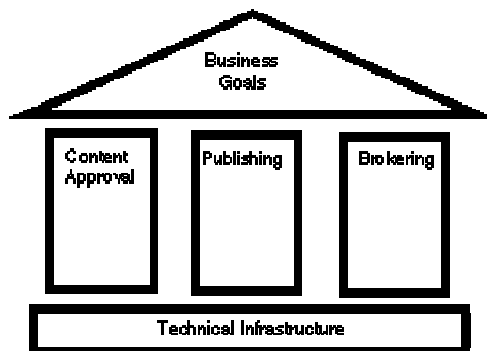




IntraNet Methodology™

Concepts and Rationale



Steven L. Telleen, Ph.D.
Director, IntraNet Solutions
Amdahl Corporation

stevet@iorg.com

Copyright (c) 1995 All Rights Reserved

IntraNet is a trademark and Amdahl is a registered trademark of Amdahl Corporation.

Contents

- [Introduction](#)
 - [History and terms of the Internet and World Wide Web](#)
 - [Universal Client concept](#)
- [Focus On Internal Effectiveness](#)
 - [The Info-SuperHighway frenzy](#)
 - [Client-Server Information management](#)
 - [Why management frameworks are important](#)
- [Benefits Of Intra-Enterprise Use](#)
 - [Reducing Information Overload](#)
 - [Empowering the individual](#)
 - [Efficient Transfer of Information](#)
- [The IntraNet\(tm\) Information Framework](#)
 - [Users/Authors/Brokers/Publishers](#)
 - [Framework Overview](#)
 - [Deliverables](#)
 - [Rapid framework implementation](#)

Introduction

History and terms of the Internet and World Wide Web

The Internet began in the 1960s as a project of the U.S. Department of Defense. The growing importance of computers gave rise to multiple challenges both in sharing information among diverse sites and networks and in keeping the information flow intact during potential disruptions at individual sites.

The Internet is based on a set of protocols developed to allow these distributed networks to route and pass information to each other independently, so if one site is down, the information can be routed through alternate sites to its final destination. The protocol developed for this purpose was called the Internetworking Protocol or "IP" for short. When you see the acronym TCP/IP, it is this same Internetworking Protocol that is being referenced in the second part of the acronym.

The IP protocol came into widespread use in the military community as a way for researchers to share their computerized information. Since the military had numerous research projects underway at universities around the country, and it provided an effective way to move information across diverse networks, the protocol quickly spread outside the defense community. It also spread into NATO research institutions and European universities. Today the IP protocol, and hence the Internet, is a ubiquitous standard World Wide.

By the late 1980s a new problem had emerged on the Internet. In the early days information consisted primarily of mail or traditional computer data files. Various mail protocols and file transfer protocols had evolved to handle these requirements. However, newer types of files were beginning to emerge on the Internet. These were multimedia files, that contained not only pictures or sound, but also hyperlinks that allowed users to jump around inside files in a non-linear way, or even jump to other files containing related information.

In 1989, the European Particle Physics Laboratory, CERN, began a very successful internal project that led to an effort to create standards for passing this new kind of information around the Internet. The basic components consisted of a standard for creating the multimedia, hypertext files, and a standard for serving these standardized files when they were requested. The file standard is called the HyperText Markup Language or HTML. It is actually a much simplified subset of another standard called the Standard General Markup Language or SGML. The server standard is called the HyperText Transfer Protocol or HTTP. A server running the HTTP demon will send HTML files over the Internet to clients requesting them.

These two standards provided the basis for a whole new kind of access to computerized information. Creating multimedia files in a standard way allows client software to be built that not only can retrieve the files from an HTTP server, but also open them and

display them as part of the request. And, since the file can contain hyperlinks to other files (even when they reside on other computers), a user now has the ability to navigate information with a point-and-click interface from what looks to them like standard printed documents. This technology takes away the complexity of accessing information on distributed computers.

The multimedia files served and retrieved in this way are commonly called "pages" by those using this technology. The information sent to the client machine for each request is known as a page. This is because the information that normally would be a document on paper often is broken into smaller units that are hyperlinked together. This separation gives the user flexibility in deciding what she wants to see, and saves time and network bandwidth by not transferring a lot of information that is not of interest.

Most browsers can be set to automatically retrieve a specific page when they are started. Usually this page contains an index hyperlinked to other pages the user retrieves frequently. This page is called the home page, and if the browser has a button called "home page" this is the page it will return to when the button is activated. The hyperlink documents mentioned in the previous paragraph generally have a page that is the equivalent of a table of contents, or at least provides a common starting point for the set of information. This page is called a home page too. In this context, the term home page is relative and can be used for any index page that provides a starting point into a specific set of information. These relative home pages generally are preceded by the name of the set of information they are representing, e.g. the Amdahl home page.

