

UNIVERSITY OF CALIFORNIA,
IRVINE

Supporting Reuse: IT and the Role of
Archival Boundary Objects
in Collaborative Problem Solving

DISSERTATION

submitted in partial satisfaction of the requirements for the degree of

DOCTOR OF PHILOSOPHY

in Information and Computer Science

by

Wayne Garrett Lutters

Dissertation Committee:
Professor Mark Ackerman, Chair
Professor Paul Dourish
Professor Alladi Venkatesh
Dr. Steven Poltrock

2001

© 2001 Wayne Garrett Lutters
All rights reserved.

This dissertation of Wayne Garrett Lutters is approved
and is acceptable in quality and form
for publication on microfilm:

Committee Chair

University of California, Irvine
2001

DEDICATION

To

the engineers at GTS-West,

exceptional professionals who keep those birds in the air, day after day.

To

my parents

for their unfailing interest and support.

SOLI DEO GLORIA

TABLE OF CONTENTS

	Page
<i>LIST OF FIGURES</i>	<i>x</i>
<i>LIST OF TABLES</i>	<i>xi</i>
<i>ACKNOWLEDGEMENTS</i>	<i>xii</i>
<i>CURRICULUM VITAE</i>	<i>xiii</i>
<i>ABSTRACT OF THE DISSERTATION</i>	<i>xv</i>
<i>CHAPTER 1: The Promise and Problem of Reuse</i>	<i>1</i>
1.0 The promise of reuse in organizations	1
1.1 The problem of reuse in organizations	5
1.2 Research response	7
1.3 Overview of the dissertation	8
<i>CHAPTER 2: Knowledge Work and Collaborative Problem Solving</i>	<i>12</i>
2.0 Introducing Computer-Supported Cooperative Work	12
2.1 Lessons from Computer-Supported Cooperative Work	13
2.2 Social use of information as mediated by technical artifacts	15
2.2.1 Processes of Contextualization	18
2.2.2 Boundary Objects	21
2.3 Social use of information as mediated by people	25
2.3.1 Hubs in the knowledge network	26
2.3.2 Social mediation in technical support	28
<i>CHAPTER 3: Boundary Objects and Reuse</i>	<i>35</i>
3.0 CSCW and Knowledge Management	35
3.1 Suitability of boundary objects	35
3.2 Addressing the alternatives: coordination mechanisms	38
3.3 Research questions	39

<i>CHAPTER 4: Study Design / Methodology</i>	42
4.0 Research approach	42
4.1 Site selection	43
4.1.1 Technical support venues	43
4.1.2 Extreme environments	44
4.1.3 High-reliability organizations	45
4.1.4 High-liability organizations	45
4.1.5 Global Airframe: Technical Support Western Division	46
4.2 Study Design	46
4.3 Data collection techniques	48
4.3.1 Understanding the People	48
4.3.2 Understanding Information Flows and Artifacts	53
4.3.3 Understanding the Work	60
4.4 Data analysis techniques	62
<i>CHAPTER 5: Organization, People, Systems, and Environment</i>	64
5.0 Requisite background information	64
5.1 Organization	64
5.1.1 The corporation: Global Aerospace	64
5.1.2 The corporate division: Global Airframe	65
5.1.3 The functional division: Technical Services	66
5.1.4 The Facility: Global Technical Services-West	67
5.1.5 GTS-West: Departmental overview	68
5.1.6 GTS-West: Airframes	70
5.1.7 GTS-West: Analytic support teams	72
5.1.7.1 Stress	73
5.1.7.2 Other teams	76
5.2 People	77
5.2.1 Director	77
5.2.2 Senior management team	78
5.2.3 Group managers	78
5.2.4 Team leads	79
5.2.5 Engineers	80

5.3	Systems	83
5.3.1	GlobalCOM	84
5.3.2	ROC system	85
5.3.3	GlobalCOM / ROC system compared	87
5.4	Environment	88
5.4.1	Productive culture of conflict	89
5.4.2	Union issues	90
5.4.3	The current dilemma	91
 <i>CHAPTER 6: Understanding Routine Work</i>		 94
6.0	Defining the work of GTS-West	94
6.1	Collaboration essentials	94
6.2	Service expectations, constraints, and abilities	96
6.3	Classes of routine actions	98
6.4	Prioritization	100
6.5	AOG scenario walk-through	101
6.6	Additional Complexity	105
6.6.1	FAA	106
6.6.2	Global Airframe	108
6.6.3	Media, NTSB and public accountability	109
 <i>CHAPTER 7: Boundary Objects in Collaborative Problem Solving</i>		 111
7.0	Boundary objects revisited	111
7.1	The primacy of the ROC as boundary object	112
7.2	ROC as boundary object illustrated	114
7.2.1	An example	115
7.2.2	An example unpacked	116
7.2.3	Boundary object diagramming technique	119
 <i>CHAPTER 8: Coordination and Collaboration, Inter-group Boundary Spanning</i>		 122
8.0	Embeddedness of boundary objects in routine work practice	122
8.1	Revisiting the nature of the Stress / Structures boundary	122
8.2	The ROC as foundational inter-group facilitation mechanism	124

8.3	Completeness	127
8.4	Social control structures for the ROC	132
8.4.1	Need for reliability	132
8.4.2	Need for control over archival process by Stress	133
8.4.3	Need for control over customer presentation by Structures	134
8.4.4	Resolution achieved through split views	134
8.5	Managing cross-group expertise	136
8.5.1	Job continuation	138
8.5.2	Perceived expertise	138
8.5.3	Compliance with historical precedent	139
8.6	Impact of ROC Visibility	140
8.7	The interdependence of the Stress Team Lead and the ROC	143
8.7.1	Enactment of Perceived Authority	144
8.7.2	Job assignment	145
8.7.3	Availability	147
8.7.4	Defender	148
8.8	Integration of external expertise	149
8.8.1	Within GTS-West	149
8.8.2	Within Global Airframe	151
8.8.3	Outside of Global	151
<i>CHAPTER 9: Negotiation & Translation, Interorganizational Boundary Spanning</i>		<i>154</i>
9.0	Boundary: GTS-West / Operator	154
9.1	Negotiating to completeness	156
9.2	An Example	156
9.3	Opaqueness of routine work	160
9.4	Real and perceived operator abuses	161
9.5	Information elicitation	162
9.6	Coping with intractable operators and absent data	163
9.7	Due dates	165

CHAPTER 10: <i>Recontextualization and Reuse, Boundary Spanning Across Time</i>	167
10.0 The notion of boundary objects across time	167
10.1 Introducing the practice of reuse at GTS-West	168
10.2 Culture of reuse	169
10.3 Methods of reuse	171
10.3.1 Precise matching	171
10.3.2 Repurposing elements	172
10.3.3 Precedent setting	174
10.3.4 Addendum: Procedural Reuse	176
10.4 Common breakdowns and resolutions	178
10.4.1 Decontextualization	178
10.4.2 Recontextualization	179
10.4.3 The special case of aging records	181
10.5 Example	183
10.4 Formal / informal archives	188
10.4.1 The Folder system	189
10.5 Keys embedded in the environment	191
10.5.1 Correspondence identifiers	191
10.5.2 Serialization	192
CHAPTER 11: <i>Implications for Design</i>	194
11.0 Introduction	194
11.1 Designing for negotiation points	194
11.2 Designing for recontextualization	196
11.2.1 ROC Completeness	197
11.2.2 Resource coverage	199
11.2.3 Similar cases	200
11.2.4 Awareness of knowledge/expertise network	201
11.3 Designing for inter-organizational boundary spanning	202
11.4 GlobalCOM integration revisited	204

<i>CHAPTER 12: Conclusions</i>	206
12.0 Reuse at GTS-West	206
12.1 Review of key findings	206
12.1.1 Integrating boundary objects for knowledge management	206
12.1.2 Leveraging boundary objects to enhance compliance	208
12.1.3 Gatekeepers	208
12.1.4 Methods of enhancing recontextualization	209
12.1 Extreme organizational stressors revisited	210
12.1.1 High-Reliability	210
12.1.2 High-Liability	211
12.2 Beyond boundary objects	212
12.2.1 Representing the embedded nature of boundary objects	212
12.2.2 Representing the trajectory of use for boundary objects	213
 <i>REFERENCES</i>	 215

LIST OF FIGURES

Figure 1: GTS-West organizational chart	69
Figure 2: Stress workspace	82
Figure 3: Structures workspace	82
Figure 4: Sample GlobalCOM request	85
Figure 5: Sample ROC request	87
Figure 6: Collaborating teams	96
Figure 7: Information flows	102
Figure 8: Information flows	106
Figure 9: Information flows	116
Figure 10: Boundary object diagram	121
Figure 11: Information flows	183
Figure 12: Boundary object diagram	187

LIST OF TABLES

Table 1: Comparison of Technicians	31
Table 2: GTS-West Departmental Overview	68
Table 3: Comparison of GlobalCOM and ROC Systems	88
Table 4: Prioritization	101
Table 5: Problems with Decontextualization	179
Table 6: Problems with Recontextualization	180

ACKNOWLEDGEMENTS

This research has been supported by grants from the CISE/IIS/CSS Division of the U.S. National Science Foundation and the NSF Industry/University Cooperative Research Center (CISE/EEC) to the Center for Research on Information Technology and Organizations (CRITO) at the University of California, Irvine. I am pleased to have been a participant in the engaging community of multi-disciplinary scholars at CRITO over the years.

First, I am thankful for Global Airframe's interest in this study and for the exceptional hospitality of the GTS-West engineers. That they were able to welcome a student researcher poking and prodding for over a year, despite their intense workload, stands as a testament to their passion for improvement and ideals of service.

