Architecture?

One approach to architecture posits a hierarchy of tasks.[2] Physical architecture begins with a statement of requirements (the "bubble drawings" that rough out room arrangements) that is successively refined through architectural drawings, construction plans, subcontracting processes, and bills of material. Information systems are, likewise, successively broken down from task requirements to implementation along three parallel paths: data structures, algorithms, and communications networks.

The Internet was built on entirely different principles, even if few people remember that its first purpose was to share computing power too expensive to duplicate everywhere. Data structures (except for addressing), algorithms, and even routing are now entirely beside the point. The Internet does not *do* anything. It is an infrastructure used by a great many people to do a wide variety of things. The Internet does have a straightforward and explicit set of rules that govern how networks become members and format their message envelopes. Otherwise, if architecture is defined as who-can-say-what-to-whom, the Internet has very little of it.

The Internet works famously, so why bother with architecture at all? Why not simply let everyone make information according to their abilities and take it according to their needs? Why limit the resulting conversation? There are several reasons.

First, the Internet is primarily a transport mechanism and hence only part of an information system. To accomplish specific tasks within specified parameters, organizations have had to impose an information architecture on top of it. Even the research community, where the ideals of the Internet are most fully realized, has architectural elements embedded into its culture, such as how material is published.

Second, the Internet's survival for most of its life has been based on shared norms and the kindness of strangers. As it has opened itself up to literally everyone over the past five years, it has faced tough

---

[2]J. A. Zachman, "A Framework for Information Systems Architecture," reprinted in *IBM Systems Journal,* Vol. 38, No. 2, 1999, pp. 454–470 (originally printed in 1987).

problems in privacy protection, intellectual property rights, spamming, and poor security.

Third, even if the principles of the Internet and World Wide Web are the goal of an information infrastructure, such ideals must be made manifest in a world where truly open organizations are rare. Yet not every organization has an information architecture per se. RAND does not have one for its research, although it does have a culture. RAND's ability to get work done depends on researchers acquiring information from outside the corporation and analyzing it using tools they know very well. In effect, RAND rides on the information architectures of its clients and the research community.

Architecture Follows Culture?

How cultural norms affect people's participation in the information era is a big and well-studied subject. But it helps to point out some general questions that suggest there are large variations in how different cultures use information. Intuition says that cultural factors can help or hurt an organization's use of information systems—especially people's willingness to trade information, seek out potential disquieting knowledge, undertake honest analysis, and base decisions on the results of that analysis. Other factors, such as the tradeoff between horizontal and vertical flows, between public and private credibility, and between oral and written communication, ought to affect information system design as well. Part of what DoD (or any organization) ought to do in developing an enterprisewide information architecture is to examine its own cultural assumptions and the ways in which they may feed into architectural decisions—a process akin to warmup exercises before undertaking a long run. To understand any given culture, one should ask

- Are people more likely to hoard and then trade information than they are to share information based on trust?[3]
- Are people more likely to believe what they get from public sources or what their friends tell them?

---

[3]Frank Fukuyama argues that cultures can be characterized as high- or low-trust. See his *Trust: The Social Virtues and the Creation of Prosperity,* The Free Press, New York, 1995.

- Does more information flow vertically or horizontally? Is vertical collaboration accorded the same weight as horizontal coordination?
- Is information passed via the written word or the spoken word? Which form conveys more authority, granted that this often depends on personal preference?
- Is there a bias for acting and letting the facts flow from the results rather than undertaking analysis before acting?
- Must decisions be justified by facts and arguments, or do authority, experience, and/or charisma suffice? To what extent are formal credentials taken seriously as a source of status?
- Will people seek out knowledge even when it may contradict their earlier judgments—especially publicly offered judgments?

Underlying the premise that information systems based on sound architectures can help an organization is the requirement that an organization be sound and coherently organized. A few test questions to make this determination might be:

- Is the organization in the right line of work? Is it looking to the right customers and solving the right problems? (Solving the wrong problems more efficiently is of only modest benefit.)
- Is the achievement of organizational goals accorded a higher priority than the achievement of any one faction's goals? of personal goals?
- Is honesty accorded respect? Can the organization accept bad news without shooting the messenger or inventing fantasies?

If the answers are affirmative, information and thus a good information architecture can help. DoD seems to score above average on all three questions.

DoD as an Institution in Its Own League

DoD differs from other institutions, so many of the issues raised concerning its architecture may be of less consequence to those other institutions. The differences are many.

DoD is a hierarchical organization that relies on command-and-control (especially in the field) irrespective of the distribution of knowledge, and it is likely to remain so regardless of how technology evolves. Warfighting puts people at personal risk, giving individuals a potential motivation that can be at great odds with their organization's motivations. Militaries are built to engage in contests and to serve the national will. Most other organizations, by contrast, respond to whoever will pay them; most customers are individuals. DoD has multiple layers, each with its own assets and requirements, and each strongly biased toward owning its own information sources.

