

The Second Layer

How Contextual Tagging Transforms
File Systems into Knowledge Networks

The Second Layer

How Contextual Tagging Transforms
File Systems into Knowledge Networks

Michiel A.M. Arens



My thanks go to Ruud Duivenvoorden and Johan Goosen.
Without their critical attitude and their incredible dedication to the work
they carried out for this book, it would not have come into being.

Author: Michiel A.M. Arens
ISBN 9789403881348
Second edition
July 2026
© Michiel A.M. Arens
Publisher: Uitgeverij Jupiter
Email: uitgeverij@jupiternederland.nl

CIO Foreword

By Paul Wyman, Chief Information Officer

In my career, I have overseen multiple digital transformation initiatives. Most promised efficiency. Some delivered incremental improvements. Few addressed the structural foundation of how organizations manage knowledge.

This book does.

For decades, enterprise file systems have relied on hierarchical containment as their primary organizing principle. That approach is not flawed — but it is incomplete. It was designed for a world in which information moved slowly, projects were more linear, and organizational boundaries were clearer.

That world no longer exists.

Today's enterprises operate in networks of interdependence. Strategic initiatives cut across departments. Regulatory requirements intersect with operational execution. Regional decisions influence global risk exposure. Yet our digital infrastructure still assumes that information belongs in one place.

Tagging introduces a contextual layer that complements hierarchy. It does not dismantle structure; it enriches it. By embedding multi-dimensional attributes directly into documents, organizations gain visibility without restructuring their entire archive.

From a CIO perspective, the value is threefold:

First, tagging reduces structural fragility. Strategy shifts no longer require wholesale reorganization of folder structures.

Second, tagging improves traceability. Audit and compliance readiness become systemic rather than reactive.

Third, tagging lays groundwork for intelligent systems. Artificial intelligence relies on structured metadata. Contextual tagging provides that structure.

Technology alone does not modernize organizations.
Architecture does.

This book reframes tagging not as a feature, but as architectural evolution. It challenges leaders to think beyond containment and toward contextual integration.

That shift is overdue.

Contents

EXECUTIVE SUMMARY	1
Who This Book Is For	5
Core Ideas of the Book	7
PART I: Foundations	9
1 The Purpose of Work.....	9
2 The Invisible Friction Within Organizations.....	13
3 From Documents to Meaning	15
4 The Architecture We Inherited.....	17
5 The Hidden Assumption	21
6 The Question That Changed Everything	25
7 The Second Layer	29
8 Epilogue.....	33
PART II: When structure starts to fail	35
9 The Quiet Chaos of Digital Files.....	37
10 The Illusion of Order.....	41
11 The Hidden Cost of Hierarchy.....	43
12 What Problems Lead Organizations to Implement Tagging? 49	
13 The Breaking Point	53
PART III — The Idea of the Second Layer	55
14 How the Brain Retrieves Information.....	57
15 Tagging in Practice: From Email Labels to Enterprise Metadata.....	65
16 Tagging Versus Categorizing.....	71

17	From Location to Context.....	77
18	From Hierarchies to Networks	81
	PART IV — Designing the Second Layer	85
19	Mapping the Existing Landscape	87
20	Starting Small: Introducing Your First Tags.....	93
21	Running the Second Layer	97
22	Designing a Personal Tagging System.....	103
23	Governance for Tags: Security, Compliance, and Risk.....	107
24	Tagging in Daily Practice.....	111
25	Advanced Workflows and Collaborative Tagging	115
	PART V — Building the Second Layer	119
26	Under the Hood: Architecture of a Second-Layer System.	121
27	AI in the Second Layer: From Manual Tags to Assisted Context	127
28	Smart Remediation Rules	133
29	Bringing Tagging Into the Windows Environment.....	139
30	Setting Up Tagging in Windows the Right Way	145
31	Why Tagging Feels Like Extra Work (At First)	151
	Part VI — Making the Second Layer Work	157
32	From Legacy Folders to a Second Layer	161
33	Making Tagging Stick	167
34	Overcoming Resistance	173
35	The Return on Context	179
36	ROI Snapshot	185
37	From Intention to Implementation: A Migration Strategy	189