This work matured intellectually through interactions with Mark Ackerman, Paul Dourish, Steve Poltrock, Alladi Venkatesh, Jonathan Grudin, Gloria Mark, Jim Danziger, David Redmiles and John King. I thank them for their interest and challenging questions. In addition, the doctoral consortium participants at both AMCIS '00 and CSCW '00 contributed to the evolution of this work.

I would especially like to thank Mark Ackerman for his intellectual partnership these past eight years. I have appreciated his unique perspective on the world, maturity as a scholar, and patience with our sometimes problematic projects. It has been an honor to have been his student and colleague.

I am very grateful for the numerous friends I have in the Social Worlds research group, especially David McDonald, who boldly scouted the path for me, and Jack Muramatsu, who is ever ready for a good debate, the CORPS program both past and present (esp. Leysia Palen, Young Kim, Rob Klashner, and Madhu Reddy), around the department (esp. Cathy Blake, Peter Kammer, Michael Kantor and Shilpa Shukla), and all those Taco Tuesdays regulars.

Suzanne Schaefer, my exceptional officemate, inspirational writing partner, and black-belt gremlin fighter, is due tremendous thanks. Both I and this dissertation owe her much for her generously shared friendship, laughter, and wisdom.

My grandmothers, Sophie Cutler and Agnes Lutters, are deserving of special thanks for their financial and prayerful support through the years. Yes, I am *finally* done.

Lastly, I am grateful for my parents, Ken and Carol, and their consistent support in my pursuit of this dream. (Special thanks as well for my father's outstanding "just in time" editorial contributions.)

CURRICULUM VITAE

Wayne G. Lutters

EDUCATION

Ph.D., Information and Computer Science University of California, Irvine	September, 2001
M.S., Information and Computer Science University of California, Irvine	December, 1995
B.A., Cognitive Science and History Connecticut College	May, 1992

HONORS

Regents' Dissertation Fellowship, University of California, Irvine	2001
Regents' Fellowship, University of California, Irvine	1994
Phi Beta Kappa, Connecticut College	1992
Hannah Grace Roach Prize in European History, Connecticut College	1992
Frederick H. Sykes Scholar, Connecticut College	1989

RESEARCH INTERESTS

Computer-supported cooperative work (CSCW), human-computer interaction (HCI), computer-mediated communication (CMC), knowledge management, social analysis of computing, electronic social spaces, computerization in domestic environments and history of technology.

RESEARCH EXPERIENCE

Assistant Professor, Information Systems Department University of Maryland, Baltimore County	2001-Present
Research Associate, Center for Research on IT and Organizations University of California, Irvine	1999-2001
Research Associate, Department of Information & Computer Science University of California, Irvine	1995-1999
Usability Engineer (Intern) Microsoft Corporation	1997-1998

PUBLICATIONS

Lutters, Wayne G., Mark S. Ackerman, James Boster and David W. McDonald. (2000). "Mapping Knowledge Networks in Organizations: Creating a Knowledge Mapping Instrument." *Proceedings of the Americas Conference on Information Systems*, pp. 2014-2018.

Lutters, Wayne G. (2000). "Supporting Knowledge Reuse: A Field Study of Service Engineers in a High-Reliability Organization." *Proceedings of the Americas Conference on Information Systems*, pp. 2129-2131.

Lutters, Wayne G. (1998). Book review of "Virtual Culture: Identity and Community in Cyberspace" (Steven G. Jones, Ed., Thousand Oaks, CA: Sage Publications, 1997) for *Library and Information Science Research*, 20(1), pp. 106-107.

Lutters, Wayne G. and Mark S. Ackerman. (1997). "A Collectivity in an Electronic Social Space: Examining Mouse World." *SIG-Group Bulletin*, 18(1), pp. 41-43.

ABSTRACT OF THE DISSERTATION

Supporting Reuse: IT and the Role of Archival Boundary Objects
in Collaborative Problem Solving

By

Wayne G. Lutters

Doctor of Philosophy in Information and Computer Science

University of California, Irvine, 2001

Professor Mark S. Ackerman, Chair

Critical information artifacts are often the product of complex social and organizational processes. However, much of the context surrounding these processes is routinely lost in the archiving of these artifacts, significantly impacting later reuse. How then can organizational practice and IT be best configured to support knowledge management?

This dissertation examines knowledge work in a high-reliability organization. Specifically, it explores the collaborative problem solving behavior of service engineers and their analytic support teams for a world-class aircraft manufacturer. This ethnographic field study focuses on the organizational memories, information flows, and boundary objects which facilitate the routine handling of technical support requests from airlines. Special attention is given to the expertise required to successfully navigate the complexities of this information-intensive environment. High-level requirements are generated which inform the socio-technical design of appropriate knowledge management systems.

Primary contributions of this work to the fields of Computer-Supported Cooperative Work and Knowledge Management include a deeper understanding of the social factors which contribute to successful reuse in high-reliability, high-liability organizations, an advanced requirements analysis for improving IT support for reuse, and a critical assessment of the analytic concept of boundary objects.

CHAPTER 1: The Promise and Problem of Reuse

1.0 The promise of reuse in organizations

Despite the near universal experience of reuse,¹ the issue has come to the fore in organizational analysis and systems design within the past decade. Multiple disciplines are engaging the issues surrounding reuse with empirical programs ranging from organizational learning to expert systems to organizational memory to knowledge management. There are many motivations for the new focus. Since understanding this confluence of social trends is important framing for the arguments presented in this dissertation, a few of the most relevant will be introduced here.²

The first trend is increasing globalization. As our world metaphorically shrinks, competition has increased in nearly all dimensions of human enterprise, from service to art to manufacturing. All of this affords exciting new opportunities, which individuals and organizations struggle to capitalize upon. This increased awareness and competition is concomitant with a faster pace. The accelerating rate of obsolescence for ideas, styles,

¹ The term reuse is frequently used in the abstract, general sense throughout this dissertation. In the real world reuse is fluid concept; few knowledge workers consciously shift from data reuse to information reuse, they simply reuse. It is only in academic analysis that reuse is dissected into its constituent parts. Of course there is value to this, but there is also great utility in examining the behavior holistically. For the record, this work adheres to the standard definitions in the information science community: data are details devoid of context, information is data that has been contextualized, knowledge is information augmented with human reasoning, and meta-knowledge, or wisdom in some circles, is an evaluative understanding of knowledge. The former two may exist independently in information systems while the latter two can exist only in the human mind.

² Many of the observations in this section are derived from Prusak's preface to (1997) and from (Davenport and Prusak, 1998).

fashions, products, and even business plans, is staggering. Staying ahead is a daunting challenge.

In attempting to cope with these ever-increasing rates of change, many organizations have found tremendous value in the collective knowledge of their employees.³ It is this knowledge that allows for quick, creative responses to complex, ever-shifting conditions. It is what provides the organization with a critical, competitive edge. As a result, knowledge has come to be viewed as a key organizational asset, one that, like all other organizational resources, requires understanding and active management.

The second trend involves the fluidity of workforce. In recent years one particular segment, knowledge workers, has become increasingly more mobile. This is partly in response to the instability of organizations and gyrations in markets brought about by globalization and partly by dramatic improvements in communication technologies. Organizations are left struggling to retain their collective knowledge and identity in the face of ever-higher rates of employee turnover. In prior decades, much of this knowledge remained steady with a solid, stable workforce. Prusak notes that in such settings, “knowledge becomes embedded in the firm’s routines and culture. New recruits learn from old hands purely by working alongside them, an exposure and seasoning is a far more important learning mechanism than training.” (Prusak, 1997, p. x). In the modern,

lean-enterprise world of perpetual outsourcing, these traditional knowledge retention and transmission mechanisms are beginning to stress and fracture. Intelligent organizations are paying close attention to these forces and attempting to better understand the role of knowledge in their processes in order to stem potential collapse.

The third trend is a direct result of the “information revolution.” Simply put, in many domains, symbol manipulators, those who create value and meaning from pattern manipulation, have become predominant. Both qualitatively and quantitatively, this is a very different kind of work than physically manipulating artifacts of wood, ceramic or steel. In these information disciplines, information itself has become the very substrate of work. As organizations rely more and more upon this segment of traditionally semi-autonomous workers, they need to understand the nature of knowledge in action so that they can better support their employees.

The fourth trend involves the rapid expansion of information technologies (IT). There are two key elements here. The first is the exponential growth of storage and processing capacities of IT. The amount of data and information that organizations can accumulate and analyze has grown lock step with the technological ability to do so. The process of sifting through and assigning value to this vast stockpile to meet the current task-oriented needs of the organization is inherently a knowledge-based task. Although even more

³ For many organizations this discovery was poignantly accidental. In their headlong rush to remain competitive in the global economy, many firms embraced business process re-engineering and total quality management programs which radically reconstructed staffing and processes. It was only when the knowledge network was perturbed by reformulated processes (losing embedded tacit knowledge) and

noteworthy than these capacity increases is the rapid ascension of the digitally networked world. As networked computing environments have spread through organizations, the value of their communication capacity has far outweighed their computational capacity. Through such informal groupware systems as electronic mail, instant messaging, and message boards, and through more formal workflow or organizational memory systems, individuals are interacting and collaborating with more people over greater distances than ever before. Knowledge transmission is an inherently social activity and increasing the potential social network for any given individual is advantageous. The use of the Internet, or intranets, has not only increased the amount of information available to knowledge workers, but increased the potential for building shared contexts.

The ability to create, maintain, and exchange knowledge in the service of reuse, or knowledge management, offers significant hope to address the issues raised by these trends. Organizations are right to look long and hard at their knowledge resources. A chorus of authors and consultants sing the praises of what can result: prevention of knowledge loss, cost savings, improved decision making, enhanced collaboration, smooth workflows, increased innovation, adaptability and flexibility to change, new competitive advantages, asset development, product enhancements, superior customer management, advances in training, and better use of human resources. (Radding, 1998, pp. 59-61)

downsizing (flushing expertise out of the organization by way of the pink slip) that its importance was evident. It became most visible in its absence, in its silence.