Because DoD needs to know the whats and wherefores of uncooperative foes, information collection is expensive and problematic. Sensors are expensive, deception is the norm, and militaries are obsessed about information security. Businesspeople at least can rely on the force of law to inhibit mischief by competitors; militaries work in an anarchic milieu and thus cannot.

Because DoD works outdoors, it relies heavily on radio-frequency links, often established in austere regions (e.g., at sea). Nodes are often at risk of destruction; communications is ever subject to jamming, so transmission and reception equipment must be hardened. Bandwidth constraints pinch DoD more than they typically pinch commercial enterprises. Combat communications are often urgent, and many of those who have most need of information have limited attention to give to what they see or even hear.

Finally, DoD cannot rely on outside sources for all of its education or institutional learning. Much of what it teaches has no outside counterpart, at least in this country, and often not anywhere. Bill Joy, the cofounder of Sun Microsystems, once observed that regardless of who you are, most of the smart people in the world work for someone else. DoD can occasionally be the exception to this observation. Most of the people who best understand, say, stealth technologies work for DoD either directly or indirectly. This also means that if the state of the art in areas of DoD's interest is to advance, those areas must be deliberately resourced by DoD itself.

Perhaps the most interesting facet of DoD's unique requirements is its imperative to optimize not efficiency, but adaptability—particu-

larly now, when it is so difficult to make reliable statements about where or against which foe the U.S. military will have to practice its craft next (as September 11 and its aftermath have once again proven). To illustrate this point, Table 3.1 categorizes warfare according to whether mass or precision matters more, and whether the United States gets to fight from standoff range or has to move close in. Every square holds a different form of war, each of which demands different features from DoD's architecture.

For the historic, conventional type of warfare—close-in mass operations—commanders need data from the field to build and share almost-instant situational awareness. This information fosters command dominance, which, in turn, permits its possessors to execute decisions faster than enemies can react. The military, however, requires an even larger quantity of *internal* data to exercise its core competence of conducting highly complex maneuvers. Carrier battle groups, theaterwide air tasking orders, and maneuvers in corps are all examples of how the U.S. military can orchestrate the actions of 10,000 to 100,000 individuals better than any other entity can.

For strategic warfare—standoff mass operations—commanders need a more analytic and synoptic understanding of the enemy's pressure points. One method is to learn how the enemy's economy, political structure, and infrastructures (both military and civilian) are wired so as to be able to identify the nodes. This invariably becomes a large modeling exercise.

For hyperwar (DoD's preferred mode), the most pressing problem is to locate, identify, and track mobile targets so that they may be struck quickly from afar. The important qualities are the ability to scan large battlespaces, sift the few interesting data points from the mass of background, sort the targets of potential by priority, and strike those targets while they are still visible. This search for in-

Table 3.1

A Two-by-Two Categorization of War

	Close In	Standoff
Mass	Conventional conflict	Strategic conflict
Precision	Mud warfare	Hyperwar

formation is akin to finding the classic needle in a haystack—or, because targets continually move, finding the snakes among the worms.

Finally, if despite everything DoD is stuck in mud warfare and working within dense milieus, it needs a precise understanding of its environment so that it can distinguish the malevolent from the irrelevant. A precise situational awareness also permits threat patterns to be perceived.

These are, admittedly, gross generalizations. But they illustrate the wide range of functions that DoD's information architecture must satisfy. They also illustrate the potential folly of building the GIG architecture from the top down. If the U.S. military cannot be sure of when, who, or how it will fight, should not its architecture be optimized less for any specific scenario and more to accommodate the fluidity of a real-time, high-density, rapidly changing and restructuring world?

For instance, is the GIG to be understood as a command, control, and communications (C3) system that ties people to each other (and, incidentally, to some interesting services) or as an intelligence, surveillance, and reconnaissance (ISR) system that uses communications to accommodate a distributed workforce? The two facets ought to be two sides of one coin, but the cultures they support differ greatly. Communicators like to share information; intelligence types tend to hoard it. Communicators want to know who talks to whom; intelligence types want to know what information is fused with what.

Who is the ultimate customer: the White House or the foxhole? The intelligence community has been facing this question since the Cold War ended. Initially, both urgency and the scarcity of good information oriented intelligence to the U.S. president. Now, with the Cold War over and open-source information more available, intelligence oriented to the foxhole makes more sense—but how much, and how fast? This question is felt with great keenness in peace operations. Is the primary point of military operations to generate an outcome consistent with U.S. interests or to serve clients (e.g., those who live in an erstwhile war-torn land)? The answer sets up a broader question: Should such a system be designed to spread information down and around or to filter it up?

The issue of who owns information raises command and control (C2) issues. Ownership duly and dually implies both responsibility and control. But is the responsibility to collect information logically connected to the right to control not only who gets to see the information, but also the form in which it is released? That case is hard to make in the age of the Internet and the Web. It would seem that information should be generally releasable, subject to no-more-than-necessary security constraints, broadly accessible, and formatted in the most interoperable way—which often means with the least amount of unnecessary processing. But can potential users counter the collectors' argument that data cannot possibly be released to the rest of the world without being subjected to thorough and time-consuming analysis?