PART VII: Strategy, scale and longevity	193
38 Tagging as Structural Modernization in Networked Organizations	195
39 Tagging as a Professional Discipline	199
40 Where Tagging Creates Business Value.....	203
41 Tagging as Structural Evolution: From Hierarchical Classification to Networked Organization	215
42 Tagging as Institutional Memory	233
43 Future-Proofing Your Tagging System	235
44 The Tagging Maturity Model	237
45 Tagging as a Strategic Response to Digital Complexity	239
46 Governance at Scale	241
47 Enterprise Knowledge Systems	243
48 Tagging in the Age of Artificial Intelligence	249
49 Tagging vs SharePoint Metadata vs Knowledge Graphs ...	251
50 Why Most SharePoint Metadata Projects Fail.....	257
51 The Enterprise Metadata Maturity Curve	263
52 Tagging for Project Managers	269
53 Tagging for CEOs.....	277
54 Hostile Board Q&A Simulation	285
55 The Pros and Cons of Tagging.....	289
56 Simulated skeptical CIO attack	293
57 Why the File System Is the Last Unmodernized Layer of Digital Work.....	295
58 Conclusion	301

APPENDIX A: Quick Start Guide	303
APPENDIX B: Sample Tag Frameworks	305
APPENDIX C: 90-Day Rollout Template	309
APPENDIX D: Use cases	311
References.....	315

EXECUTIVE SUMMARY

Tagging as a Structural Upgrade for Modern Knowledge Work

The Problem

- Modern professionals operate in environments characterized by:
- Overlapping projects
- Cross-functional collaboration
- Increasing regulatory layers
- Rapidly shifting priorities
- Growing volumes of digital information

Yet most digital file systems remain organized primarily through hierarchical folders — a model inherited from physical filing cabinets. Hierarchies are effective for containment but inadequate for multi-dimensional retrieval. They force exclusive placement, increase cognitive friction, and obscure cross-boundary connections.

This can result in:

- Time lost searching
- Repeated structural reorganization
- Duplication of files
- Reduced visibility across initiatives
- Increased cognitive load

The issue is not information volume. It is structural misalignment.

The Core Argument

This book argues that Tagging should be understood not merely as a feature, but as a structural upgrade.

Tagging introduces a second layer of organization based on context rather than location.

While folders answer:

“Where does this belong?”

Tags answer:

“What is this related to?”

This shift enables:

- Multi-dimensional classification
- Recognition-based retrieval
- Cross-boundary visibility
- Reduced decision fatigue
- Greater adaptability over time

The optimal system is hybrid:

Folders provide stability.

Tags provide flexibility.

Together, they reflect the reality of modern networked work.

Theoretical Foundation

Tagging is supported by multiple research domains:

- **Cognitive Load Theory (Sweller):** Reduces extraneous decision effort.
- **Spreading Activation (Collins & Loftus):** Aligns with associative memory.
- **Recognition vs Recall (Tulving):** Supports more efficient retrieval.
- **Network Theory (Barabási, Watts):** Mirrors interconnected systems.
- **Behavioral Economics (Kahneman, Thaler):** Explains adoption barriers.
- **Knowledge Management Research (Alavi & Leidner; Davenport & Prusak):** Emphasizes metadata-driven retrieval and cultural reinforcement.

Tagging aligns with both cognitive architecture and organizational complexity.

Practical Implementation

This book provides:

- A step-by-step introduction strategy
- A five-dimension tagging framework:
 - Entity
 - Initiative
 - Stage
 - Time
 - Theme
- Guidelines for avoiding over-tagging
- Governance principles for collaborative environments
- Integration methods within Windows-based workflows
- Advanced portfolio and risk management use cases

The emphasis is gradual adoption, minimal complexity, and consistent habit formation.

Strategic Implications

For leaders and organizations, tagging delivers:

- Faster executive preparation
- Portfolio-level visibility
- Improved audit readiness
- Cross-functional transparency
- Reduced structural reorganization
- Lower cognitive friction across teams

Tagging transforms the file system from static archive into contextual intelligence layer.

Final Insight

The central insight of this book is simple but structural:

Folders are not wrong.

They are insufficient alone.

In an era defined by digital complexity and networked interdependence, organizations require systems that combine hierarchical stability with contextual flexibility.

Tagging is that second layer.