1.1 The problem of reuse in organizations

In the face of such grand expectations, the end results of the vast majority of knowledge management projects fall woefully short. They often simply cannot deliver on these promises of reuse and if they do deliver, it is rarely in proportion to expectation. A managerial backlash is starting to form and knowledge management is beginning to be viewed as a passing fad in some circles. This is an irrational response, as an increased reliance on reuse does hold very real and tremendous promise. There are fundamental problems though, which are deeper than a simple matter of an over-hyped concept.

At its root, the knowledge management predicament is the result of a host of deep conceptual flaws underlying most projects. Intellectually born of the expert systems and decision support experiments of the late 1980's, knowledge management has suffered from many of the weaknesses of those programs, namely an overemphasis on computationally facile artifacts (e.g. manuals or e-mail), a disproportionate weight given to explicit knowledge over tacit knowledge, and a tendency to assess the most value to repositories of contextualized information instead of human knowledge networks. In its worst cases this has led to attempts to replace human experts with static information sources; in its best cases it has resulted in a tendency to conflate information management with knowledge management.

Fahey and Prusak (1998), in a retrospective survey of five years of knowledge management projects, came to the conclusion that "efforts to manage knowledge are themselves based upon faulty knowledge principles." In one of the most reasoned

critiques to date they identified a core set of conceptual concerns, dubbed the “eleven deadliest sins.” While all of these eleven are addressed to some degree by this dissertation, the following four are core concerns.

- “Emphasizing knowledge stock to the detriment of knowledge flow” – Knowledge management efforts have a lingering preoccupation with knowledge as a thing, instead of as a web, process, or flow.
- “Not understanding that a fundamental intermediate purpose of managing knowledge is to create shared context” – The ability of knowledge to assist in building shared understandings among participants may be its greatest contribution to decision making and collaborative problem solving. However, little effort is made to encourage these communicative activities.
- “Disentangling knowledge from its uses” – Knowledge, infusing information with meaning, is inherently contextual. However, many knowledge management efforts seek to abstract knowledge so far from its context that it is rendered worthless. (The most common way of doing this is removing the “knowledge” from the very individuals who can successfully interpret and use it.)
- “Focusing on the past and the present and not the future” – Knowing the future audience and their intended use for knowledge being captured today, with a high degree of precision, is nearly impossible. However, this is precisely where knowledge

management is the most powerful. So it is useful to try and project future use based on current trajectories.

In addition to these conceptual concerns, many knowledge management efforts have been beset by a host of practical problems as well. Three common examples are *employee participation*, where difficulty in adoption precludes reaching the necessary critical mass for the system to be truly useful, *organizational incongruity*, where the system is at odds with the organizational culture or incentive structure, and lastly *relevance*, where in a frenzied desire not to let any knowledge leak out of the organization, many firms have archived massive amounts of data, hopelessly burying useful information in the heap.

1.2 Research response

Fahey and Prusak's descriptive examination of the conceptual problems with knowledge management efforts, their prescriptive suggestions for improvement, and the bevy of additional practical concerns, all warrant deeper, empirical investigation. As such they provide the motivation for this study.

Though reuse is a complex and nuanced phenomena, it is occurs naturally, every moment, in the daily activities of knowledge workers. People seamlessly interpret current information, integrate prior knowledge, and negotiate professional relationships to solve contemporary problems. This study seeks to offer a deeper understanding of natural reuse behavior by examining a class of canonical knowledge workers, service engineers in a technical support department. The intent of performing an extended qualitative field

study is to ground these examinations firmly in context and study their unfolding within the routine of daily work practice. The ultimate goal is to better inform the future design of socio-technical systems that support knowledge management.

The next section will provide an overview of this dissertation document, outlining how the exploration of these issues unfolds throughout its twelve chapters.

1.3 Overview of the dissertation

This dissertation is organized as follows.

Chapter 1: The Promise and Problem of Reuse. This chapter frames the motivation for the study within the context of modern knowledge management efforts.

Chapter 2: The Role of Reuse in Collaborative Problem Solving. This chapter reviews the relevant literature from the CSCW research community. In two central sections it examines the notion of routine information use as mediated by both technology and people. The first section introduces the key analytical construct of boundary objects and the processes of contextualization, while the second addresses the roles of particular individuals in an organizations' information flows. The chapter concludes with an in-depth look at four studies of technical support venues.

Chapter 3: Boundary Objects and Reuse. This chapter returns to address the challenges of knowledge management as outlined in the first chapter and assess the

relevance of the boundary object concept to the study of reuse. The core research questions are also presented.

Chapter 4: Study Design / Methodology. This chapter presents the core methodological stance of the research project. It provides the rationale for the site selection criteria and details the selection process. The extended discussion of study design and data collection techniques is structured around the three central themes of the research: understanding the people, the technology, and the work itself.

Chapter 5: Organization, People, Systems, and Environment. This chapter provides the requisite background information to understand the context of the study. It examines the organizational structures, roles, information systems, and operating environment.

Chapter 6: Understanding Routine Work. This chapter introduces the routine work practices at the site. Through an example of a typical high priority call, the various processes, constraints, expectations, and dependencies are explained. The notion of information flow mapping is presented and this central example is diagrammed.

Chapter 7: Boundary Objects in Collaborative Problem Solving. This chapter acquaints the reader with the concept of boundary objects as presented in Chapter Two and adapts it for use in this study. The value of this construct is illustrated through the unpacking of another example of a resolution to a routine service request. A boundary object diagramming technique is also introduced.

Chapter 8: Coordination and Collaboration: Inter-group Boundary Spanning. The first of the core boundary object findings chapters, this examines in detail the mediating effects of a particular boundary object, the ROC, on inter-group coordination and collaboration. It also focuses on the mutual reinforcement of perceived authority between this artifact and the role of the team lead.

Chapter 9: Negotiation and Translation: Inter-organizational Boundary Spanning. This chapter examines the role of the ROC in mediating collaboration between the site and its customers. It focuses on the various forms of required negotiation across this boundary. It introduces situations where the boundary object concept provides inadequate theoretical coverage.

Chapter 10: Recontextualization and Reuse: Boundary Spanning Across Time. This chapter addresses the central behaviors of reuse by engaging the temporal characteristics of the ROC through an extended analysis of another routine action item resolution. Common breakdowns in the processes of decontextualization and recontextualization are highlighted, as are the organization's routine work arounds.

Chapter 11: Implications for Design. This chapter explores the design implications of these core findings. Central issues discussed are the value of designing for negotiation points, the role of lightweight preservation of context in aiding recontextualization, and the interplay between formal and informal information.

Chapter 12: Conclusions. This chapter revisits the key findings of the study. It addresses the unique stressors of high-reliability and high-liability, and posits future trajectories of research that are suggested by the outcome of this study. Lastly, it speculates on expanding the theoretic concept of the boundary object to more adequately describe the observed reuse behaviors.

CHAPTER 2: Knowledge Work and Collaborative Problem Solving

“Knowledge-intensive work processes often concern collaborative problem-solving...”
(Carstensen and Snis, 1996, p. 171)

2.0 Introducing Computer-Supported Cooperative Work

Given the commonality of reuse in organizations, it is not surprising that the behavior has been studied by numerous academic disciplines, such as computer science, information science, management, and organizational behavior. They have examined domains as diverse as archival concerns in library science, executive information systems in management, and the design rationale capture problem in software engineering. Clearly not all of these literatures can be surveyed here, so the one that is of greatest assistance, computer-supported cooperative work (CSCW), will be summarized.⁴

CSCW is multi-disciplinary research community drawn together by a shared interest in building a deep understanding of the work practices of small groups, with the intent to inform the design of socio-technical systems to better support those groups.

Methodologies are largely drawn from the social sciences (notably sociology, anthropology and psychology), and the focus is resoundingly on human actors. When

⁴ Within the CSCW research community, knowledge management issues are generally covered by the term “organizational memory.” For an introduction to this notion and an excellent survey of the term’s use throughout multiple disciplines, please refer to Walsh and Ungson (1991). For an update of some core organizational memory themes in CSCW, please see Bannon and Kuutti (1996).

examining reuse, CSCW affords a healthy shift away from the computational, artifact-centered viewpoints of the other dominant disciplines, and towards people and processes.

2.1 Lessons from Computer-Supported Cooperative Work

Over the last decade, CSCW researchers have developed a core set of findings regarding the design and use of collaborative technologies.⁵ Some of the ‘lessons-learned,’ most relevant to this study, will be summarized here.

1.) Knowledge work is a highly situated phenomenon

2.) Exception handling is the norm in routine work

Suchman’s early work (1983;1987) uncovered the underlying problem with attempts to automate knowledge work (e.g. expert systems). Knowledge work is inherently highly situated and exception driven, two cognitive elements that are notorious limitations for computing technologies. With this understanding, design interest within CSCW has shifted away from automating these tasks and toward building better support tools to augment human capabilities.

3.) Awareness of the activities of others is valuable

There are a host of studies within CSCW examining the role of awareness in routine work. Most are focused on the impact of specific communication technologies, especially when supporting distance collaboration. Others focus on

⁵ For a summary of the key sociological findings please refer to Ackerman (2000) and for core high-level design criteria, see Grudin (1988;1994).

awareness behaviors in natural settings, such as (Bently et al., 1992) and (Heath and Luff, 1992).

4.) Co-evolutionary development cycle of socio-technical systems

One of the foundational findings in CSCW is that all systems have social and technical elements, which co-evolve over time via a complex feedback loop.

Consciousness of this allows developers to address the interplay of the social and the technical components in iterative or participatory design (O'Day et al., 1996).