Like the Internet, on which all this operates, a specific page might be found and accessed via several different routes. This global collection of pages and hyperlinks on the Internet is known as the World Wide Web (or WWW or W3). It is completely distributed, and an author (or publisher) of generally available information has no way of knowing about all the index pages that may have pointers linking people back to her page. It is important to note here that the server of that page can be set up to keep track of who accesses that page, just not what other pages are pointing to it. For those of us used to paper media this is almost a reversal of the kind of information we had available. In many research fields in the paper world there are periodical indexes that list all the articles referencing a specific article, but there is no way to tell when someone reads, that article. With this medium we can tell who is accessing (reading) the page, but not who is referencing it.

The other side to this is that it has become impossible to keep track of all the information available on the Web. New information appears and disappears regularly, independent of any central management. However, before we become too alarmed by this fact, we should consider that this is no different than the world of paper and print. If we find this problematic, then it probably comes from restricting our view to the sheltered world of computers where, in the past, information was so hard to get in and out that we had a clear gate and a limited set of information to deal with. This technology widens that gate and makes electronic information as easy to create and as easy to access as print information.

Universal Client concept

The client software that retrieves and displays the HTML files is called a browser. The first of the graphical browsers was Mosaic, from the University of Illinois. Many of the browsers today are based on MOSAIC. However, since the information is standard, it does not matter which browser is used, as long as it supports the standards.

Browsers exist for most major client systems that support intelligent windows. These include MS/Windows, Macintosh, X-Window systems, and OS/2. There also are browsers for non-windowing systems that can display the text portions of the accessed documents. The availability of browsers across these diverse platforms is significant. The operating environments of the author, the server, and the client viewer are independent of one another. Documents built using the HTML and related standards and served by an HTTP server can be accessed and viewed by any client regardless of the operating environment on which they were created or from which they were served.

HTML also supports forms development and return functionality. This means that the user interface extends beyond point and click to both query and data input. A number of sites, including Amdahl, have written interfaces between HTML forms and legacy applications, creating a universal client front-end to those applications. This opens the possibility of writing client-server applications where the author does not have to be concerned with client-side coding. In fact, new applications already are emerging where the client is assumed to be a browser, Oracle's WOW interface which substitutes for Oracle Forms and Oracle Reports is an example.

While this technology is still very young, it appears to be having an analogous effect on information management to what the invention of solid state electronics and microprocessors had on computer hardware. It is modularizing functions and simplifying applications allowing us to move to a new level of integration that is closer to the business functions of the enterprise.

Focus On Internal Effectiveness

The Info-SuperHighway frenzy

Perhaps it is because the Web technology evolved in a cross-organization environment that the primary focus of its application as it moves into the commercial arena has been on inter-enterprise and direct marketing functions. The ease-of-use aspect of the universal client has not been lost on marketeers who believe that this technology finally provides the friendly interface that will open up network computing to the masses. The headlong rush into this arena has created quite a flash, that in retrospect no doubt will provide rich material for researchers who study business and marketing strategies.

The most hyped early application for the Web was as a new channel for selling directly to customers. While a few companies are beginning to realize revenues directly from Web

sales, at this writing, the majority of business uses still revolve around activities that do not involve the direct transfer of money (or credit). Many companies market their products and services on a home page, much like they might advertise in a magazine. Software patches and support information also can be dispensed successfully today.

One of the most successful business uses of the Web today is supporting informational queries that otherwise would have been handled by a customer services representative. These range from checking personal bank account balances and recent transactions via an encrypted transaction ([Wells Fargo Bank](#)), to checking the location and status of specific overnight packages ([FedEx](#)), to finding job opportunities in the [State of Florida](#), to identifying the times and [local theaters](#) where a movie is playing. In a similar vein, a number of politicians have discovered the Web as a new channel for distributing their ideological products and [public service information](#). More elaborate Electronic Data Interchange (EDI) for business to business use has been suggested, but if much has happened in this arena, it has not been widely shared.

But the external Web is only one potential use of the Web technology. Its use for improving communication and coordination inside large organizations has the potential to generate even more profound effects on the way we do things than the external commerce receiving most of the publicity.

Client-Server Information management

Many large enterprises today have the same distributed and heterogeneous information problems inside their organizations that the Internet and World Wide Web were designed to solve on a global basis. Additionally, many have an "Internet" backbone already in operation. And, many of these organizations have no trouble identifying real benefits that would accrue from replacing current paper processes with electronic information flow. Because the networks are internal, access and security standards can be enforced that would be impossible on a public network.