When implemented thoughtfully, it is not a minor efficiency gain. It is architectural modernization.

Signature Idea

The fundamental problem of modern digital work is not the volume of information we manage. It is the structure we use to manage it.

Hierarchical folders were designed for a world in which information moved slowly and belonged to clear categories.

Today's work operates as a network of relationships—projects intersect, initiatives overlap, and documents connect to multiple contexts simultaneously.

When we force this networked reality into a rigid tree structure, friction appears.

Tagging introduces a second layer of organization that reflects how work actually happens: not as a set of isolated containers, but as a web of meaningful connections.

Core Thesis

Folders are not wrong. They are simply insufficient for the complexity of modern work.

As knowledge work becomes increasingly networked, dynamic, and multi-dimensional, hierarchical file structures can no longer carry the cognitive load we place upon them. Tagging introduces a second layer of organization — one based on context rather than location — that aligns digital systems with the associative way human memory actually functions.

This book argues that tagging is not merely a feature, but a structural upgrade in how we think, retrieve, and manage information.

Who This Book Is For

This book is written for professionals who feel that managing digital information has become harder than it should be.

It is for people who spend time searching for documents they know exist, restructuring folders that once made sense but no longer do, and navigating file systems that struggle to reflect the complexity of modern work.

The central idea of this book is simple: the way we organize digital information has not evolved at the same pace as the work we now perform.

Professionals Managing Complexity

This book is particularly relevant for professionals whose work spans multiple contexts simultaneously:

- project managers coordinating cross-functional initiatives
- consultants managing multiple client engagements
- knowledge workers handling complex documentation
- analysts and strategists working across teams and departments

In these environments, documents rarely belong to a single category. They connect to projects, clients, timeframes, strategic themes, and regulatory frameworks at the same time.

Traditional folder systems struggle to represent this complexity.

Tagging introduces a second layer of organization that makes these relationships visible.

Leaders Responsible for Information Architecture

The ideas in this book are also relevant for leaders responsible for digital structure within organizations:

- CIOs and IT leaders designing information environments
- knowledge management professionals
- digital transformation leaders
- operations and governance teams responsible for document systems

For these roles, tagging is not merely a convenience feature. It represents a structural improvement in how organizations retrieve, connect, and interpret information.

Individuals Building Personal Systems

Finally, this book is useful for individuals seeking to create more effective personal knowledge systems.

Even without enterprise platforms or complex software, the principles of contextual organization can dramatically improve how professionals manage their own files and archives.

What This Book Is Not

This is not a technical manual.

It does not attempt to replace existing enterprise systems, nor does it propose abandoning hierarchical folders entirely.

Instead, it explores how a contextual layer can complement hierarchy—transforming file systems from static storage structures into flexible knowledge environments.

The Core Question

If modern work operates as a network of relationships rather than a simple hierarchy, then our information systems must reflect that reality. This book explores how tagging can provide that missing layer.

Core Ideas of the Book

Folders are not wrong. They are simply insufficient for the complexity of modern work.

Hierarchical organization provides stability, but modern information environments require an additional layer of contextual connection.

Tagging does not replace hierarchy; it complements it.

Folders provide containment. Tags provide context. Together they create a hybrid structure capable of handling complexity.

Information rarely belongs to one place. It belongs to multiple contexts simultaneously.

Tagging allows those contexts to coexist without forcing artificial prioritization.

The true purpose of organization is not storage. It is retrieval.

Systems should be designed around how people find information, not merely where they place it.

Tagging transforms the file system from a static archive into a contextual intelligence layer.

When documents carry meaningful attributes, patterns emerge across projects, teams, and time.

PART I: Foundations

Understanding How Information Should Serve Human Work

Prologue

Every generation inherits not only the accumulated knowledge of those who came before, but also the methods by which that knowledge is organized, preserved and made accessible, and although these methods often appear so natural that they remain largely unquestioned, they profoundly influence the way individuals think, organizations operate and societies develop.

Throughout history, humanity has continuously searched for better ways of bringing order to an increasingly complex world. Libraries developed catalogues that enabled readers to navigate vast collections of books, governments created archives that preserved administrative records, museums introduced classification systems for their collections and businesses adopted filing cabinets that allowed documents to be stored and retrieved with remarkable efficiency. Each of these systems represented an important achievement because each reflected the technological possibilities and practical realities of its own time.