ELEMENTS OF ARCHITECTURE

In this context, eight categories can be used to describe architectural issues related to information: (1) collection, (2) access, (3) presentation, (4) networking, (5) knowledge maintenance and management, (6) security, (7) interoperability, and (8) integration. Each is discussed in turn.

Collection

How an organization gathers data depends on technical considerations as well as its judgment about what kind of information is needed by whom.

What, for instance, should be the mix between hunters and gatherers? Hunters define information requirements and then collect in accordance with them; their questions are specific, and their tools are often focused and may be used intermittently. Gatherers collect large amounts of data and then thresh through what they have for anomalies, telltale changes, and other interesting tidbits; their tools tend to be general and used continuously. The intelligence community has both types of users. The imagery business tends to be filled with hunters, with tools that are continuously busy (e.g., reconnaissance satellites) and tools that are employed only in discrete missions (such as U-2s). The signals intelligence business tends to be filled with gatherers.

Is information to go up or down the logic tree—the oft-cited path connecting data to information to knowledge, understanding, and wisdom? Some use induction, starting with examples and winding up with generalizations. Others prefer deduction: they start with rules and then exploit them to develop differentiations. The first approach is good for figuring out the enemy's doctrine, the second for distinguishing an enemy from a bystander. Inductors start their post-processing late in the game; deductors early.

What integration metaphor should be exploited? One approach (and many may be needed) is scan-sift-and-sort, which is particularly useful for engaging mobile targets amidst clutter. Another challenge is building a synoptic picture from the coordination of distributed sensors. Managing by maintaining parameters (e.g., the odds that a district is secure, the likelihood that the enemy will do X under stress) requires that new facts be consistently integrated into old equations.

How are data validated and reconciled? The Army says the village has 100 enemy holed up, the Marines say 50—which number goes in the database? What criteria are used to point users toward one or another estimate? What metrics are used to label facts as being of greater or lesser validity?

For DoD, many of these questions are arising as its GIG becomes more deeply networked and more inundated by sensor data. The greater the bit flow of any one phenomenology, the greater the logic to move from hunting to gathering successively lower-grade ores simply because it can be done. Many of the more intractable problems of warfighting, such as tracking targets in clutter or mobile emitters, are problems of gathering. Issues of data integration and validation arise as the volume of data grows apace and such techniques as cue-filter-pinpointing, automatic sensor coordination, systematic parameter maintenance, and even the computer-aided marriage of scattered knowns and needs all become more feasible.

For DoD, data collection is increasingly an issue of sensor deployment and management rather than, say, human reconnaissance.

For instance, should sensor systems be designed to collect broad knowledge and to cue weapons that then find the target, or should sensors collect knowledge precise and timely enough to guide weapons to their target? As Chapter Four, "Incorporating Information

Technology in Defense Planning," suggests, the choice between man-guided, seeker-guided, and point-guided weapons is not simply one of engineering, but of empowerment as well. A third party given man-guided weapons has to be operationally competent to make use of them. Given seeker-guided weapons, the third party can do what it wants. With point-guided weapons, it will have to depend on those who supply the points and tracks, so if that "who" is the DoD, the United States retains a great deal of leverage over how the weapons are used.

A parallel question is whether DoD should lean toward using a few expensive sensors (such as billion-dollar spy satellites, Aegis radar, Joint Surveillance Target Attack Radar System [JSTARS]) or many cheaper ones (such as micro unmanned aerial vehicles [UAVs]) or disposable unattended ground sensors. If expensive sensors are used, there will be contention over who gets to use their capacity. A system of cheaper but more numerous sensors lets all operators have their own capability, but what happens when there are more sensors than humans can manage? How far can and should the sensors be designed to manage themselves?

Finally, should sensors output (1) raw data, (2) data cleaned up by some automatic and semi-automatic processing, (3) only data that are completely exploited and have passed quality-control tests, or (4) only data that have been fully fused with all other comparable data? How close should the coupling be among sensors, sensor data storage, sensor post-processing, and sensor data display?

Access

Over the last half-century, management attitudes on information have shifted from need-to-know to right-to-know. Many companies, for instance, have found labor unions to be more tractable if workers are shown detailed information on the company's financial performance. As a general rule, the military has lagged on this issue, although DoD is gradually giving people more access to its information

both internally and externally.[4] For example, the end of the Cold War reduced many classification levels to Secret.

Many reasons have been offered for limiting information. Two of them—the cost of making legacy systems interoperable, and limited bandwidth—have technical solutions. The security argument is a shibboleth to the extent that almost all career military personnel are cleared to the Secret level; *operational* information is rarely classified higher unless sources and methods are involved, but they can usually be scrubbed out of the data.

That leaves one real reason for restricting information: it might be distracting. DoD has gotten great mileage out of warfighters with modest education by carefully developing a hierarchy of task structures to which operators are closely trained. Information requirements are then carefully generated for these tasks. Unfortunately, the narrow task specialization that pervades the national security establishment (only 1 percent of all the intelligence community's workers are all-source analysts) is increasingly at odds with how information is handled elsewhere: workers are given an increasingly broad view of their task within an organizational context.

