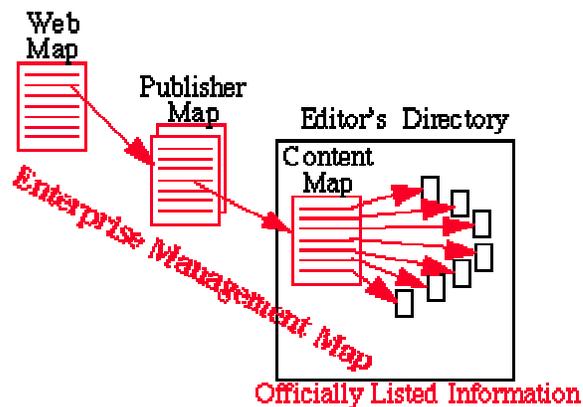




The IntraNet Architecture™:

Managing information in the new paradigm



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

stevet@iorg.com

Copyright (c) 1995 All Rights Reserved

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

Table of Contents

1. [Introduction](#)
2. [Preliminary Concepts](#)
 - o [Intranets](#)
 - o [Three Sources of Information](#)
 - o [Two Types of Pages](#)
3. [The IntraNet Infrastructure™](#)
 - o [Management Roles](#)
 - o [The Enterprise Map](#)
 - o [Diagram of the Enterprise Map](#)
 - o [Shadow Maps](#)
 - o [General Brokering](#)
 - o [Brokering Summary](#)
 - o [Workflow Management](#)
 - o [Access to Database Information](#)
4. [Conclusion](#)

Introduction

Information is the way organizations coordinate their activities and achieve their goals. The technologies for managing and distributing information have changed over time, but the functions required for human organization have remained fairly constant. The effort to find ways to authenticate electronic requests and submissions is merely attempting to meet the same needs that seals, signature comparison, and notary publics met in the paper world. The need to secure information on networks is the same need that led to sealing wax and armed guards in previous eras using paper media.

However, security requirements are not the only functions that stimulate organizations to manage their information. If the information does not enable some further use that provides value, there is no need to secure it. Organizational information generally carries content that enables action leading to a gain or loss of resources. An organization amplifies its ability to control those resources by dividing among multiple individuals the work required to reach a goal. For the organization to be effective, activities and progress must be coordinated. An important reason for sharing information within organizations is the agreement on and coordination of these goals and tasks.

A requirement for successful coordination is consistency of information. It is not very efficient if the existence or location of important information remains unknown to an individual who needs it. It also is not very efficient if a team tries to reach a consensus when each member is operating from a different information base that may be incompatible or inconsistent with the others on the team. Some information gets stale and requires attention to keep it current. Most of today's organizational structures and processes have been refined over centuries to solve these problems for paper-based information.

Information currency and integrity is a much simpler problem when the content does not change often, activities being coordinated are not large or complex, and the information is centrally collected and distributed. However, these are not common characteristics of most enterprises today. The distributed environments more commonly found today need to be able to coordinate information in a different way, and this requires a different set of management structures and processes than most organizations have inherited.

An Intranet offers new options for more effective coordination of organizational activities in a distributed decision-making environment. However, today's paper-based infrastructure inhibits exploitation of many Intranet enhancements and does not effectively control others. What is needed is an infrastructure that uses the Intranet to meet the requirements for organizational coordination by supporting the management of content in the distributed decision-making environment.

Constructing an effective Intranet infrastructure requires attention to three distinct areas: management, technical and content. Management consists of the roles, policies, processes and organization needed to manage the life cycle of formal Intranet content. The technical infrastructure consists of the networks, hardware and software required to

support content development, publishing and access. And, content requires processes to be developed to support special needs such as initial conversions, creation of database or application interfaces and development of "glossy" pages for high impact.

Most of the attention to Intranets has been focused on implementing the technical infrastructure and development of special content. The former is an extension of systems and network administration and the latter an extension of traditional technical publications, PR and programming processes and resources. The unique aspects of managing the creation and maintenance of content generated by [Adaptive Innovation](#) in an Intranet environment to a large extent has been ignored. This paper focuses on these management aspects of an Intranet and presents an organization and information architecture to support the new environment.

Preliminary Concepts

Before discussing the IntraNet Architecture™ a few background concepts need to be introduced.

Intranets

An Intranet is a communication infrastructure. It is based on the communication standards of the Internet and the content standards of the World-Wide Web. Therefore, the tools used to create an Intranet are identical to those used for Internet and Web applications. The distinguishing feature of an Intranet is that access to information published on the Intranet is restricted to clients in the Intranet group. Historically this has been accomplished through the use of LANs protected by Firewalls.