When computers gradually replaced paper, we naturally carried these same principles into the digital world. Filing cabinets became folder structures, drawers became directories and paper documents became electronic files, while the underlying assumption—that every document should occupy one logical location within a hierarchical structure—remained virtually unchanged. Technology evolved dramatically, but the conceptual model upon which it rested continued almost exactly as before.

For many years this approach proved entirely sufficient. Organizations became increasingly efficient at creating, storing and distributing digital information, and every new technological development appeared to reinforce the belief that larger storage capacities, faster search engines and more sophisticated document management systems would eventually solve the challenges associated with finding information.

Today, however, we find ourselves in a fundamentally different world.

Knowledge has become the primary resource of almost every organization, projects increasingly span multiple disciplines, employees collaborate across departments, organizations and continents, and artificial intelligence is beginning to participate in work that until recently depended entirely upon human expertise. Information no longer exists within isolated collections of documents, but forms a continuously evolving network of relationships, contexts and meanings that cannot easily be represented by a single hierarchical structure.

It is therefore becoming increasingly clear that the difficulties organizations experience are not primarily technological. They arise because we continue to organize digital information according to conceptual principles that were originally developed for a world in which information existed on paper, documents could occupy only one physical location and relationships between different pieces of knowledge had to be reconstructed manually by the people who used them.

This book begins with a simple observation that gradually developed over more than thirty years of working with organizations of different sizes and in different industries. During those years I repeatedly noticed that people did not become frustrated because information was missing; they became frustrated because the information they needed could not be connected to the context in which they needed it. The knowledge existed, the experience existed and the documents certainly existed, yet the relationship between them remained hidden behind structures that had originally been designed to solve an entirely different problem.

The purpose of this book is therefore not to introduce another software solution, nor to argue that existing folder structures should be abandoned, but rather to invite the reader to reconsider one of the most fundamental assumptions that has quietly shaped information architecture for more than half a century. My intention is to explore whether the next step in the evolution of information management lies not in replacing the structures we already have, but in complementing them with an additional layer that organizes information according to its meaning rather than merely according to its location.

1 The Purpose of Work

Every meaningful piece of work begins as an intention and ends as reality.

When people leave their homes each morning to begin another working day, they rarely do so because they look forward to navigating complex organizational procedures, searching through endless folder structures or spending valuable time locating information that should already be available to them. They leave because they want to contribute to solving problems, to creating something meaningful or to help another person, and although the nature of that contribution differs enormously between an engineer, a project manager, a teacher, a nurse or an entrepreneur, the underlying motivation is remarkably similar. Every one of them begins the day with an intention.

That intention may be as modest as completing a single task before lunch or as ambitious as delivering a project that has occupied months or even years of work. Whatever its scale, every meaningful activity starts with a picture of a desired outcome that exists first in the mind before it gradually becomes part of reality. Work, therefore, should not primarily be understood as a collection of tasks or responsibilities. At its deepest level, work is the continuous process through which human intention is translated into action until it eventually becomes reality.

This seemingly simple observation has accompanied me throughout my professional life and has gradually become the foundation beneath everything I have built. Long before I became interested in information management, I became fascinated by a much more fundamental question. Why do some working environments allow people to perform their work with energy, satisfaction and creativity, while others gradually drain that same energy even when the people themselves are talented, motivated and committed?

The answer, I discovered, was seldom found in the capabilities of the people themselves. More often it was hidden within the countless small interruptions that quietly accompany almost every working day. Some of these interruptions are inevitable because reality is unpredictable. Customers change their minds, markets evolve, machines fail and unexpected events require immediate attention. Such interruptions belong to the nature of work itself.

Many other interruptions, however, are entirely different. They are not created by reality but by the way organizations choose to organize themselves. Employees wait for approvals that add little value, attend meetings that produce few decisions, recreate work that has already been completed elsewhere or spend considerable amounts of time searching for information that they know must exist somewhere within the organization. None of these activities contribute directly to the purpose with which they began their work. Instead, they consume attention, interrupt concentration and gradually reduce the satisfaction people experience from accomplishing meaningful tasks.