5.) Critical mass is an elemental criterion for the adoption of groupware systems

There is a utility threshold with groupware systems. If they do not achieve a consistent critical mass of users, then the utility plummets for those who are using the system (Grudin, 1988;1994).

6.) Proper alignment with incentives structures is critical

Systems that are misaligned with an organization's culture and incentive structure are prone to resistance and problematic adoption. (Orlikowski, 1992; Grudin, 1988;1994)

7.) Workplaces are socially heterogeneous; ignoring this fact in development is costly

Workplaces are diverse, socially complex, and politically charged environments.

Often workgroups contain members from multiple communities of practice

(Wenger, 1998) and as such have different goals, mental models, norms and even

vocabulary. Systems which do not align with particular user groups will likely face resistance (Markus, 1983), while those which support multiple viewpoints and translation are more likely to enjoy successful adoption (Schmidt and Simone, 1996).

Each of these themes plays out in this research program presented in this dissertation and will be revisited in context.

When examining the use and reuse of information in collaborative environments, the CSCW literature has been bifurcated between studies which emphasize the mediating effects of technical artifacts and those that focus on the mediating effects of people.

While there are very few studies that blindly adhere to either perspective, and a growing corpus of work exists which actively seeks to integrate them, for the purposes of this review the historical split is a valuable framing device. Each will be addressed in turn below.

2.2 Social use of information as mediated by technical artifacts

There have been a number of studies over the past decade that examine the use of information as mediated by technical artifacts. For example, Harper in his ethnography of the International Monetary Fund (IMF) focused almost exclusively on the critical role documents play in structuring routine policy work (1998). At the IMF, documents are the physical instantiation of organizational processes and perspectives. Through examining the creation of one particular staff report for an international mission, Harper details the

transformations required in the evolution of the final report. His work aptly concretizes the exception handling and post-hoc rationalization described by Suchman above, and the issues of context preservation and interpretation presented in the following section.

Other studies in CSCW have examined comparable roles played by information artifacts. Both Bently et al. (1992) and Mackay (1999) have examined the role that paper flight strips play in coordinating the work of regional air traffic controllers, specifically focusing on the strips' role in directing the controller's attention during complex cognitive tasks. By actively engaging the controller with the information on the slips and by forcing him to physically interact with the environment and his colleagues, the controllers have developed a system with remarkable error catching capabilities. Heath and Luff found similar behaviors with rail controllers annotating paper schedules in conjunction with observations on a shared control representation screen and radio links with rail operators (1991;1992).

Carstensen and Snis (1996), in their field study of knowledge workers in a quality control department of a Danish pharmaceutical company, move beyond examining collaboration to engage issues of knowledge exchange. They closely studied the use of two organizational memory systems in support of reuse behavior, noting the interplay between the formal and informal. They have three core findings relevant to this work.

The first is their observation that “the artifacts used [in organizational memory systems] are seen as the knowledge-mediating mechanisms for explicit knowledge exchange.” (p.

177). Of Nonaka's four modes of interaction (1995), these artifacts most clearly support the processes of combination, where "one piece of information is combined with another" (Carstensen and Snis, 1996, p. 176) in the continued creation and growth of organizational knowledge. These artifacts form an intersection point between different knowledge sources.

Second, Carstensen and Snis found that by leveraging the information byproducts of routine knowledge work in an organizational memory system (e.g. archives of report drafts or approved government documents), employees were more likely to engage in knowledge transfer. Considering the problems that most knowledge management projects have regarding adoption, this is a significant finding. However, they do not elaborate on this finding. Areas of clear interest and extension include understanding how documents were valued, how selection criteria were generated, how documents were used, and most importantly, how they were modified from their intended use to also support reuse.

Lastly, Carstensen and Snis explicitly engage the issues of systems design to better improve these knowledge management processes. While their nine requirements are very general and high-level, they are a valuable starting place for continuing research such as this dissertation. Few ethnographies of organizational memory systems continue on to this next stage, feeding their findings back into design principles. This has been a major criticism of similar ethnographically informed studies (p. 172).

Further research examining related behaviors will be presented within the context of discussing boundary objects later in this chapter.

2.2.1 Processes of Contextualization

A central theme running through most of these studies is the role of context in interpreting information. Simply put, information is viewed differently based on contextual cues, whether it is a flight strip, a speech draft, or a mechanics log entry. Recognizing the role of context in shared information is critical to interpreting its potential for reuse (as is appreciating the knowledge required to make these evaluations). Two early studies develop these issues of context.

Confronted with the observation that colleagues in his research lab successfully used two different shared calendaring systems simultaneously, Dourish et al. (1993) set about ferreting out the underlying rationale for this apparently redundant behavior. The obvious observation from a human-computer interaction (HCI) perspective, that these were merely two different interfaces to the same data tuned appropriately to the task at hand, did not satisfactorily explain the puzzle. The deeper reason involved context. One system was designed for a specific known audience, personified by a human caretaker, routinely archived, supported a long-range view of activities, and ingrained in the lab's work practices (e.g. annual reporting). It was best suited for broadcasting formal activities. The second system had a more restricted audience, distributed ownership (though each record had an identification tag), non-archival status, focus on immediate activities, and

decidedly protean and ephemeral qualities. It was best suited for maintaining informal entries or tracking tentative activities.

From observing the use of these two calendaring systems, specifically how individuals evaluated calendar entries, the researchers deduced a set of contextual components: ownership and responsibility, medium and mutability, timelines, and organizational status and relevance. The interactions among these components are complex, but not so complex as to preclude some general guidelines for designing for interpretation, namely by better supporting the integration of contextual cues into the information artifacts themselves.

In his evaluation of six organizations using an experimental organizational memory system, Ackerman (1993;1996) observed similar behavior regarding the role of contextual cues. Instead of focusing on the affordances of these cues, however, he emphasized the constraints they placed on the interpretability of information. A key additional contextual cue he observed was the importance of the expected trajectory of the information. Different from only understanding the intended audience, this ties the information to the organizational processes and goals that spawned it.

Ackerman's central contribution though was in understanding the temporal nature of contextualized information. He focused more heavily on the processes surrounding the maintenance of context rather than on the results of reuse. In doing so he outlined two core processes: decontextualization and recontextualization.

When storing information for later reuse, clearly not all context can or should be preserved. In addition to preserving details about authorship and intended audience, work is done to make the information more generalizable, more accurate, or better tailored for a different trajectory of expected reuse. In describing this process of decontextualization, Ackerman notes that managing the appropriate level of context can be tricky business: “Too much contextual information results in a sea of extraneous detail and hinders generalization of the information. Too little context results in the users’ being unable to understand the information and the authors’ formalization of their material” (1993, p. 8).

In balance, recontextualization is the process by which context is reconstituted from these preserved cues in order to evaluate the appropriateness of the information to the task at hand. Both of these processes are complex, but critical. Poorly preserved information is useless as it cannot be successfully adapted to later information needs. This complexity increases dramatically when information crosses organizational boundaries.

It is specifically in its ability to address this complexity that the theoretical concept of a boundary object is so useful. It affords the capability to track the context requirements and interpretive perspectives of multiple groups engaged with a common artifact. The concept will be presented in the following section and elaborated through a set of relevant CSCW studies.

2.2.2 Boundary Objects

The concept of boundary objects was originally presented within the artificial intelligence research community to better understand distributed decision making (Star, 1989).

However, Star and Griesemer quickly applied the concept to social settings in their historical study of heterogeneous scientific work at Berkeley's Museum of Vertebrate Zoology (1989). In their canonical definition, they present boundary objects as being:

“objects which both inhabit several intersecting social worlds and satisfy the informational requirements of each of them. Boundary objects are objects which as plastic enough to adapt to local needs and the constraints of the several parties employing them, yet robust enough to maintain a common identity across sites. They are weakly structured in common use, and become strongly structured in individual use. These objects may be abstract or concrete. They have different meanings in different social worlds but their structure is common enough to more than one world to make them recognizable, a means of translation. The creation and management of boundary objects is a key process in developing and maintaining coherence across intersecting social worlds.” (p. 393)

Star and Griesemer suggest a preliminary taxonomy of boundary objects based on their role in translation. These four initial classes are: repositories (“ordered ‘piles’ of objects which are indexed in a standardized fashion”), ideal types (“an object... which in fact does not accurately describe the details of any one locality or thing. It is abstracted from all domains and may be fairly vague. However, it is adaptable to a local site precisely because it is fairly vague; it serves as a means of communicating and cooperating symbolically.”), coincident boundaries (“common objects which have the same boundaries but different internal contents”), and standardized forms (“objects devised as methods of common communication across dispersed work groups”) (pp. 410-411).

As the concept of boundary objects has matured in the work of Star and Bowker, they have focused more on their role in “developing and maintaining coherence across intersecting social worlds” as described above. As is logical then, they have tended to investigate objects primarily of the fourth class, specifically examining boundary objects that mediate at the institutional level: classification schemes and standards development (Bower and Star, 1999).

However, others have taken the concept and developed it in the opposite direction, focusing on specific instantiations of information artifacts that mediate a set of very local relationships. Berg and Bowker have done interesting work with patient medical records in hospitals (1997). Working in a high-reliability, high-liability environment they uncovered a number of special constraints which are unique to this environment such as: the need for a legal paper trail, management of formal and informal records (the creations of “accounts”), and objects inscribed with work flow information. While they also engaged the notion of time with regard to boundary objects, they did not address the concomitant issues of life cycles, maintenance, or reuse.