Earlier technologies applied to these corporate problems were centralized in nature and difficult to keep current. More recent groupware technologies have attempted to provide a distributed approach, but users have complained that they are not intuitive to use, and if they want to use an application not already integrated with their proprietary groupware package, the custom integration takes a significant effort by an expert.

The Web technology is different. It is easy to create, publish, and access information; so easy that anyone with access to the TCP/IP backbone can publish information inside the enterprise that is accessible on a wide variety of client platforms. Additionally, many off-the-shelf applications can be quickly and easily added as helper applications. Even interfaces to custom legacy applications are much easier to create than with traditional programming approaches. An analogy that comes to mind is the explosion of desktop systems that occurred after the microprocessor revolution. In this case we are facing distributed information management rather than distributed computing.

Why management frameworks are important

Managing distributed systems provides interesting challenges that are not found in centralized environments. The biggest challenge is moving from an attitude of control to an attitude of enabling independent decisions and actions. Without some standards organizations lose their ability to communicate effectively and coordinate their activities. Without some level of support, knowledge workers become too involved in low level maintenance activities at the expense of the high leverage functions that most benefit the enterprise. The challenge is meeting the needs for coordination and efficiency without destroying the independence of decision making and action that make enterprises strong and flexible.

The Web technology makes the creation and publishing of information easy. It also makes the retrieval and viewing of the information easy. What is not easy is finding the relevant information that is created in this independent environment. What Web technology is missing is a marketing and distribution channel for the information created. Without a marketing and distribution channel there is no efficient way for information to be found. Our paper systems have information marketing and distribution channels that have been refined over time. In most organizations it is so integrated with the way we operate that we don't think of it as an information marketing and distribution channel, it is viewed as the processes and procedures we use to do business. It is the management framework.

Often the Web usage in organizations began as a skunk works project started by technical experts. While today many of these projects enjoy higher visibility than during their inception, most still are managed as a skunk works, separate from the mainstream business. Much of the material on corporate home pages consists of traditional collateral translated to Web standards. The technology has not been incorporated into the business infrastructure, where it has the most potential to transform the way the enterprise manages itself.

An enterprise can treat Web technology as a separate or "extra-step" technology, or it can integrate the technology into its day-to-day operation. Once Web technology is effectively integrated into the internal operation, its effective use for external interactions becomes a natural and easy extension. Without the internal infrastructure, external interactions will always be strained and limited.


Benefits Of Intra-Enterprise Use

Reducing Information Overload

Information overload is a malady of our time. Technologies that were supposed to help the problem seem only to have made it worse. This is not surprising if one looks at the in-baskets (paper or electronic) of the average knowledge worker. Even discounting the inevitable pile of "junk" mail advertisements, the majority of information is sent to that


worker "just in case" he might need it. Add to this the information that is "out of phase", that is will be needed, but not until later, and the majority of the in-basket is accounted for. The knowledge worker probably files half of the just-in-case information (just-in-case) and all of the out-of-phase information. When information is needed, he is faced with a high-volume, low-density personal information system, that may have additional complexities of multiple formats or media.

The advent of xerography machines made just-in-case information much worse. Instead of being limited to a small number of carbon copies, the copy list could be expanded. Electronic mail has carried this to greater heights. Today a "publisher" of information can store personal mailing lists and, with a single command, send a nearly unlimited number of copies "just in case" they might be needed. Some of these information loaders recognize that their lists are inappropriate. But instead of tailoring the list, they put an identifier at the front of the message that goes something like this: "If you are not interested in ... delete this message now." The message still clutters the mailbox, and the knowledge worker still has to spend time evaluating and disposing of it.



Just-in-Case
 Irrelevant (data)
 Out of Phase (timing)
 High-Volume/Low-Density
 Inconsistent formats/interfaces


The alternative to "just-in-case" information is "just-in-time" information, or information on demand. This has been a promise of computers and networks that up to now has failed to meet its potential. Historically there have been two major approaches to just-in-time information delivery. The first left information distributed across applications and systems. To access the information the user had to learn and navigate multiple, complex, access procedures. Once accessed each application required a different interface. Faced with this level of complexity, users generally ignored just-in-time information. They may have learned how to access one or two applications, but the rest were left to languish.