During the years in which I founded and led several companies, I slowly became aware that these seemingly insignificant interruptions had a surprisingly profound influence on the way people experienced their work. They did not merely reduce productivity; they reduced the natural flow through which intention becomes reality. Whenever that flow was interrupted unnecessarily, frustration appeared almost immediately. Conversely, whenever people were able to continue their work without unnecessary obstacles, a different atmosphere emerged. People became more relaxed, more engaged and, perhaps most importantly, happier.

This realization eventually led to the guiding principle that would accompany my later work: Work Smart, Be Happy. Many people interpreted these words as a slogan promoting efficiency or productivity. For me they always meant something entirely different. They expressed a belief that organizations have a responsibility to remove unnecessary friction wherever it appears, because every unnecessary obstacle diminishes the opportunity people have to accomplish the work they genuinely came to do.

The central question therefore became remarkably simple, although its implications continue to shape everything that follows in this book. How do we create environments in which people can naturally transform intention into reality with as little unnecessary friction as possible? Every page that follows is, in one way or another, an attempt to answer that question.

If work is the transformation of intention into reality, then we should perhaps ask why that transformation so often becomes unnecessarily difficult.

2 The Invisible Friction Within Organizations

Organizations rarely notice the friction they gradually learn to accept.

One of the most remarkable characteristics of organizational life is our extraordinary capacity to become accustomed to inefficiency. Behaviours that would initially be regarded as temporary inconveniences gradually become accepted as normal, and after sufficient repetition they eventually disappear from conscious attention altogether. Employees adapt to them; managers stop questioning them and organizations quietly begin treating them as unavoidable aspects of everyday work.

This process of normalization fascinated me for many years because it occurred almost everywhere I looked. People accepted that searching for information might consume valuable time. They accepted that colleagues often possessed knowledge that official systems could not provide. They accepted that documents were frequently recreated because existing versions could no longer be located. Above all, they accepted that asking another person had become a more reliable method of retrieving information than consulting the organization's own information systems.

None of these behaviours appeared extraordinary to the people performing them. On the contrary, they had become so common that they were rarely questioned. Searching was simply considered part of work.

I never accepted that conclusion.

The reason was not that searching required time, although it certainly did. My concern was that searching interrupted the very activity people had originally intended to perform. The engineer who spends fifteen minutes locating the latest drawing is not engaged in engineering. The project manager searching through dozens of folders is not managing the project. The financial specialist trying to identify the correct spreadsheet is not analyzing financial performance. In every case, meaningful work temporarily stops while attention shifts toward locating the resources required to continue.

It gradually became clear to me that organizations were unknowingly introducing friction into activities that should have flowed much more naturally. The problem was not that people lacked motivation or competence. The problem was that organizational systems repeatedly diverted their attention away from the work they actually wanted to accomplish.

The consequences extended far beyond lost productivity. Every unnecessary interruption subtly influences the emotional experience of work itself. People become tired more quickly, lose concentration more easily and gradually lower their expectations of what an effective working environment should feel like. Eventually they no longer recognize these interruptions as problems because they have become part of the accepted reality of organizational life.

Perhaps the most dangerous aspect of invisible friction is precisely this process of normalization. Once organizations begin believing that searching, waiting and recreating work are unavoidable, they stop looking for fundamentally different solutions. Instead, they concentrate on making inefficient processes slightly less inefficient, while leaving the underlying assumptions completely untouched.

History repeatedly demonstrates that genuine innovation rarely emerges from improving existing assumptions. It emerges from questioning whether those assumptions remain valid at all.

Among all the forms of friction that organizations create, none is more persistent than the difficulty of finding the information people already know exists.

3 From Documents to Meaning

Nobody searches for a document. People search for what the document allows them to understand.

Among the many forms of organizational friction that I encountered during my career, one particular pattern returned with remarkable consistency. People were rarely unable to find information because the information did not exist. Far more often, they were unable to find the document that happened to contain it.

At first sight this distinction appears insignificant. Upon closer examination, however, it reveals one of the most fundamental misunderstandings underlying contemporary information management.

Organizations generally believe they are managing information. In reality they manage documents. Word files, spreadsheets, presentations, reports and PDF documents become the primary objects around which entire information systems are designed. Folder structures determine where these documents are stored, document management systems regulate their versions and permissions, and search engines attempt to retrieve them as efficiently as possible.