Another possible reason is that warfighters would not risk their lives in dire circumstances if they understood how bad things really were. Such a logic may apply to other militaries, but the U.S. military takes public pride in the competence and professionalism of its warfighters, who expect to be told the truth and take responsibility for keeping focused.

What gets shared beyond DoD? Many of the arguments made against sharing data with coalition partners follow the logic above. For longtime partners, such as Britain and Canada, this argument rings hollow. For those who are made coalition members for political rather than military reasons (e.g., Syria in the Gulf War), a certain hypocrisy that separates the promise to share from the actual information supplied can be expected. How much of DoD's information should

---

[4]After September 11, 2001, however, public access to civilian information appears to have been restricted because of the fear that terrorists could use it to prepare attacks; see Ariana Eunjung Cha, "Risks Prompt U.S. to Limit Access to Data," *The Washington Post,* February 24, 2002, p. A1.

be publicly accessible? In the last five years, DoD Websites criticized as too revealing have cut back. "Guarded openness"[5] may permit U.S. forces to acquire the cooperation and facilitate the coordination of third parties, or at least to get their input for intelligent preparation of the battlefield.

Presentation

In many cases, information flows and needs cannot be matched unless they are formatted in some standard way, which, in turn, means they must be encapsulated and categorized according to some semantic and cognitive structure.

The choices information providers face regarding how to forward information to the field are rich ones. E-mail, pop-up alerts, and monitors are examples of push processes; Websites and query capabilities typify pull processes. A mix of the two types of processes is possible. In a guided tour, for instance, a broad request yields a linked set of answers, elements of which can be invoked based either on the user's interests or on abilities made explicit by the user. A more sophisticated version would infer what the user wishes to know and push selected information (both stored content and news as it develops) based on the specific inquiry, past inquiries, and perceived user habits.

Information architects endlessly repeat the mantra "the right information to the right person at the right time," also adding "with the right presentation and the right security." But who is to judge what is "right"? Some analysts believe that military missions and tasks can be decomposed so that the right information pops up as needed, much like an automobile navigation system builds trip routes and flashes left or right arrows as key intersections are encountered. Users of good operating systems, such as that of the Palm Pilot™, may delight in having the system bring up the right menu of choices whenever one or another action has taken place. There is an ele-

---

[5]Carefully vetting what is released, but then releasing it without restrictions. See John Arquilla and David Ronfeldt, "Information, Power, and Grand Strategy: In Athena's Camp—Section 2," in John Arquilla and David Ronfeldt (eds.), *In Athena's Camp: Preparing for Conflict in the Information Age,* MR-880-OSLVRC, RAND, 1997, pp. 413–438.

gance to good design; this is why people hire architects and try not to draw their own maps.[6] It helps eliminate clutter; it also permits information providers to draw out implicit narratives from media, such as maps, which only the skilled can navigate. Pushing information at users relieves them from having to make explicit choices. It can be valuable when users are overworked or under stress.

To the user, *how* information is presented is a critical adjunct to what information is presented. Simply by highlighting certain information or certain relationships among data elements, a person or program can push everything else to the cognitive background. In a complex environment, data that are not highlighted might as well not exist. But who decides? "Shared situational awareness" is one thing if it means that everyone can access the same information. But should there be a standard method of presentation? If there is, how easy should it be for users to redraw their own map, so to speak, if they can even do so at all? The Common Operational Picture (a GCCS application that indicates the presence of major units on a battlefield map that everyone shares), for instance, permits certain predesignated layers of information to be highlighted or hidden, but it provides no macro language that lets intelligent users (or their information aides) manipulate data fields of unique importance to them. People do, after all, see information in different ways. Having everyone see everything alike has two advantages: it promotes efficiency in conducting complex operations, and it minimizes complaints such as "How come I can't see this?"

But what if the problem is not complex coordination but complex recognition, as may be the case for modern warfare against an elusive and well-hidden foe? Users might usefully be allowed to peer into the weeds in multiple ways on the theory that it may take only one person's burst of insight to find the elusive foe or make everything fall in place for everyone else. Free play exploits the vast differences in how people perceive the world. Making everyone look at the world the same way may not hurt those whose perspective is sufficiently near the official norm, but it may reduce the contribution of those who think differently. Even if people thought similarly, there

---

[6]See Edward R. Tufte's *Envisioning Information,* Graphics Press, Cheshire, CT, 1990; *The Visual Display of Quantitative Information,* Graphics Press, 1992; or *Visual Explanations: Images and Quantities, Evidence and Narratives,* Graphics Press, 1997.

might still be a point to having them learn how to manipulate their perspective on their own until the world comes into focus. They may learn more about perception and pattern recognition for having discovered as much on their own. In an era of information overload, simply going through the exercise of determining which few features of an environment together tell the story is a valuable lesson in recognizing the essential. So presentation, like its close cousin access, also depends on who decides.