Multi-Source Infor. on Demand
 Complex access procedures
 Different access procedures
 Different user interfaces


To solve this problem, some enterprises attempted to collect all the distributed information into one master system. This gave the user a single access and single interface. However, because they tried to manage all requirements in the enterprise centrally, these systems tended to become large and complex. After more than a decade many still are not fully populated with information because the cost of inputting and maintaining it was too great. There were other problems as well. The complexity of these unified systems made them difficult to modify and to use. The tools to manage them were designed to support the discrete data of transactional processes. Over the past decade we have moved much more toward complex data to support informational processes. The shift in information needs combined with the difficulty of change has caused these large centrally managed systems to lag the enterprise requirements.

The Web technology offers a new approach for delivering information on demand. Because it supports distributed information authoring, publishing and management, it does not require the complexity of the old centralized systems. Information is authored and managed by those



Single Source Infor. on Demand
 Size makes it difficult to manage
 Complexity makes it difficult to modify
 Forces a transactional model
 Users find it time-hibie & difficult to use
 Expense to maintain

who create it, without having to rely on programmers to create data entry and reporting programs. With the new browsers, a user can retrieve and view information from distributed sources and systems using a simple, uniform, interface without having to know anything about the servers they are accessing. These simple changes in the technology will revolutionize our information infrastructures, and change our organizations.



The Universal Client

Platform independent

Common interface to information & applications

Point & click access to heterogeneous information sources

Supports independent, heterogeneous information management

Allows deployment of client server without client side coding & testing

Empowering the individual

The key characteristic of this technology is its ability to shift control of information flow from the information creators to the information users. If the user has the ability to easily retrieve and view the information when they need it, the information no longer needs to be sent to them just-in-case. Publishing can be separated from automatic distribution. This applies to forms, reports, standards, meeting minutes, sales support tools, training materials, schedules, and a host of other documents that flood our in-baskets on a regular basis.

Making this work requires not only a new information infrastructure, as discussed above, but a shift in attitude and culture. As creators of information we must retrain ourselves to publish without distributing. As users we must retrain ourselves to take more responsibility for determining and tracking our changing information needs, and actively and efficiently acquiring the information when we need it.

More detail on the roles required to support these changes will be discussed in the section on the IntraNet™ Information Framework

Efficient Transfer of Information

The sections above have discussed the inefficiencies of distributing information through "documents" just in case it might be needed. Another way we transfer information is through education and training. Traditional face-to-face training tends to have a high content of just-in-case and out-of-phase information. The content becomes like the filing cabinet where neither the student or the instructor know which pieces will actually be needed on the job, so it's all crammed in: just-in-case.

If information can be efficiently found and assimilated on demand, it changes the education and training model. Traditional training can focus on skills, while specific content can be provided on demand, as it is required.

The on-demand model has other advantages. Traditional information transfer technologies are batch oriented. A batch of information is created for delivery through documents or traditional education and training methodologies. Because the information is packaged as a batch, when one portion changes it often is not updated immediately. As time progresses the batch gets out of date. Some material becomes obsolete and new material is not yet included. To fix the problem, the entire batch would have to be re-released every time a change occurs, and the released batches all have to be updated. Several questions begin to emerge. How often can the enterprise afford to update the batches? How much time can the employees afford to spend in training? And, how much information can an employee absorb at one time?

As mentioned above, the Web technology supports distributed information creation and management, so each portion can be updated as it changes. If the proper infrastructure is in place, current information content can be found and accessed when and where it is needed. This reduces both the time needed for training and the amount of information an employee is required to absorb at one time.

The IntraNet™ Information Framework

Users/Authors/Brokers/Publishers

Key to understanding the IntraNet™ information framework is the definition of several roles. These truly are roles rather than positions, because an individual is very likely to play more than one in the course of doing their work. Understanding the differences helps the individual and the organization navigate the requirements for success.

Users access and view the information. There are many reasons why a user may be accessing the information, and the reasons will vary even from session to session with the same individual. The important point is that the user is where the value is created and the ultimate requirements are defined. If the difficulty of accessing the information exceeds the information's value to them, they will either not use it or will find another way to get it.

Authors create the information. In traditional media, authoring and structuring are intimately tied. The access paths are primarily linear, and all the related information is bound together and delivered together. If a different access path or a different combination of information is required, it frequently is more efficient to build the structure into a new document, replicating the common information, than to use external indexes to tie together pieces of different documents.

In the new world of hyperlinks, the structuring aspect of authoring will change. Several factors are driving this change. First, hyperlinks allow the user to pick and choose the order in which they want to access information. So while the author will necessarily continue to provide structure, the function of that structure will shift to helping users determine for their current need which information is valuable next rather than attempting to determine their needs *apriori* for them.