The people working within those organizations, however, rarely seek documents for their own sake. They seek understanding. They seek the decision that was made during a previous meeting, the lesson learned from an earlier project, the technical specification that explains a design choice or the contractual agreement that clarifies a particular obligation. The document itself merely functions as the vessel through which that knowledge is temporarily expressed.

This distinction changes everything.

A document occupies one location within a folder hierarchy, but the meaning contained within that document simultaneously belongs to many different contexts. An engineering drawing may relate to a project, a client, a supplier, a geographical location, a contract, a maintenance program, a specific risk category and numerous technical disciplines at the same time. The hierarchy forces the organization to select only one of these perspectives, while the people using the information naturally think in terms of all of them.

It is therefore hardly surprising that employees gradually develop their own methods for overcoming this limitation. They create personal folder structures, maintain collections of shortcuts, remember where colleagues store information and increasingly rely upon personal relationships rather than formal systems. These behaviours should not be interpreted as resistance against organizational procedures. They are entirely rational attempts to restore the context that the information architecture itself has failed to preserve.

This observation ultimately led me to a realization that would shape everything that followed. Organizations do not primarily struggle because they possess too little information. They struggle because the relationship between information and its many possible meanings has become hidden behind an architecture that was originally designed for storing documents rather than supporting human understanding.

To understand why our systems continue to organize documents rather than meaning, we must first understand the architecture we inherited.

4 The Architecture We Inherited

Every generation inherits the information architecture of the previous one.

It is easy to assume that the systems we use every day were deliberately designed for the world in which we now live. In reality, almost every information system reflects the assumptions, limitations and possibilities of an earlier age. What appears obvious to one generation often originated as an ingenious solution to problems that no longer exist.

This has always been true throughout history. Cities continue to follow the routes of ancient roads long after the original reasons for those roads have disappeared. Legal systems still contain principles that were established centuries ago because they solved important problems in their own time. Languages preserve expressions whose original meanings have long since been forgotten, while organizational structures often continue to reflect realities that disappeared decades earlier.

Information architecture is no different.

Long before computers existed, organizations had to solve a practical problem. Information was recorded on paper, and every piece of paper could physically exist in only one place at any given moment. A contract could not simultaneously be stored in two filing cabinets without creating duplicate copies. Engineering drawings occupied large physical archives. Personnel files filled entire rooms. Libraries devoted enormous effort to determining the single location where every book should be placed, because once that decision had been made, readers could reliably find it again.

Hierarchy was therefore not an arbitrary invention. It was an elegant response to a physical limitation. When information can exist in only one location, every organization must decide where that location will be.

The filing cabinet became one of the greatest organizational innovations of the twentieth century because it introduced consistency into an increasingly complex world. Once everyone understood the underlying structure, documents became easier to store, retrieve and manage. For its time, hierarchy represented a remarkable achievement.

When computers gradually entered the workplace, they inherited this architecture almost unchanged. Paper folders became digital folders. Filing cabinets became directory structures. Archives became servers. Although the physical limitations disappeared almost overnight, the conceptual model remained exactly the same. Every document still occupied one logical location within a hierarchy, not because digital technology required it, but because this was simply how organizations had learned to think about information.

For many years this approach continued to serve organizations remarkably well. Digital storage increased rapidly, computers became faster and search capabilities improved continuously. Every technological advance appeared to confirm that the existing architecture merely required better software, more storage capacity or more sophisticated search algorithms.

Few people questioned the architecture itself.

This is hardly surprising. Successful systems rarely invite fundamental questions while they continue to produce acceptable results. Most innovations therefore focus on improving the existing model rather than reconsidering the assumptions upon which that model was originally built.

Yet organizations themselves were quietly changing.

Projects increasingly crossed departmental boundaries. Customers expected integrated solutions rather than isolated products. Teams became multidisciplinary, while digital collaboration allowed people to work together across countries and continents. Information no longer belonged exclusively to one department, one project or one discipline. Instead, every piece of knowledge gradually became connected to an expanding network of people, activities, responsibilities and relationships.

Reality had become dynamic.

The architecture remained static.

This growing difference between the way organizations actually worked and the way information continued to be organized explains many of the frustrations that employees gradually accepted as unavoidable.