Networking

Access often depends on how much bandwidth is available. In a world of thin 9,600 bits-per-second (bps) lines (the top speed of the new Single Channel Ground and Airborne Radio System [SINCGARS]), warfighters can demand neither imagery nor access to databases. Since compression of one sort or another has to take place at the information source, or at least at the last cache before the fiber ends, the prejudice against user choice is easy to explain away as a technology-driven requirement. But is it?

Mobile and fixed access to the Internet differ sharply. Europeans have taken to mobile phones faster than Americans, in part because they settled on digital cellular standards earlier. The race is on to serve them Internet and Web access through small phone screens, although Palm VII™-like devices may be part of the mix. Because cell phones know where they are (to within roughly 300 m), investors are excited about tailoring content to location (pass an ice cream shop, and a pop-up message about new flavors appears on the screen). Americans, by contrast, are more likely to surf while sitting still; they have more screen space and bandwidth to match. Content cannot be presented the same way to both the desktop and the handset. Presentation for the latter must be more parsimonious, favoring text over images, sound over sight, and voice over tactile input.

The differences in information accessibility for mobile versus immobile users are critical for the U.S. Army. Division XXI, the Army's digital initiative, provides a rich array of information devices from the corps headquarters down to the battalion tactical operations center (TOC) using million bit-per-second connectivity for systems that support intelligence, maneuver control, and fire support.

Beyond the battalion TOC, however, soldiers are linked, at best, at 9,600 bps via an Appliqué—a limited terminal to which some information is pushed and from which certain numbered queries can be sent. If the ability to wield power is enhanced both by being close to the field and by being able to access information, the digitized Army appears to pivot around the battalion TOC commander, who alone has a good enough view of both worlds. Is this what the Army really wants? Is this necessarily the optimal configuration for the sorts of wars the Army is likely to fight? Urban combat, peace operations, and new theories of swarming[7] may demand that control be exercised by company commanders, one full level down. The Marines pivot around the squad and the platoon. Will reifying Army doctrine in the hard-to-change hardware and software of its information systems promote jointness or the ability to carry out combined operations with overseas counterparts?

Perhaps the issue is less bandwidth than how bandwidth scarcity is managed, for there are choices other than limiting the functionality of terminals to fit bandwidth limitations. Appliqués could be full-fledged clients that accommodate reduced bandwidth by squeezing information flows, transmitting images with less clarity, less color, less detail, and/or fewer updates, or even abjuring imagery in favor of symbolic transmission. Or, each radio's capability could be enhanced but bandwidth managed explicitly as a battalion-shared resource, so as to use more bandwidth for each radio but employ explicit contention mechanisms to allocate spectrum.

A similar architectural issue comes from the breakdown between the field headquarters and those back home who have copious access to information. If the major source of battlefield information comes from space, and space information comes in gigabytes, will there be an inevitable shift from the field toward those with sufficient access to see everything coming in from space? How will the nascent shift from space to field-supported UAVs affect this balance?

Traditionally, C2 was the metric for distributing bandwidth. Fat pipes connected the commanders in chief (CINCs) with the United States, and these pipes grew successively, if inconsistently, smaller as they

---

[7]For background, see Sean Edwards, *Swarming on the Battlefield: Past, Present, and Future,* MR-1100-OSD, RAND, 1999.

skittered down the ranks. Actual commands or reports are unlikely to overfill the pipes as much as imagery will, but the demand for imagery is not very correlated with rank. The current thinking tries to take operational architecture (who says what and how much to whom) as its foundation, but the calculations have proven vexing. The STU-3 (a telephone handset for making secure phone calls) was said to be the great weapon of Desert Storm because it allowed direct contact between the Pentagon and the field, and, in many ways, circumvented traditional routes of influence. There may be understandable reluctance to accord a patched-up set of relationships equal status with wiring diagrams on paper.

Knowledge Maintenance and Management

Knowledge tells one how to carry out tasks. But in a military context, knowledge also includes broad estimates of reality (e.g., this village is or is not safe) and operational rules of thumb (e.g., a village with few young men visible is generally a hostile one). In deciding how analytic judgments (such as how safe the village is) are to be maintained, one must decide who will have authority over what variables. If automatic processes rather than people maintain judgments, the choice of which variables are fed in and how becomes important.

In any complex environment, there will be many newly discovered facts, including new events, and parameters that change to reflect those facts. The methods used to feed new facts to parameters (and the algorithms by which new facts are ingested) are thus important. Sensors, for instance, may pick up new jeep tracks, a finding that could be used to update the odds that an adversary is working in the area, the specific equipment the adversary has, perhaps how much equipment, and maybe even the adversary's doctrine for using it (why were no efforts made to hide the tracks?).

Many computer specialists are starting to think about how parameters may be maintained by software agents—i.e., processing components passed among machines to perform functions for their owners. To use a biological metaphor, an agent (for a parameter: in this case, evidence of a country's interest in nuclear weapons) may wander among nodes that have receptors for various kinds of events—say, the hiring of engineers with nuclear specialties, market reports of certain chemical sales, and suspicious travel patterns. A match might

cause the agent to emit signals that stimulate other agents to activity (to review where the country is sending its graduate students) or lull others into relative inactivity (giving less weight to the country's acquisition of fertilizer plants).