Three Sources of Information

At least three sources of content quickly emerge on enterprise Intranets: formal, project/group, and informal.

The **formal information** is the officially sanctioned and commissioned information of the enterprise. It usually has been reviewed for accuracy, currency, confidentiality, liability and commitment. This is the information with which the formal management infrastructure is most concerned.

Project/group information is intended for use within a specific group. It may be used to communicate and share ideas, coordinate activities or manage the development and approval of content that eventually will become formal. Project/Group information generally is not listed in the enterprise-wide directories and may be protected by passwords or other restrictions if general access might create problems.

Informal information begins to appear on the Intranet when authors and users discover how easy it is to publish within the existing infrastructure. Informal information is not necessarily the same thing as personal home pages. A personal folder or directory on an Intranet server can serve as a repository for white papers, notes and concepts that may be shared with others in the enterprise to further common interests, for the solicitation of

comments or for some other reason. Instead of making copies, the URL can be given to the interested parties, and the latest version can be read and tracked as it changes. This type of informal information can become a powerful stimulus for the collaborative development of new concepts and ideas.

Two Types of Pages

There are two basic types of pages: content pages and broker pages. Content pages contain the information of value required by a user. Broker pages help users find the content pages appropriate for their current requirements.

Content pages can take many forms. They may be static pages, like the ones you are reading here, or they may be active pages where the page content is generated "on the fly" from a database or other repository of information. Content pages generally are owned by an individual. Over time expect the "form and sense" of content pages to change as more experience is gained in the areas of non-linear documents (hyperlinking), multimedia, modular content and integration of content and logic using applets.

Broker pages also come in more than one form, but all have the same function, to help users find relevant information. Good broker pages serve an explicitly defined audience or function. Many of the pages with which we already are familiar are broker pages. A hyperlink broker page contains links to other pages, in context. It also may have a short description of the content to which it is pointing to help the user evaluate the possibilities. On the other hand, a search oriented broker page is not restricted to the author's scope, but it also does not provide the same level of context to help the user formulate the appropriate question.

Combination search and hyperlink broker pages are common today. Search engines return the "hits" as a hyperlink broker page with weightings and first lines for context, and hyperlink broker pages sometimes end in a specific category that is refined by searching that defined space. It is unlikely that hyperlink broker pages ever will be generated entirely by search engines and agents, because the context that an expert broker provides often contains subjective or expert value in its own right. After all, not all content is of equal quality or value for specific purposes, and even context sensitive word searches cannot provide these qualitative assessments. As the amount of raw content increases, we will continue to need reviewers to screen which competing content is most useful, or the official source, for workers in our enterprise.

A special use of broker pages is for assisting with the management of web content. There are several specific instances of these management pages. We call one instance the "Enterprise Map" because collectively these broker pages form a hyperlinked map of all the formal content in the organization. Other sets are used for project management, functional management and to support content review cycles. The use of broker pages for each of these management functions is discussed in more detail in the next section.

The IntraNet Infrastructure™

Management Roles

The IntraNet Infrastructure relies on four distinct roles for managing the formal content: the Web Administrator, publishers, editors and authors.

The **Web Administrator** is responsible for facilitating cooperative opportunities among the various organizations in the enterprise and administering the enterprise content management infrastructure. By contrast, the Webmaster is responsible for the technical infrastructure. The same person may serve in both roles, but to do so requires that she have both of the distinctly different skill sets and enough time to carry out both sets of responsibilities. The Web Administrator chairs the Enterprise Web Council.

Publishers determine what kinds of formal information will be created and maintained by their organization. They represent their organization on the Enterprise Web Council and may create and chair an Editorial Board within their own organization. The publishers own the processes and policies that both the enterprise and their organization require officially sanctioned information to follow. In larger organizations, they may delegate the monitoring and implementation to editors, but the responsibility remains with the publisher.

Editors are found in organizations that have multiple product lines or service areas. For example, Human Resources might have editors for Benefits, Compensation, Equal Opportunity and Staffing. In a line of business, the editor often is the primary marketing person for each product line. The editor determines what official information will be created for specific activities and manages the information creation and update process, including the formal review cycles.

Authors create the content.