It also explains why organizations repeatedly reorganized their folder structures. Every organizational change required the hierarchy to be adjusted, because the hierarchy attempted to represent a reality that had already begun changing again.

The problem, therefore, was never that hierarchy was poorly designed.

The problem was that it was designed to solve a different problem.

It solved the problem of storing information in a world where every document occupied one physical location.

It could never fully solve the challenge of helping people navigate knowledge that simultaneously exists within many different contexts.

Recognizing this distinction does not diminish the importance of hierarchy. On the contrary, it allows us to appreciate its historical significance while also acknowledging its natural limits. Every successful architecture eventually reaches the boundaries of the problem it was originally designed to solve. When that moment arrives, genuine progress rarely comes from abandoning the existing architecture altogether. More often, progress begins by asking whether the original architecture should be complemented by something new.

If every hierarchy was originally created to solve one particular problem, we must now ask ourselves a more fundamental question. What assumption lies hidden within every hierarchy, and does that assumption still correspond to the reality of modern work?

5 The Hidden Assumption

Every hierarchy solves one problem while creating another.

The realization that hierarchy was solving the wrong problem did not come to me while reading a book. It came during an ordinary management meeting.

Every architecture begins with an assumption, although the people using that architecture are rarely aware of its existence. Assumptions gradually disappear from view because successful solutions become so familiar that they are no longer questioned. Eventually, the solution itself begins to appear inevitable, while the original problem that gave birth to it quietly fades into history.

Information architecture is no exception.

For more than half a century organizations have organized their digital information around hierarchical folder structures, not because hierarchy represents the only possible way of organizing knowledge, but because hierarchy once provided the most practical answer to a very specific question.

Where should this document be stored?

It is worth pausing for a moment, because this question has shaped almost everything that followed.

Notice that the question is not concerned with the information itself. It is not concerned with the knowledge contained within the document, the decisions it records, the people who contributed to it or the many situations in which it may later become valuable. The question concerns only one thing: location.

Where should it be placed?

Once that question becomes the foundation of an architecture, every subsequent decision follows almost automatically. A hierarchy is created. Categories are defined. Departments receive their own folders. Projects receive their own subfolders. Naming conventions are introduced.

Governance is established. Permissions are assigned. Everything becomes focused on ensuring that every document occupies one clearly defined place within an increasingly elaborate structure.

For many years this approach functioned remarkably well because it solved precisely the problem it had been designed to solve. When every document occupied a single physical location, deciding where that location should be was both necessary and logical.

The hidden assumption was never questioned.

It simply became accepted that every document should have one logical place.

Digital technology inherited this assumption without hesitation.

Although computers removed the physical limitation that required every document to exist in only one location, our thinking remained remarkably similar. We continued asking the same question our predecessors had asked for generations.

Where shall we store this document?

Yet something subtle had changed.

The physical world had imposed that question upon us.

The digital world no longer did.

A digital document can be copied instantly. It can appear simultaneously on multiple screens. It can be accessed from different countries at exactly the same moment. Artificial intelligence can analyse thousands of documents without ever considering the folder in which they happen to reside.

The limitation disappeared.

The assumption remained.

This distinction is more significant than it first appears.

Once organizations begin asking where information should be stored, they naturally optimize their systems for storage. Every improvement strengthens the hierarchy. Folder structures become more detailed. Naming conventions become more precise. Search engines become faster. Metadata becomes richer.

Yet the original question remains unchanged.

Where should we store this?

Perhaps that is no longer the most important question.

Suppose, instead, that we began somewhere else.

Suppose the first question was not about storage but about work.

Suppose we asked:

How will people naturally want to retrieve and use this information?

Immediately the architecture begins to look different.

The same engineering drawing may be relevant to the project manager planning tomorrow's activities, the procurement specialist ordering materials, the maintenance engineer responsible years later for servicing the installation, the legal department reviewing contractual obligations and the financial department evaluating project costs. None of these people think primarily about the folder in which the drawing resides. Each approaches exactly the same information through an entirely different context.

The information has not changed.

The question has.

That single change transforms the architecture.

Instead of organizing information solely according to where it belongs, we begin considering why it matters, to whom it matters and under which circumstances it becomes valuable. The emphasis gradually shifts from