Another highly relevant application of the boundary object concept has engaged similar issues in a very different setting. As part of the GMD’s POLITeam project with the German federal government, Mambrey and Robinson examined specific instances of documents as boundary objects (1997). The core of their paper is a detailed trace of a single document, a draft of a speech for a government minister, throughout its information career. Aware that boundary objects tend to experience perturbations when

they cross boundaries, they followed this speech throughout the routine information flows in the ministry's office, noting the alterations which occurred at the borders. Mambry and Robinson then developed an understanding of these border changes and discussed the emergent dynamic of work flows as represented by these objects.

Their approach is provocative and promising. However, as their focus is primarily on understanding work flows and how to design systems to better support them, boundary objects were only a tool used to highlight critical junctures in the organization's information flows. They do not return to further develop the concept of boundary objects.

Kovalaninen, Robinson, and Auramaki (1996;1998;2000) also examined boundary objects in a different setting, among technicians servicing a paper mill. In this impoverished communication environment (deafening machinery, low staffed shifts), the technicians found value in communicating via informal electronic diary entries. These entries served as boundary objects across time, bridging understanding across shifts. Eventually they also were opened up to others outside of their technical team, namely to management. While this study brings to light the success these diaries had in crossing barriers of time, the authors also do not engage the standard temporal issues regarding information: decontextualization, recontextualization, maintenance, decay, and obsolescence.

Others studies have focused the role boundary objects play in developing shared understanding for a particular problem. Henderson's work with design engineers (1991a,b,c;1998a,b) centered on appreciating how engineers use diagrams, drawings, and blueprints as points of negotiation. She focused specifically on the changes, both positive and negative, occurring as the CAD revolution shifted these artifacts from paper to digital form. Bechky (1999) also attended to the role that drawings play in negotiations among engineers. However, she focused on drawings that explicitly span social world boundaries (e.g. moving from design to manufacturing).

Lastly, Kim (2000) has advanced the concept of boundary objects beyond that of specific information artifacts and toward boundary instances. In his study of DRAM fabrication engineers, he found points of negotiation across community-of-practice boundaries not only at the expected places (e.g. a particular machine or technical specification) but also around challenging problems. It was the problem itself that became the locus for negotiation: its definition, its meaning, and its resolution. A limitation of this study though is that, while valuable, the concept of a boundary instance is relatively static and isolated within the organization's information flows.

In examining all of these boundary object studies, the next step in maturing the concept is clearly to examine boundary objects/instances in very heterogeneous groups, in constrained information environments, and across time (addressing all that this entails), while being sensitive to their trajectories of expected use and embeddedness in the complex information flows of the organization.

2.3 Social use of information as mediated by people

Turning now to the other category of CSCW perspectives on information use, the following studies focus more on the social factors that mediate information flows.

A standard example of this class of studies is that of Ciccourel (1990), who examined the processes of medical diagnosis in a major hospital as collaborative problem solving among attending physicians, medical residents, nurses, and laboratory technicians. The cognitive processes of diagnosis are understandably complex, however Ciccourel foregrounds the even more complex social practices. Few physicians have the expertise to render a correct diagnosis independently; instead they must rely on a web of information sources from those as simple as printed lab results to others of increasing complexity such as consultations with specialists. They navigate this web through a series of social interactions. The rationale for doing so, instead of querying a database for example, is that the decision maker must integrate these disparate information sources to form a coherent and accurate diagnosis, and the key element in doing so is understanding the credibility of each source. This routine, though complex, process of authenticating and evaluating information is partially embedded in the organizational structure of hospitals which have strict hierarchies, allowing the decision maker some safe latitude to assume that the higher in the hierarchy the source, the higher the degree of accuracy for the provided information.

The processes by which individuals, such as the attending physicians or medical residents of Ciccourel's study, seek out expertise in an organization are a core concern of this research literature stream. Next, a key concept in understanding these expertise networks will be presented.

2.3.1 Hubs in the knowledge network

In his retrospective analysis of studies of technological innovation in companies, Allen focuses on the communication patterns of engineers (1977). Tracing these formal and informal communication networks within the organization, he discovered that there are significant hubs in the network. These hubs are individuals who have a particularly extensive social network and rich understanding of information external to the organization. Basically, the individuals are well read and well connected, and as such routinely introduce new ideas to their colleagues. Allen coins the term "technological gatekeeper" for this organizational role and calls for greater appreciation of their critical contributions to the processes of knowledge creation and diffusion. While a valuable contribution to understanding knowledge networks, his analysis is limited by focusing solely on the gatekeeper's role in mediating information flows and by an overemphasis on their introducing information from domains external to the organization.

Tushman et al. addresses the latter point in his further development of the technological gatekeeper concept through a series of studies on research and development laboratories (1977;1980;1981). While still engaging the external mediation role, Tushman additionally highlights some gatekeeper's ability to navigate the information environment

within their organization. Those who could best leverage these internal relationships were dubbed Boundary Spanning Individuals (BSI). These individuals are often not the formal cross-departmental liaisons one would expect, but rather informal “communication stars” who are respected by their peers as reliable sources of internal information. He found that BSI have the ability to cross-fertilize ideas between groups as diverse as marketing, engineering, and scientific staff.

Tushman’s studies are hampered though by their sole focus on the communication of content, as was Allen’s. Clearly, these hub individuals are providing a role far more complex than simply sharing the latest article that they read with those gathered around the water cooler.

McDonald and Ackerman (1998; 2000) address this point and continue to develop the gatekeeper concept with their examination of behaviors of expertise finding among developers and technical support staff in a mid-sized, medical software development firm. Moving beyond the content of the information exchanges performed by these hub individuals, McDonald examines their meta-knowledge of these exchanges, specifically how they know the best referrals for particular content. The researchers identify this role as an “information concierge,” one who maintains a sophisticated map of ‘who knows what’ within the organization (2000, p. 130). The authors’ interest is in leveraging these individuals to gain a deeper understanding of the knowledge network within an organization in order to better design an automated expertise recommendation system.

While instrumental in identifying the varying roles these hub individuals perform within an organization's information flows, all of these studies do not address the transformative effects these individuals have on the information they are presenting. For example, what translations must a boundary spanner accomplish before information is understandable to his team? How is this accomplished? Or are all expertise recommendations from gatekeepers or information concierges viewed with the same authority?

2.3.2 Social mediation in technical support

Other researchers in CSCW have found technical support environments, such as McDonald's medical software firm, to be rich venues for examining the role of expertise in collaborative problem solving as mediated by social interaction. Three related studies will be presented below.

The first of these is Orr's ethnography of copier repair technicians (1996). His study focuses on the expertise of technicians in collaborative problem solving, specifically how the technicians leverage information from the customer, other technicians, and the machine itself to produce successful resolutions to service requests. While this study has many strong parallels to the one documented in this dissertation there are interesting points of divergence.

Knowledge in this workplace clearly resides in the heads of the technicians. It is routinely shared through the telling of war stories over informal breakfast meetings within the context of reviewing the day's job assignments. These meetings are the critical

knowledge diffusion activity in the organization, yet they have a very low visibility and level of support from higher management whose focus is on endlessly improving manuals (“directive documentation”) in the hopes of controlling organizational knowledge and allowing for a more deskilled workforce in the future. Although an all too common experience, such attempts at fully automating technical support are futile, as Orr explains. Building on Suchman’s claims regarding the impossibility of designing formal plans to cover all aspects of highly contingent knowledge work, Orr adds that “knowledge relevant to the job of diagnosis cannot be precisely defined.”(p. 107) However, as Suchman also claimed, these efforts at collecting and formalizing knowledge are not without merit. While they may not supplant the technician, they are valuable in augmenting her skills. The technicians themselves are aware of this, Orr notes that “they grant the documentation some utility while denying it complete credibility...saying that the machine is far too complex to anticipate correctly all of its possible failures.” (p. 111) The documentation is useful when they are encountering a problem for the first time or are working on a machine they are unfamiliar with, but it is only used as a starting place. They know that the solution is far more complex and nuanced than can ever be represented in the documentation alone.

The core of Orr’s analysis focuses on just these processes of diagnosis that happen outside of the formal procedures. He likens the activity of diagnosis to *bricolage*, “the technician’s task in diagnosis is to create a representation of the problematic situation that is sufficiently complete to indicate a course of repair” (pp. 114-115) However, the information they have with which to build this representation is sparse: a confused and

contradictory problem report from the customer, a few abnormal clicks and whirs from the machine, and a lingering story from a colleague who had seen something similar once. The assemblage of such scattered cues by an expert to form a coherent diagnosis is truly the work of a *bricoleur*.

Orr's technicians are savvy enough to know that a repair does not end with a reconstituted machine, it continues with training the customer. Their mantra is "don't fix the machine, fix the customer!" (p. 79). Whether using inappropriate paper, toner or supplies, misunderstanding features, or losing staples in the sheet feeder, sharing some of the organizational memories from the repair organization with the customer benefits both parties – reducing the frequency of malfunctions for the customer, and easing the technician's future interactions with that customer. (These efforts at "managing the customer" are routinely centered on the designated contact within the customer's organization, the key operator. This individual not only has personal responsibility for the machine, but also has a higher degree of understanding about its workings. This increases the degree of shared understanding with the technicians.)

A key limitation of Orr's study is the high degree of personal involvement by the copier repair technicians. This situation is actually quite rare in technical support work. A quick comparison between his field technicians and a common telephone support hotline operator brings these differences to light.