The Enterprise Map

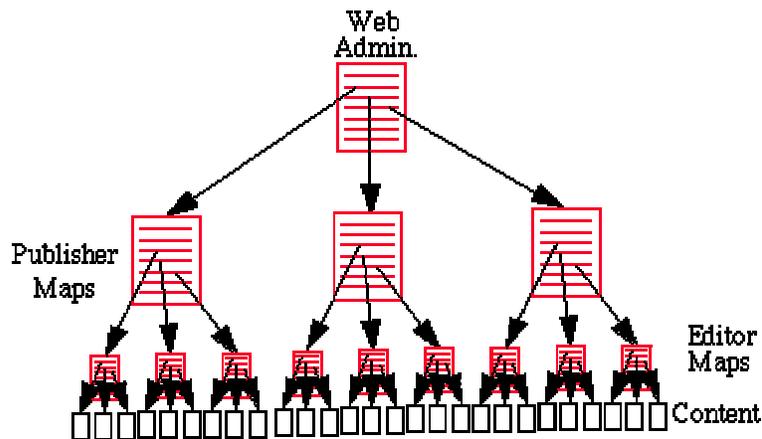
A structured set of broker pages can be very useful for managing the life cycle of published content. We call this the Enterprise Map, and while the primary audience for this set of broker pages is management, we have discovered that end users frequently use the Enterprise Map for browsing or to find content when their broker pages have failed them.

With the exception of the content pages at the bottom of the map, the Enterprise Map pages consist only of links. Each page corresponds to an organization committed to the creation and quality of a set of content pages. In today's organizations, commitments tend to aggregate into a hierarchical pyramid, but the mapping technique also could be applied to most any organizational model. The Enterprise Map also does not have to be based on organization. It could be a logical map where the top level is the mission, the next level the major focuses required to accomplish the mission, and so on, down to the content level. Since most large organizations are starting from a pyramidal accountability structure, that is the form of the example that follows.

Using the terminology above, the Enterprise Map begins with a top page, owned by the Web Administrator (representing the CIO and /or CEO). This page consists of a link to the Map Page of each line of business and major support organization in the enterprise. This page is owned by the publisher for that organization. The Publisher Pages, in turn, consist of links to each of their Editor's Pages. The Editor's Pages may have additional pages or structure below them created and maintained by the editor that help organize the content, but ultimately these pages point to the actual content pages.

This model can scale to governments or large diversified companies. In a government organization, the Administrator's Page would point to all the Agencies, and the map would follow each agency structure to the content level. Since each agency may be a large organization, each may have its own Administrator and Web Council. A major advantage of this mapping architecture is its flexibility. It can originate from the top down or the bottom up. If several government agencies developed their Intranets independently, with this type of Enterprise Mapping structure, they can be linked together at any time in the future by creating the next level map page. None of the existing Maps need to be changed. This flexibility is a result of the distributed decision making central coordination model on which the architecture is built.

The Enterprise Map



The Map provides a commitment (or accountability) view of all the formal content in the enterprise. Management can start at their point in the map and follow the links to all the content which supports the functions for which they are responsible.

An Enterprise Map has several interesting characteristics. Once it is in place, authors and editors can self publish, and the information automatically shows up in a logical structure. Also, content categories and even editor level functions generally are not affected by reorganizations, because major product lines and service areas generally are not added or deleted. Most reorganizations shift responsibilities at higher levels in the Map. This means that when a reorganization does occur, the Map can be adjusted quickly, by the new managers, by changing one or a few links. Content does not need to be moved around. The result is a very low maintenance path to all the formal enterprise

content, without forcing publishing through a central authority that can quickly become a bottleneck.