Second, when information is referenced or recombined, the original can be hyperlinked rather than copied. This concept will be covered in more detail in the Broker discussion below. However, it does have an impact on authoring. Much like the database world, the goal will become to reuse information rather than maintain separate copies that have to be continually synchronized. This means authors will become more concerned with structuring information into reusable modules than with creating specific linear structures.

Brokers are the key to finding information. This is true in the world of paper and will remain true as Web technology is adopted. Any technology that allows both prolific, independent creation of information and easy access to it will quickly generate inefficiencies for users trying to find information for specific needs. The information broker provides this service. Information brokers supporting our traditional information infrastructures are ubiquitous. We are using information brokers when we go to a phone book, to the TV Guide, to a librarian, or to an edited anthology. In our businesses we have people whose job it is to broker competitive information, to broker benefits and Human Resources information, and to broker product information. Many of the knowledge worker jobs are in fact information brokering jobs.


It is the brokers who will be the most affected by this technology. They can use it to search for and screen information for their constituencies. They can use it to deliver their results. But the most dramatic impact is the switch from delivering repackaged content to delivering information access pathways. This is an important distinction, because it changes both the deliverable and the focus of the information broker.

The deliverable in the past might have been a 300 page bound document. With the Web technology it might be a single page of hyperlinks. In creating the 300 page document, the broker spent a large portion of the time "cutting and pasting" and editing together information from other documents. With the Web technology this time will shift toward understanding the users' decision processes and structuring access paths through the content to better support those decisions. The Web technology replaces the time consuming effort of collecting and republishing text with single lines hyperlinked to the original text.

Publishers make information available. Organizations require information content to be managed, coordinated, and communicated in predictable and relatively efficient ways. Certain kinds of information are required for directed activity. Information becomes old and out of date. Some information requires approvals and validation. This is true of electronic files just as it is of their paper counterparts. A major difference is that we have 400 years of infrastructure built up to manage information conveyed on the paper medium. In fact, many of our bureaucracies and organizational processes are nothing

more than infrastructures for moving information on paper, and many managers of white collar workers are in effect publishers of official information required by the enterprise.

Publishers have different information access requirements from users or brokers. They need a structure to help them efficiently manage the information content for which they are responsible. For this reason the IntraNet information framework supports two distinctly different kinds of pages and page structures. One supports users and brokers, the other publishers and managers. We call the second structure a Management Map, and the pages it contains map pages. Map pages are structured according to rules, and once created they do not require significant maintenance except when management changes occur. Even during these times the changes are straight forward, restricted to the areas changed, and not time consuming activities.



Framework Overview

The diagram here provides an overview of the key components of the IntraNet™ information framework. The roof depicts the business goals, which provide the reason for creating the framework. The technical infrastructure at the base provides the necessary foundation for the development, management, and access of the information. It consists of the hardware,

software, networks, protocols and standards required to implement the remaining structure. The middle three columns provide the support for the business goals. They include both the publishing and the brokering functions mentioned above and the additional function of the content approval processes that support the publishing of "official" information. Content approval processes are important for coordination and efficiency and to protect the enterprise and its employees from liability and loss.

Deliverables

The IntraNet Methodology leads to the establishment of the framework described above within a client enterprise. The enterprise receives four basic deliverables:

1. an established framework and implementation methodology,
2. the IntraNet Methodology Information Base, a hyperlinked, comprehensive set of text, multimedia files, and templates supporting the methodology,
3. targeted orientation and workshops for key executives, administrators, publishers, and authors, and
4. implementation support and guidance from experienced personnel.

Rapid framework implementation

One of the major benefits of adopting the IntraNet Methodology is the rapid implementation of a framework designed specifically to manage information using Web technology. The basic management framework can be implemented quickly using our

supplied templates, prototypes, and design models. It employs a process that encourages rapid evolution toward self-customization in the client's enterprise. The Web technology is used to deliver a wealth of information immediately that supports the framework and captures and reinforces the customization results.

The implementation methodology used by the Amdahl team amplifies the effects and quickly engages the client organization's personnel in the management of the new framework. This process both shortens the implementation cycle and encourages independent innovation so the benefits begin accruing sooner.

Author: Steven L. Telleen, Ph.D.

Initial Version: January 1995

Last Updated: May 1996

Comments about this material may be sent to the author at stevet@iorg.com

**Copyright © 1995-1996 Amdahl Corporation, Sunnyvale, California, USA.
All Rights Reserved.**

IntraNet Methodology, IntraNet Architecture and IntraNet InfoBase are trademarks of the Amdahl Corporation