Agents may be similarly used to generate alerts. News that a pilot has been shot down, for example, may trigger a requirement for local intelligence, spur greater liaison with friendly forces behind the lines, allocate more channels for the pilot's signals, prepare for undertaking search and rescue as well as medical evacuation, and establish stay-away zones for certain operations. Some second-order effects may need to follow automatically (e.g., telling rescue teams which substances the pilot is allergic to).

Engineering a system for agents raises difficult questions. How are they authenticated, made safe for circulation, and kept free from contamination? What access to databases, processor time, and bandwidth are they granted? How explicitly should their logic be documented? Under what circumstances can they trigger other processes, and how can their influence be overridden?

Knowledge management is the search for ways to transfer knowledge. Traditionally, analysis determined the one right way to do things, and that method was then fed to everyone. Although this made sure everyone worked to standards, it stifled initiative, demoted the acquisition of tacit knowledge, and left organizations inflexible to change from the bottom, acquired, in part, through front-line contact with new realities. A contrasting approach would be to liberate everyone to learn the best way to do things for him/herself, but this would lengthen learning curves, retard the diffusion of new knowledge, and lead to problems when skilled workers and engineers retire.

Enterprise architectures are also being asked to support knowledge management—i.e., the art of circulating know-how. Knowledge management was a key factor, for instance, in the U.S. Navy's ability to chase U-boats from the East Coast in 1942.[8] Each commander's

[8]See Eliot A. Cohen and John Gooch, *Military Misfortunes,* especially "A Failure to Learn: American Antisubmarine Warfare in 1942," Free Press, New York, 1990, pp. 59–94.

contacts were collected and used to build a database of enemy tactics and successful countertactics. What mechanisms are there to tell those with specific knowledge of problems about acquired knowledge? Writers on knowledge management emphasize the human factors (e.g., trust, random interactions) and the sense that the transfer of tacit knowledge requires physical contact.[9]

Can militarily relevant information be systematically organized? After all, its domain is limited (compared to knowledge as a whole) and routed through well-understood task structures. Much information is generated solely to do specific jobs and can be inculcated through training. Those who repair C-17 engines, for instance, can be handed all the information they need. Nevertheless, modern militaries are continually hiving off more specialized offices, few of which are sufficiently well defined to be exclusively self-contained. Even well-established entities generate new information sources as they wander afield from their original tasking. Also, not every useful data point is contained within the manual; the Web has revealed the power of horizontal communications as a device for getting advice, perspective, and solutions to unique facets of problems. If and as defense systems resemble their commercial counterparts—a clear trend in command, control, communications, computers, intelligence, surveillance, and reconnaissance (C4ISR) equipment but less so elsewhere—a growing share of knowledge on how to manage them will be outside DoD. Finally, open-source intelligence is growing as a percentage of all intelligence. Warfighting is an outside activity often in and among third parties. The ability to ferret information from the environment may come to depend on foraging through a thickening forest of digital data. The ability to harvest such information—regardless of its mixed quality—may provide insights denied to those who abjure such disorganized chatter.

To circulate information intelligently, one needs a way to find where it resides, which gets harder as complexity complicates the task of pigeonholing information into predefined categories. If information is organized in many ways (and much information can only be categorized and catalogued approximately), the possession of a reli-

---

[9]See, for example, Thomas Davenport, *Working Knowledge,* Harvard Business School Press, Boston, MA, 1998; and Nancy Dixon, *Common Knowledge,* Harvard Business School Press, Boston, MA, 2000.

able discovery mechanism is important. Although there are many philosophies for building and maintaining such engines for finding random information (especially machines that can be taught to extract meaning from text), the explicit tagging of a document's subject and key data appears to be growing popular.

Security

Information systems have security architectures that govern

- Who can read from or write to which files (or invoke which processes);
- How legitimate users are authenticated;
- How files and processes are protected from mischief;
- What procedures are used to detect when a system is under attack, respond to limit the damage, and recover functionality when control is restored.

Security's first question is, as always, How much? Systems with insufficient security necessarily depend on the kindness of strangers. But excess security costs money, hassles users, often denies service to the legitimate, and is prone to failure if users react by subverting its rules. Systems that deal with the public (e.g., e-commerce sites) must begin by assuming all users to be legitimate, only rejecting such claims after seeing bad behavior. This rule makes them heir to a flood of messages that are individually benign but collectively choking (e.g., the February 2000 distributed denial-of-service attack). Because DoD systems rarely need open themselves to others, they, presumably, can be locked tight. But collecting information from third parties, winning their cooperation, and coordinating everyone's activities with theirs may sometimes require that at least some military systems become transparent and accessible to those third parties. Barriers will then be needed between those systems and the more closed components of the GIG.