Table 1: Comparison of Technicians

Copier Repair Technician	Technical Support Hotline Operator
On-going relationship with customer	Sporadic interaction with customer
Frequently deals with customer's local expert – "key operator"	Regularly confronted by shifting customer contacts of unknown or unpredictable expertise level
Information about the device is highly reliable, being collected by self or local expert	Information quality varies, being collected by contact of undetermined skill
Performs diagnosis in a rich information environment, physically inspecting the device in its native context	Performs diagnosis in an impoverished information environment, only able to work with details provided by the customer
Repairs the machine directly, affording the ability to fine-tune solutions through feedback (esp. tactile, visual, auditory, olfactory)	Relies on the customer to implement the solution, has only one shot at a repair and has limited or no feedback about the repair process as it unfolds
Unlimited repair options	Repair options are constrained to what the customer can perform
If solution is not successful, can iterate, building upon knowledge of prior case	If solution is not successful, the returning customer will likely be assigned to a different service representative, whose knowledge of the prior case, if any, is limited
Relies on social capital to maintain relationship with customer, builds a history	Very limited ability to build social capital, most interactions are independent

In addition, Orr's copier repair technicians are dealing with a moderately complex mechanical device, while other technical support engineers must troubleshoot infinitely more complex and ambiguous systems, such as software. Lastly, while the required degree of accuracy and reliability for copier repair is high, these are frequently not mission critical systems. It is a costly nuisance for a copier machine to suffer additional downtime, while it is a catastrophe when a corporation's accounting system is on the blink.

The findings in Orr's study are robust, however his work clearly needs to be extended to engage more complex and meaningful work situations. The additional two studies address some of these issues, while others are the addressed by the research presented in this dissertation.

In the second study, Pentland examined similar knowledge work being done by service engineers on software support telephone hotlines for two different companies (1992).

Again, he found that real knowledge is dynamic and highly contextual, stating:

“To the extent that we adhere to a pragmatic theory of knowledge, it is essential that actions become the object of our investigations into organizational knowledge structures. The critical insight is to stop treating knowledge as a static entity that resides somewhere, like a book in a library, and start treating knowledge as an active, situated phenomenon.”
(p. 545)

Differing from the other similar studies, he used situations, not individuals or groups, as his unit of analysis, noting that “the actions of members are always shaped, to some degree, by the situation in which they find themselves.” (p. 529)

In examining the resolution of support calls, he found that “moves” were the critical component in the technician's routine knowledge work. He defined a move as the transformation process of a problem, from instantiation to resolution (where resolution need not be equated with a final solution). Such moves are constrained by three structural restraints: physical (the limitations of the environment), ritual (social or cultural limitations), and competence (the limitations of the organization's distributed knowledge, both explicit and tacit). His resulting taxonomy is quite useful and its generalizability to technical support setting is borne out by this dissertation.

As a quick aside, Katzenberg et al. (1996) had examined related move behaviors among nurses in a breast cancer clinic. They examined the quasi-formalization of move rules, along with more standard work action scripts, into protocols which were mediated by a workflow management system. While their implementation at the clinic was a successful development and adoption story, they highlight the tension between the value to automate and the need to maintain significant flexibility, sometimes on a case-by-case basis, for their boundary object-like protocols.

The final study, Ackerman and Halverson's examination of technical support staff at a human resources telephone hotline for a large software development corporation, specifically addresses the nature of Pentland's call to view "knowledge as an active, situated phenomenon" (1998;1999;2000). In response to prevalent assumptions of organizational memory as a monolithic and static entity, they show that organizational memories are instead exceedingly diverse, distributed, dynamic, and often ephemeral. They examined the staff's routine use of the heterogeneous organizational memories for the successful resolution of support calls, highlighting both the diversity of these information sources and the relative ease with which the representatives navigated this complex information space. The core of their analysis is a deconstruction of a routine service call, revealing the use of "many small memories," reliance on a plethora of resources, and exercises in subtle forms of reuse.

Ackerman and Halverson's study extended many of the issues previously presented in this chapter: knowledge does not exist in a database, reuse is situated, memories can be

both objects and processes, and that there are elements of trust and authority in evaluating material for reuse. They expand on this latter point to introduce the notion of memories as boundary objects, stating that “as memory crosses between groups or even across time, it becomes a boundary object, attempting to serve the needs of both creator and reader but lacking the full context of either.” (1999) These objects are often the point of negotiation both intra-organizationally and inter-organizationally. They purport that the central problem in reuse is the recontextualization of these boundary object mediated memories.

Their approach to framing and exploring this central problem is the most promising in the literature to date and as such provides the foundation for this dissertation research. The following chapter will briefly revisit the knowledge management problems presented in the introduction in light of these CSCW findings. It will specifically address the appropriateness of applying the boundary object concept to these knowledge management issues as done by Ackerman and Halverson.

CHAPTER 3: Boundary Objects and Reuse

3.0 CSCW and Knowledge Management

The preceding review of relevant literature from the field of computer-supported cooperative work indirectly addresses many of the deep concerns in knowledge management projects, as presented in the introductory chapter. While those important problems provide the motivation for this research, it is clear that CSCW offers the most appropriate theoretic foundation.

Many of the CSCW concepts already presented provide the intellectual underpinning for this dissertation and will resurface as appropriate throughout the document. However, one central concept, that of boundary objects, requires further attention. It will be directly addressed in relation to the motivating knowledge management problems.

3.1 Suitability of boundary objects

Recall Fahey and Prusak's (1998) critique of the conceptual foundation for most knowledge management projects, their "eleven deadliest sins." These were then distilled down to the four most relevant, fundamental conceptual flaws. Each of these will be briefly readdressed here with regard to the potential contribution of applying the boundary object theoretic concept to their analysis and eventual redesign.

- "Emphasizing knowledge stock to the detriment of knowledge flow" – The boundary object concept only partially addresses this critique. It still suffers from an artifact-

centric view of knowledge, rather than a more process oriented perspective.⁶

However, the boundary object concept is a sizeable step in the right direction when compared to database or rule-based systems. With boundary objects there is at least an awareness of flows in which they are embedded and potential support for encapsulating informal, dynamic, and exception-based information. They also provide for multiple perspectives on this shared information.

- “Not understanding that a fundamental intermediate purpose of managing knowledge is to create shared context” – The ability to answer this critique is one of the primary strengths of the boundary object concept, where every object is viewed as the result of a complex negotiation process among multiple parties. It is around these common negotiation points that shared understanding is built among often very diverse participants. Recall some of the boundary object examples from the literature presented in Chapter Two: patient medical records, speech drafts, electronic diaries, and engineering drawings. Each of these acted as a locus around which individuals from different communities of practice negotiated to a shared understanding about the patient, the speech, the paper press machinery, and the items to be designed and manufactured, respectively. These are precisely the types of information artifacts that should be understood and managed in any knowledge management effort. Finally, the boundary object concept makes explicit the varied stakeholders in the information artifact, allowing easier identification of and support for those diverse individuals.

⁶ This shortcoming will be explicitly addressed in this dissertation, pinpointing where the concept breaks down and hypothesizing about how to expand the notion to include a greater sense of flow within the knowledge network.

This is a critical understanding for designing systems to support knowledge management.

- “Disentangling knowledge from its uses” – As highlighted in the examples recalled for the preceding point, the boundary object concept aptly represents the diverse users of the captured information. As the negotiation point around which shared understanding is developed, the “information fingerprints” of the participants cover the boundary object. There is little danger of over-abstracting the information artifact, for once it is completely removed from its context it no longer functions as a boundary object.
- “Focusing on the past and the present and not the future” – There is nothing about the boundary object concept which explicitly engages temporal concerns. Although, the fact that boundary objects are highly situated does assist in ascertaining likely future audiences and scenarios for reuse. Ackerman and Halverson have successfully adapted the boundary object concept to address issues of organizational memory. One of the primary contributions of this dissertation research will be to continue to develop this concept in this direction, noting both where it succeeds and where it falters.

Routine boundary objects may exhibit different traits, with some of the most common in knowledge work oscillating between ideal and standardized.

3.2 Addressing the alternatives: coordination mechanisms

Seeing that the boundary object concept is a good, but not perfect vehicle with which to address the core conceptual concerns in knowledge management, raises the question of the existence of other, more appropriate, theoretic concepts. In all regards there are very few other concepts that are as robust as boundary objects, however there is one suitable candidate for competition: coordination mechanisms.

Remember the Carstensen and Snis (1996) study of Danish pharmaceutical quality assurance workers and the archival documents they leveraged in routine reuse. To understand the mediating roles that these organizational memory artifacts played in the knowledge exchange, the researchers found the greatest explanatory power in the theoretic concept of coordination mechanisms. Schmidt and Simone define the concept as follows:

“A coordination mechanism is a construct consisting of a coordinative protocol (an integrated set of procedures and conventions stipulating the articulation of interdependent distributed activities) on the one hand and on the other hand an artifact (a permanent symbolic construct) in which the protocol is objectified.” (1996, pp. 165-166)

The concept of a coordination mechanism is a valuable analytic construct for understanding the mediating role of certain information artifacts. However, in practice few artifacts used in routine knowledge exchange exhibit explicit coordinative protocols. While the protocols may be clearly defined for certain relations, they are often weakly defined or completely non-existent for others. (Recall here Suchman’s admonition to designers that exceptions are the norm in routine work.) As good CSCW researchers,

Schmidt and Simone do acknowledge that the protocols for coordination mechanisms are “inexorably under-specified” and that users will deviate appropriately to fit their specific context. By their very nature, coordination mechanisms mediate among multiple individuals, and thus this deviation and adaptation must involve the participation of multiple parties, not a single actor. However, they provide limited theoretic support for these requisite, on-going, multi-point negotiations. A logical extension of the coordination mechanism concept would be to better address this support. (The boundary object concept is reasonably suited to represent these exception handling negotiations.)

The coordination mechanism concept is rigorously defined and provides significant explanatory power for settings that are highly routinized. The boundary object concept is more loosely defined, and therefore affords less analytical power. However, it provides a much more realistic fit to the fuzzy reality of routine knowledge exchange.⁷ Preserving this ability to represent loose, dynamic, multi-point negotiation, while developing it in the direction of coordination mechanisms would yield a more valuable analytic concept in the end.