Shadow Maps

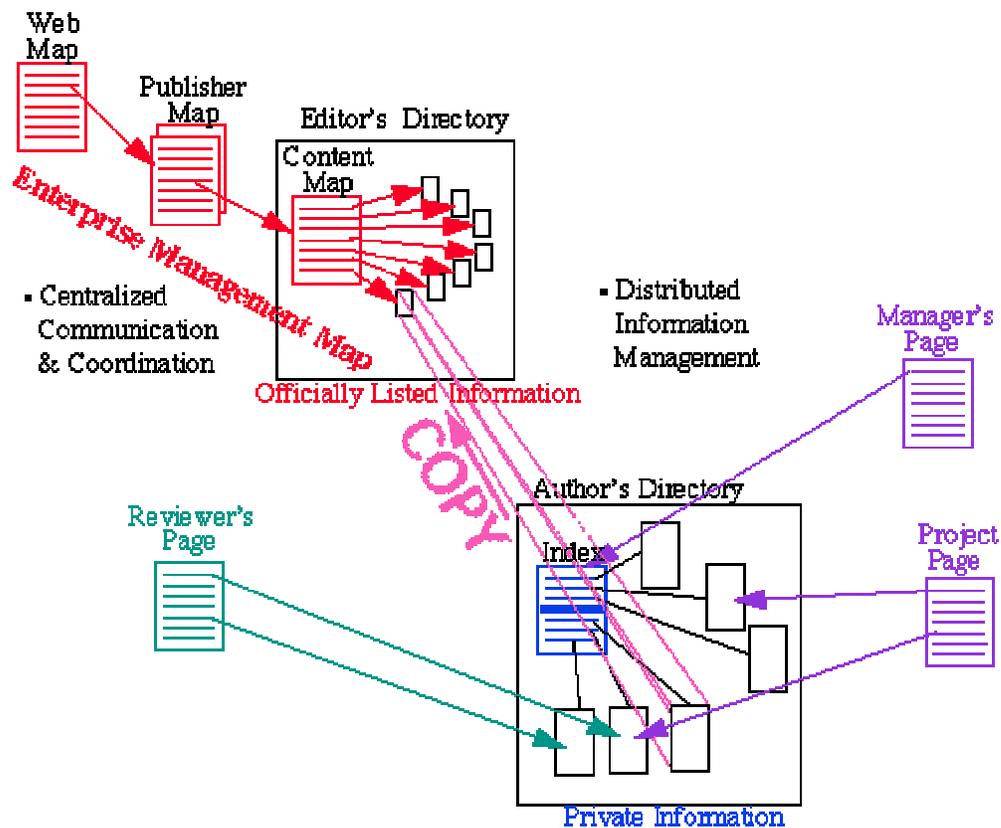
The Enterprise Map provides a management path to all the formally published content. However, management also has a need to see work in progress, formal content that is not yet completed. A Shadow Map can be constructed for this purpose. The Shadow Map works the same way as the Enterprise Map, but it is not generally advertised and can be protected by passwords or other access controls. The Shadow Map can be enhanced with a few additional Broker Pages to assist with the management of content development.

A Shadow Map continues down to the author level. In this model, the author maintains an Index Page that is divided into two sections, work commitments and work completed. When the first draft of committed content is created, the author places it in his web directory and links the item line on his Index Page to the file. As revisions are made, the author places the latest version in the same directory with the same name so the Index automatically points to the latest version. This does not preclude keeping back versions if they are required. The previous version is copied and numbered as it is moved out of the current version status. When the content completes review and goes into "production" the author moves the item from the committed section to the completed section and redirects the link to the permanent address of the published item.

A second Broker Page that can be added is a Project Page. This page is created by the Project Manager and contains item lines for all the project deliverables. When the author creates the first draft, she not only links the content file to her Index Page; she also notifies the Project Manager of the location so the content can be linked to the Project Page. Like the Index Page, as the content is revised the Project Page always points to the most current version, without explicit updating.

In a matrix organization a third Broker Page can be created by the Functional Manager. This page consists of links to the Index Page for each employee reporting to the Functional Manager. This provides a quick path to the work, both in progress and completed, of all her employees. Once again, after the structure is set up, it takes little maintenance with each person keeping his own information up to date.

Finally, Reviewer Pages can be created when the content is ready for review. Each reviewer has a "Review Page," which consists of links to all the content in their review queue. When the Editor (or whoever is responsible for managing the review process) places the content into formal review, it is added to each reviewer's page. The reviewers access their page when they are ready to do reviews, and by selecting a link can retrieve and view the content. There are numerous ways the comments and comment resolution could be handled using Internet technology. One is to funnel comments into a threaded-discussion-group format. Automated email messages can be used to notify or remind reviewers of deadlines and status.



The various Broker Pages discussed above are meant to create a model of the basic management functions and how they can be structured. Whether or not the specific model described here is used, the most effective process for managing Intranet content will use Intranet tools and approaches. When we first conceived of this model, there were no higher level tools to help create and manage the pages for a process like this. Today several tools are emerging to help manage functional sets of pages, and they can be configured to support some of these processes. Some are message-based, others are centralized, shared-database models with Web front-ends. Over time, we anticipate that a variety of vendors will offer improved tools, based on Intranet paradigms, that are specifically tuned to support the distributed, message-based management model.

General Brokering

Brokers are the main way users find information on an Intranet. A broker may serve many functions. He may provide information to users in the context of specific processes, providing structure for efficiency and consistency. He may screen large pools of content for material relevant to a large number of employees so each one does not have to duplicate the process. He may identify which information is considered official. Or, he may provide interpretation of general information in the context of the organization.

Most knowledge worker jobs involve some form of information brokering. In the paper world the broker output often is formally sanctioned by the organization and may be one of the worker's official responsibilities. The same kinds of roles will evolve in the

Intranet world. These types of formally managed broker pages can be treated as content in the map structure described above.