At the second level of protection, the question is one of compartmentation. What topics should be walled off or compartmented? What is the cost of inhibiting the cross-fertilization of ideas? Who decides these questions? Convenience and custom let data owners

choose, but is that always best? Should certain types of information have a shelf life, and, if so, by what means can the shelf life be imposed without impeding the interoperability of fleeting data with more conventionally permanent information?

Coerced insecurity creates a third set of questions. What if field operatives with a live tap into the GIG find themselves on the wrong end of an AK-47? Soldiers may be instructed to disable their units prior to capture, but they may not always be able to. Adversaries, unimpressed by the Geneva Convention, can credibly threaten people who shut down their access after capture. Do architects therefore limit what such terminals can access? If so, it is then difficult to see how to develop new warfighting techniques, such as swarming, that require warfighters to be fully and minutely informed about the location, activities, and plans of their colleagues—a prerequisite to tight and time-sensitive coordination of fires and maneuver. Another approach may be to make the GIG's servers sensitive to and thus unwilling to answer people who display the biological indicators of stress that being captured might induce. But might not the stress of combat alone then induce a cutoff of information? And what about giving users alternative passwords (as has been mooted for ATMs)? Will they remember to use them without hesitation? What kind of alternative information can be served up that is realistic enough not to anger the captors and put the captured at risk but that also does not reveal compromising details? How can adversaries be shown deceptive data without our own forces being fooled?

Trust is the fourth question. Should people be vouched for centrally or through intermediaries who presumably know for whom they vouch? Should access to the system by its components be predicated on the diligence with which they defend computers? If so, how can such diligence be effectively measured? Should architecture be transparent so that users can understand it, or opaque so that foes cannot?

Interoperability

Interoperability is needed on several levels. At the bottom is the ability to pass bits physically. Next is syntax, the ability to format bits into discrete manipulable chunks (e.g., TCP/IP, HMTL, JPEG). Its faults notwithstanding, DoD's Joint Technical Architecture has

pointed the way for interoperability at the physical and syntactic levels. The invention and widespread acceptance of XML as the markup language has resolved the syntax problem of and thus set the stage for the third level: domain semantic standards—the ability to refer to the things and concepts of the real world in mutually understood ways. DoD has endorsed a clutch of semantic standards for e-commerce, message traffic, and mapping symbology, but many domains have no standards, and others have far too many that conflict with each other. Numerous databases use similar terms in incompatible ways, a problem informed but hardly solved by the ongoing conversion from database schemas to XML schemas. The fourth level deals with process semantics: standard ways to express intentions, plans, options, and preferences so as to support negotiations.

The quest for semantic interoperability has architectural ramifications. Behind every semantic structure is an implicit data model of the world, one that makes some things easier to express and other things harder. How the structures of *human* language influence thought is a long-standing issue in linguistics, with some seeing great influence.[10] Whether similar semantic structures exist for digital data is even less clear. Computers are less flexible than people and are thus more in thrall to their limited vocabulary. But to what extent will people take their cognitive cues from their machines' semantics?

The standards *process* within an organization also reflects the organization's distribution of power. Standards can help consolidate the current GIG, which is composed of stovepipe systems managed by disparate owners. Open systems can also help DoD share information with coalition partners and their systems. Can stocking the GIG with interesting information and useful applications that use standard semantics induce others to pick up such standards for their own interactions? Or, alternatively, will various semantic constructs bubble up, with interoperability generated through explicit negotiation over what terms mean or the implicit use of inference?[11] Should

---

[10] An example is the appendix to George Orwell's *1984*. Also see Steven Pinker, *The Language Instinct,* Harper-Collins, New York, 1995; or Noam Chomsky, *Aspects of the Theory of Syntax,* MIT Press, Cambridge, MA, 1965.

[11] See Tim Berners-Lee, *Weaving the Web,* especially Chapter 13, "Machines and the Web," Harper-Collins, New York, 1999.

there be a common framework for referring to real-world objects, or should there be multiple frameworks, each called out by name?

Interoperability issues arise in great force when coalitions are being formed with partners that can build their own global information systems. How should the architecture of something akin to a NATO integrated system be developed? There are no easy choices. Having the United States and Europe build separate systems and merge them as they mature ensures that end-of-cycle integration will be long and hard. The United States could supply the global components (such as long-lines, satellites, and long-endurance UAVs) and make these the background against which local (and thus often coalition-supplied) components sit. Or the United States could supply an architecture in sufficient detail to make European systems plug-and-play. Both are technically feasible approaches, but the Europeans are unlikely to be happy playing second fiddle. Having the United States and Europe build a NATO architecture together from scratch could delay the U.S. architecture by five to 10 years. Standard, or at least explicit, interfaces in planning the U.S. GIG cannot hurt—regardless of which of these four choices is made.