3.3 Research questions

Having settled on the concept of boundary objects as the most appropriate theoretic fit for the task at hand, the core knowledge management concerns can be revisited. Defined in

⁷ The boundary object concept also provides the best fit to the data collected in this study, suitably explaining and predicting the negotiations and behaviors surrounding the core information artifact.

terms of boundary objects, Fahey and Prusak's critiques can be distilled to the following two high-level research questions:

- What is the role of boundary spanning information artifacts (boundary objects) in knowledge reuse?
- Given the local context, how can information technologies be designed to best support the mediating effects of these artifacts?

Examined within the context of my research site (to be introduced in Chapter Five), these questions are addressed by the following three grounded research questions:

- How do service engineers properly distill information for storage and later use? (processes of decontextualization)
- How is that information later successfully discovered, reconstituted, and adapted by service engineers to fit a problem at hand? (processes of recontextualization)
- How do service engineers navigate the knowledge network of their organization to find the relevant expertise to perform these tasks?

As introduced in Chapter One and re-engaged here, the promise and potential of knowledge management is great. However, most efforts have routinely met daunting

challenges. It is the intent of this study to address these challenges, as defined by the aforementioned research questions, through a qualitative field study of a technical support venue. It is to be informed by the CSCW perspective as presented in Chapter Two. The details of site selection and study design then are offered in the following chapter.

CHAPTER 4: Study Design / Methodology

4.0 Research approach

A study examining the role of reuse in collaborative problem solving is most amenable to employing a qualitative methodology. An interpretivist stance, utilizing qualitative data collection techniques and methods of analysis, is appropriate when exploring or uncovering new or not well-understood phenomena. Often these are situations where the phenomena are complex enough to inhibit reduction to a quantitatively tractable level. For knowledge reuse, there has been little descriptive work done on the behavior itself and the routine practice is fraught with complex interdependencies, rendering the identification of relevant independent and dependent variables deeply problematic.

In order to garner the deepest possible understanding of reuse behavior, with a focus on the technological and social mediating effects, the phenomena ought to be explored *in situ*. Qualitative fieldwork has a number of diverse disciplinary derivatives, however those utilizing ethnographic perspectives and techniques have proved to be very successful at illuminating complex social behaviors. Thus, an ethnographically informed field study is a wise methodological selection for this project.

It is important to note that in this type of research, conclusions emerge from the analysis of the data. Therefore, to have confidence in the ultimate findings of the study it is critical not only to appreciate the general study design but also to understand both the

techniques used to collect the data and the theoretical frameworks employed to interpret it. Each of these will be discussed in detail following a brief discussion of the field site selection process.

4.1 Site selection

In order to address the core research goals of this study, a number of criteria were critical to the site selection process. Each of these will be discussed in detail below. It is also honest to admit that a certain degree of serendipity is present in all site selection processes – being in the right place, at the right time, knowing the right individuals. Such was the case here as well.

4.1.1 Technical support venues

The choice for a general domain of study was strongly influenced by the existing literature. As outlined in Chapter Two, a number of recent studies within CSCW have focused on technical support settings as rich venues for exploring the processes of social mediation in reuse. Most notably, in the Ackerman and Halverson organizational memory study (1998;1999;2000) they were able to successfully examine both social and technical elements in the distributed cognition system underlying the collaborative problem solving behavior in these support environments. They identified some of the beneficial characteristics of technical support environments as follows. These situations have sufficient routine work to map their information processes; yet, there are always new questions and problems. Technical support work is time critical and extremely information intensive. While task responsibility is primarily personal, the routine work is

quite collaborative, often spanning intra-organizational and inter-organizational boundaries. Lastly, support groups are often lead users of collaborative information access technologies.

In selecting a technical support field site, the study in this dissertation benefits not only from these factors but also from building upon the intellectual foundation provided by these prior studies. This work also extends these studies by examining meaningful perturbations of these venues.

4.1.2 Extreme environments

When exploring understudied, complex phenomena, it is often valuable to examine it first in an extreme environment.⁸ In such a setting core attributes of the phenomena are magnified, constraints are clear, and limitations are more obvious. This often helps expose the underlying structures of the phenomena which are common across many diverse environments. Efforts to translate findings to more standard settings are not extremely difficult.

Considering the role of reuse in collaborative problem solving, two environmental stressors on an organization are of particular interest: high-reliability and high-liability.

⁸ One example has been the design of environmental systems for extreme environments such as space or deep sea exploration. Lessons learned from these settings have increased our understanding of closed systems and influenced the design of terrestrial technologies in domains as diverse as medicine, food processing and clean room manufacturing.

Each of these perturbations of a normal environment toward the extreme are briefly introduced below.

4.1.3 High-reliability organizations

High-Reliability Organizations (HRO) are organizations with no tolerance for error, where even the slightest mistake can have catastrophic consequences. As a result of their high-stakes operations, HRO's tend to have multiple layers of redundancy designed into their organizational structure and procedures. They are also often regulated and policed by external organizations out of concern for public safety. Weick most notably applied the HRO term in his examination of coordination efforts on aircraft carrier flight decks (1993). Other classic examples include control rooms for nuclear power plants, rail yards, and air traffic control centers.

4.1.4 High-liability organizations

High-Liability Organizations (HLO) are organizations which have exceptional legal responsibility for their actions. Work in this environment requires careful documentation of all processes, often routinely reporting to a regulatory agency. HLO's must be able to account for all decisions made and trace individual responsibility throughout all key processes. Common examples of HLO's are pharmaceutical R&D labs, hospitals, and the legal system itself. High-liability organizations are often organizations which require high-reliability operations as well.

4.1.5 Global Airframe: Technical Support Western Division

The site selected for the study, Global Airframe's Technical Support Western Division (GTS-West), met each of the desired criteria in varying degrees. It was immediately obvious that this service engineering group was an excellent fit for the research interests. As a technical support department, the goals at GTS-West were clear. They had a high requirement for reliability, but with growing requirements for turnaround time and cost control. Furthermore, this situation existed within a complex regulatory and legal liability web, which must be accounted for in the information processes. Finally, there was a constant concern about the public visibility of mistakes. This extreme information environment, with its time and reliability pressures, would likely bring to light key understandings about reuse behavior in collaborative problem solving situations. This would set the stage for some measurable contributions as no one before has examined reuse behavior in such an environment.

Chapter Five provides an extensive organizational overview for Global and GTS-West which may be helpful in placing the study details in context. The following discussion of the study design will assume a degree of familiarity with the site itself.

4.2 Study Design

The study of reuse behavior at GTS-West was, at its core, a thirteen month qualitative field study which occurred in three distinct phases. The transitions between phases allowed for reflective time and realignment of the study.

The first phase (4.5 months) comprised the core of the participant observation (Jorgensen, 1989). I sat with the Stress supervisors at the physical boundary between the Stress and Structures groups for the majority of the time and rotated among the other engineers for the remainder. The focus of this period was gaining an understanding of work practices and mapping the key information flows.

The second phase (6 months) centered on preliminary data analysis, requiring reduced field time. However, during this period I actively observed a business process re-engineering team, comprised of representatives from throughout the department, through their 25-week effort. This provided multiple perspectives on the same information processes.

The third phase (1.5 months) consisted of systematic data collection, examining some of the key information artifacts as boundary objects.

Throughout all phases, I attended regular staff meetings for the two groups, and interviewed a majority of their teams. I have also augmented these observations with visits to other GTS offices, interviews at a major airline repair facility, and the evaluation of considerable secondary material.

A detailed discussion of the actual study design follows. It is structured by purpose, placing the discussion of the data collection techniques employed within the context of

understanding the service engineers, the information artifacts in their world, and their work itself.

4.3 Data collection techniques

The primary set of data collection techniques utilized in this study are ethnographically informed (Emerson, et al., 1995; Lofland and Lofland, 1995; Van Maanen, 1988; Wolcott, 1990). Specifically, these are techniques which have been successfully adapted from the social sciences by CSCW researchers to address issues unique to the study of information technology use in organizational settings.

4.3.1 Understanding the People

In all field work, building a foundation of rapport, acceptance and trust with the people you are intent on understanding, in order to get the insider's perspective, is a long and difficult process. Harper, in his study of the IMF (1998), describes this process as a series of cascading ritual inductions. Far from having to "go native," an anthropological extreme, these are more akin to local rites of passage required to gain entrance to a social group.

For the ethnographer in Micronesia, these rituals may be very public, exacting moments. For the ethnographer of the white-collar office, they are often much less so. Acceptance here usually comes with tight co-presence over an extended period of time, with the researcher riding the rhythm of daily office experience together with them. I had truly outstanding access at GTS-West. In addition to their generosity and cooperation, this

resulted from a mixture of these prolonged interactions and important transformational incidents.

- **Presence** – During the first phase of the project, I was in the office regularly, five days a week seated at a high visibility workplace. I was there for the first bagel run of the morning, scurried with everyone for our forty-two minute lunch, and there at the frantic close of the day. By the end of the four and a half months my regular, routine presence was a powerful indication that I was “one of the crew.”

- **Longevity** – This is an issue related to presence. The service engineering team at GTS-West has had plenty of experience with consultants, most of it quite grim. There was a high expectation that I would blow in for a few weeks, poke around and leave declaring that I knew how everything worked and that I knew precisely how to fix it (and who to fire). The early incredulous questions of “are you still here?,” eventually gave way to viewing me as part of the team. This even went to the point of having a full week of competing “retirement” lunches at the end of the study as different teams claimed me as their own.

- **Shared passions** – Everyone understood that I was not a structural engineer. In the early days I practically reeked of academia. That was excusable. In time, if it was worth their effort, they could help me start to understand the world as a structural engineer would. What was not excusable was a lack of passion for the industry. In their mind, building and maintaining aircraft is one of the greatest engineering vocations on the

planet. If I didn't love "the birds" then I really had no part being here, I would clearly be an outsider. So, to the best of my ability, I developed the aerospace fanatic within me and it worked well enough. I passed as an enthusiastic novice. This stance afforded other unexpected data collection opportunities such as behind-the-scenes plant tours and interactions with groups outside of the department.