Most organizations also have informal broker pages that spring up. An individual may start the page for herself, and it gains a following, or she may identify an unfilled need and consciously fill it. These pages can be a valuable way to identify and quickly meet new requirements. However, until these pages are in a formal commitment (or accountability) structure, there is no guarantee that the author will keep the content current.

The Broker Directory

An Enterprise Broker Directory, sometimes called a "Yellow-Pages," organized by subject, can help users find the broker that they need. The Broker Directory generally is maintained by either the Web Administrator or the Enterprise Webmaster. Because informal broker pages are included in the Broker Directory, some mechanism needs to be included to keep the directory from filling up with outdated and abandoned pages. As with other Intranet functions, the challenge becomes providing the centralized view without imposing a central implementation bottleneck.

One way to handle this is to create a "Sunset" provision for all pages not officially managed by an organizational entity. Any broker can list their page by submitting a web form or email to the Web Administrator or an automated script. However, informal pages are only listed for 60 days. If the broker does not renew the request in 60 days, the page is removed from the Broker Directory. This allows informal brokers to "self-publish," and protects the directory from becoming a repository of links to abandoned pages.

The Enterprise Index

The Enterprise Index provides users with another way to find information. This frequently is tied to the Search Engine. As described in a [previous paper](#), the Index and Search Engine should not require pages to be published on a specific system or managed by specific management software. The Index and Search Engine should be fed by an agent (Web Crawler or Spider) that regularly searches the Intranet and catalogs the content. This is consistent with the coordination versus control model and also protects the enterprise from major conversion efforts (proprietary locks) if an alternative product or upgrade is desired in the future. This is yet another way for users to find the content they require.

Brokering Summary

Three distinct discovery paths need to be provided by the Intranet Infrastructure:

- **The Enterprise Map**
- **The Broker Directory**
- **The Index and Search Engine**

Workflow Management

Workflow management is a relatively new focus for the Intranet. Historically, a number of Internet/Web tools have been available to help with this process. Email, threaded-mail discussion groups and newsgroups provide forums for discussion and resolution of issues. The HTML "mailto:" function has been used to provide reviewers with easy connections through their browser to these forums.

What has been missing are packages that integrate the functionality of the independent tools, add routing and tracking, and provide the user with an interface that is easy to configure. This appears to be changing with the appearance of companies like [WebFlow](#), [MKS](#), [Action Technologies](#), [RadNet](#) and [OpenText](#) who now offer web-based products that support groupware, reviewer comments, routing, sign-off, and checkout-checkin functionality.

Access to Database Information

Discrete, structured information still is managed best by a database management system. However, the quest for a universal user interface has led to the requirement for access to existing database information through a web browser. Three models of access can be identified:

- **Automatic tailoring of page content**
- **User specified database requests**
- **User initiated database updates**

From a technical standpoint, there are a number of ways these interfaces can be created. What is important is that access be provided to the content providers (knowledge workers) in a way that supports the distributed decision-making, enabling model rather than the centralized expertise model. This means that relatively naive users need to be able to incorporate database managed data into their pages.

A number of tools are beginning to emerge that move in this direction. One set combines a library of cgi-scripts or Java-scripts residing on the hosting web-server with templates, wizards and "bots" incorporated into WYSIWYG authoring packages (e.g. [Microsoft's FrontPage](#)). The other set, coming from the database side, automatically converts database schemas into hyperlinked web pages that allow users to browse and access the data from their web browser (e.g. [Netscheme](#)).

When applications that merge these two functional approaches begin to appear, very powerful packages will be available to content providers who need to incorporate database information into their pages.

Conclusion

As stated in the beginning of this paper, creating an effective Intranet requires attention to the management infrastructure, the technical infrastructure and the content creation process. The focus of this paper has been on architecting a management infrastructure

that supports content creation, maintenance and use in a distributed decision-making environment. The architecture and models outlined above describe a process rather than specific tools. Since we first described and implemented this architecture, Intranet tools have evolved at an unprecedented rate. Even so, today's tools have not made the need for a management architecture obsolete, because tools provide support, not the purpose or goals that define an organization. The evolution of Intranet tools will continue to make implementation and operation of many aspects of the architecture easier into the foreseeable future.

Intranets are rapidly becoming the primary information infrastructure for enterprises. To effectively utilize this infrastructure, we must become as proficient at managing content and coordinating our actions on our Intranets as we are at managing content and coordinating our actions using paper today. The architecture and models above were put forth to provide the first few steps in this direction.

Author: Steven L. Telleen, Ph.D.

Original Version: June 1996

Last Updated: June 1996

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

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