Can access to DoD's GIG be used to induce third world nations to become allies? Although technical issues may be less daunting in linking systems in this case—DoD will give, third world allies will take—trust issues are more salient. Today's friends may not be friends tomorrow, so the case for giving them deep access is hard to make. What are the terms of the relationship? What will such friends be shown, with what restrictions, and in exchange for what? Should the United States applaud or even seek to induce changes in these partners' information systems (e.g., their forgoing the purchase of potentially competitive C4ISR systems or their adopting an architecture and thus viewpoint similar to DoD's)?

Indeed, how deeply should one nation's military information system penetrate another's? Take Walmart's relationship to Proctor and Gamble, one of its major suppliers. Walmart could have chosen to collect data on its own sales, in-store, and warehouse inventories and then calculate the restocking schedules and hand the resulting orders over to P&G. Instead, it makes intermediate data directly available to P&G and leaves P&G responsible for scheduling production so as to optimize Walmart's inventory of P&G products. In deep

partnering relationships, all partners can write to the files of the others—much as several doctors may enter information on a patient who is under the primary care of a specific physician—and can trust imported information enough to automatically fuse it with their own. Partners may contribute to a library of codes and subroutines to be liberally passed around. A further step could be for partners to populate a jointly accessed architecture with applets, servlets, agents, or knowledge rules, each capable of successively transforming raw information. When systems seek mutual intimacy, their success may depend on having a transparent architecture.[12] But all these are explicit policy choices.

Integration

Systems must work together. But designing an architecture to make them functional on day one at the expense of everything else is a recipe for growing obsolescence on day two and beyond. How and when should capabilities be added? Is it better to have everything up on day one, or is it better to introduce capabilities (or at least major components thereof) serially, shaking out the bugs one stage at a time? The second of these, incrementalism, has two problems. First, a tool such as a grammar checker must reach some threshold of capability before it is worth using at all. Switching has costs, and users forced to switch too often will simply balk. Second, incrementalism can lead to a hodgepodge. Legacy features (such as DOS and the four-byte address space of IP, not to mention two-digit years) may be hard to eradicate. Conversely, projects that aim to have everything up all at once force users to wait a long time. Only a third of all complex systems arrive on time, and failure can be devastating. And even success can be accompanied by great weariness and hence a wariness of change, or can mean a structure so finely tuned to today's problems that it resists conversion to tomorrow's. All-at-once complex systems are increasingly ill adapted to a world that is always moving on.

---

[12]For example: to provide targeting guidance to a topside gun on an Italian frigate, data from a British UAV's electro-optical sensor is linked through a U.S. network to readings from Dutch microphones, the data flows being fused with the help of a French-hosted software agent and being compared to a German-provided database of marine templates.

Another key issue is whether users feel the systems are theirs, in the sense of being a tool they understand how to use and apply to their problems. Tools, as often remarked, should feel like a natural extension of the body or mind. Good carpenters or musicians take proprietary and personal interest in the tools they use, up to the point, sometimes, of not allowing others to touch them. Tools that are balky, arbitrary, fussy, and unreliable are less likely to be used and, if forced on people, less likely to be mastered or maintained. Will the GIG be perceived by users as friendly and trustworthy enough to entrust their lives and time to? Will they come to "own" their piece of it, seeing it as something they have control over and therefore will step up to feeding and caring for?

THE NEED TO THINK NOW

Information systems reflect the organizations that buy, run, and embed them in an architecture. The architectures reified in computer and communications devices tend to replicate, with some mix of conscious design and unexamined assumptions, relationships that exist among people. But architectures also have a funny way of molding these relationships. Information systems are a tool of amplification and disenfranchisement and inevitably alter the balance of power in any organization they enter.

In their dawning, computers were strange creatures, each seemingly programmed with its own language and conventions. As computers became standardized in the 1960s around the IBM 360 mainframe, the popular expectation was that they would become instruments of top-down control. Data would be entered at the bottom, collated, organized, drawn upwards, and used to give top-level managers a much finer picture of their enterprise. The autonomy of middle managers, which was based on their possessing knowledge too variegated to be effectively amalgamated, would thus end, as would the workers' freedom of action, which was based on their being able to make their own decisions. Computerization into the 1980s did, indeed, encourage the amalgamation of disparate enterprises into conglomerates managed "by the numbers" alone.

However, when the proliferation of personal computers in the workplace that began some 15 years ago combined with the privatization and rapid spread of the Internet that began in 1992, a revolu-

tion from the bottom was sparked. Once again, there is an expansion of approaches and a sense of great change, with considerable disagreement about what these all mean for human contact. Will this prove permanent or temporary? Again, the disenfranchisement of middle management appears imminent. In all this flux, it seems right to empower everyone to search for the new best way or at least the best way for him or her.

It is too early to tell whether this empowerment represents a revolution from which there is no going back or simply the blooming of a thousand flowers before a new reconsolidation. But it is not too early to think clearly about the situation. DoD has had, does have, and will continue to have an information architecture, regardless of whether it knows or admits as much. This architecture is driven in large measure by policy choices; it is not solely a consequence of information technology. These policy choices should be made explicit and maybe open in most cases. Otherwise, DoD may either unconsciously make choices it does not like or subconsciously opt for immediate efficiency over longer-term adaptability.