• **Interest in the mundane** – One of the most important keys to acceptance was that I was interested in what really went on around the office. I was interested in filing cabinets, pet software projects, and workspace organization. I wanted to know history, the old practices, the great heroes and villains of the department. I thought that service engineering was actually interesting; no one else at the plant would ever think that. And to the last day people still expressed surprise that I was excited by weekly staff meetings. The more interest I showed in what really mattered to the engineers in their daily experience (versus what might interest Global's high-level management), the more people were willing to open up and describe the intricacies of their jobs.

• **Rites of Passage** – Beyond sharing the common trials of the engineers trying to work in a decaying manufacturing building (which had been razed by the time the dissertation was complete), there were three distinct experiences that were turning points in my crossing the boundary to be accepted as an insider. The first was my willingness to file. Filing is the most despicable of all chores in the department. The engineer's work is very paper intensive. Yet their exacting schedules keep them from managing it well, so it buries them. Many engineers lived with stacks of paper taller than they were surrounding

their desks. Early on I struck a bargain, that I would help anyone with a few hours of filing in exchange for them telling me about what I was to file. (I had to walk a very fine line between this activity and being stigmatized as the new filing clerk. Due to budget cuts GTS-West had lost their filing staff and it was sorely missed.) While this was a very worthwhile endeavor in and of itself, it was the one evening when I stayed long after closing time helping an employee clean out his cubicle on his last day that helped me cross the insider/outside boundary in the department. This action was viewed as showing support for a fellow co-worker (many others had lent a hand during his final week) and concern for the group's collective memory (the departing individual was Stress' lead expert on a number of important subsystems prompting anxiety over losing this information. By helping to integrate his filing system with the public system for the group, many of those fears were relieved).

The second rite of passage was developing a healthy fear of flying and getting over it. (I've imagined this is akin to the medical student's experience.) Engineers teased me regularly early on that once I understood what really went on behind the scenes, I would think twice about flying. Even the department's director joked my first week, "do you like to fly? Well, you won't." Sure enough, they were right. Phobias I never knew I had began creeping out of the woodwork. However, fear was not the marker of being an insider, overcoming that fear was. Between the first and second phases of the study, I had to take a transcontinental flight. When I returned a number of co-workers asked about the trip. Was my experience different? The fact that it was and that I would still fly again was verification enough that I could participate in this culture. There is a

quixotic tension in air travel between it being certain (the most tightly regulated mode of transportation in the nation with a statistically impressive safety record) and unpredictable (every one of the thousands of flights a day is a near miraculous confluence of a million different factors from weather, to air traffic control, to pilot skill, to the myriad complex systems on board, even down to the temperature induced metal fatigue of rivet number 6742R). To work in this world you need to understand that tension and internalize it. Ultimately people's lives are at stake with every decision you make, yet you are contributing to the safest means of long distance travel ever developed.

The final ritual involved "being screwed by the system." Consultants rarely have problems; things are taken care of for them, employees on the other hand usually sink in the mire of bureaucracy at one of the largest corporations in America. Co-workers took unfair delight in my routine travails with Global's bureaucracy. It took me two months to see a new employee's manual (everyone joked that that was pretty standard even for employees), three months to get a full complement of office furniture (my office space had been raided after the departure of its previous occupant), four months to get a phone and five months to get access to the information systems I was supposed to be studying (one of which I never was able access). The engineers enjoyed helping me out in my difficulty, usually while recounting stories of their own misfortune. One even let me access his computer accounts while he was at lunch or in meetings, until I could get an account of my own. All of this helped build a great sense of camaraderie but the ultimate incident occurred when I was moved.

GTS-West was a group accustomed to frequent, painful moves. They had been in practically every office complex within a five mile radius of their current location. When I found out suddenly that my office space was going to be given to a new hire (quite understandable) and that I would be generously given new office space, though in a different part of the building, many co-workers came to my rescue offering their desks for me to work at on days they would be out. For the duration of the study people would ask about my life “over there” with unrepentant glee.

One other experience helped in this regard, though with mixed results. To a much lesser extent my later participant observation of a business process reengineering team, signaled my commitment to the long term interests of the group, though it uncomfortably called into question my loyalties (to those on the project team my role as observer was clear, to those not on the team my motivations seemed questionable).

4.3.2 Understanding Information Flows and Artifacts

Another class of ethnographic data collection Harper discusses involves seeking to understand the information careers within an organization. He is looking at the life-cycle of information; “information [that] is marshaled, is worked up, reviewed, circulated, used, stored and then forgotten about. Information within organizations has, if you like, a birth, a life and a death.” (1998, p. 69) While my interest is explicitly on the reuse phase of that this lifecycle, I needed to gain a complete understanding of the full life cycle to understand the reuse in context. In this paper, I use the term *information flows*, an

analogous concept but one that emphasizes the collaborative, distributed processes of information use.

Harper notes that information may appear in multiple modalities during its life span and cautions against document-centrism on the part of the researcher. I was keenly aware of this at GTS-West as critical pieces of information would regularly shift from a piece of scratch paper, to a data base entry, to a person's head, to a formal legal document; adapting appropriately along the way to the context at hand. In general I will refer to information instantiated in a particular media as an *information artifact* and information artifacts which have significant mediating effects as *boundary objects* (instances of boundary events) in this paper.

How then to map information flows and capture the life span of information artifacts?

The obvious inaugural step was to interview the engineers, so I began the study with a round of introductory interviews on both information flows and artifacts which continued throughout the year. Other data collection activities brought additional depth to my understanding of these issues.

• **Open, contextual interviews** – The majority of the interviews conducted for this study were open-ended, unscheduled, and opportunistic. Over the course of the work day, I would often pull an engineer aside for a clarification or elaboration of an activity that had piqued my interest. Often these were concerning critical incidents (described below), but they were also frequently in response to something that had come up in a staff meeting,

the latest water cooler war story or an unanswered question from the analysis of my field notes the night before. Usually these were brief (10-20 minutes) and done right at the person's desk as not to interrupt their routine work. Understanding that the technical support work always takes precedence, I was interrupted regularly. Once the engineers learned that I understood this aspect of their jobs, they were very accommodating, often coming later in the day to find me to continue an interrupted conversation. Also, I eventually learned the cadence of the different groups and tried, whenever possible, to time my non-action item specific interactions during less hectic times of the day. Informal social situations, such as lunches with small groups of engineers, were also tremendously valuable.

• **Semi-structured interviews** – Due to the nature of the work in the department and the conditions of my access, I was only able to secure a limited number of formal interviews. By formal, I mean extended interactions, about an hour in length, away from an engineer's workspace (and interruption) loosely driven by an interview protocol I had prepared in advance. My usual goal with these types of interviews was to fully explicate a particularly important work practice. The questions were most often based on a composite of issues raised in earlier open, contextual interviews. Because we were able to leave the employee's workspace, usually relocating to a conference room, these were some of the few interactions that I was able to audio record. Lastly, due to the non-availability of individuals for extended interactions such as these, the majority of these interviews were with managers.

• **Artifact identification, collection and description** – As mentioned above, one of my early field activities was filing. I offered engineers assistance in clearing off their desks in exchange for a discussion of the items I was filing. This brought me up to speed very quickly on the common (and exceptional) information artifacts in the environment. This early introduction to the pervasive paper elements in the office would prove valuable to my later observation. After this initial push, I continued to probe engineers regarding novel information artifacts for the remainder of the study.

• **Capture of critical incidents** – The artifact survey gave me a descriptive understanding of the information resources in use, but not knowledge about their use. Also, as a known limitation of interview data, the responses I received were more attuned to the rationalized understanding of the proper uses of the artifacts, not the messy, complicated ways in which they were used every day. So to find the other side of the story, I observed critical incidents. In general, I defined a *critical incident* as any breakdown in the routine work practice whose resolution required the assistance of a team leader or senior colleague. Seated as I was, between the team leads for all four Stress groups, I had an outstanding vantage point from which to observe these interactions.

As these critical incidents unfolded in real time the actual data collected varied significantly, from a few spoken snippets per incident to complete verbatim transcripts of the interaction. For a number of the incidents I was able to obtain brief clarifying interviews with some of the participants (or with others who had observed the incident with me).

These critical incidents deepened my understanding of common breakdowns in the routine work of the group. They also provided a solid body of natural observations which would prove useful in understanding the differing perspectives of the key information flow bottlenecks raised in Phase II and form the foundation for the cases collected in Phase III.

• **ROC survey** – Early on in my artifact identification work, it became obvious that the Record of Conversation (ROC)⁹ was the focal point for collaboration on any given job. In order to get to know this one important artifact well, I surveyed an entire month of archival, paper ROC records. My initial intent had been to review the on-line ROC summaries, but was told by a number of engineers that the on-line versions are only really used for convenience as pointers into the paper stacks which held the truly valuable information. (As my time at GTS-West progressed I was able to more clearly delineate the different uses for the paper and electronic versions.)

⁹ The ROC is formally presented in the following chapter, but a brief overview is appropriate here. The ROC is a uniquely identifiable record of the resolution of an operator's service request. It is the primary object in the formal, legal paper trail required by the FAA to trace any service action requiring stress analysis. The ROC itself is a paper packet usually five to ten pages thick. A typical ROC contains a two page action summary, the original repair request from the operator, and the most relevant details pertaining to the analysis and final recommendation. These details are a distillation of work done to resolve the service request— material deemed by the engineer as critical or exceptional are included (e.g. sketches of proposed repairs) and those deemed peripheral or common are either excluded entirely (e.g. scratch work) or represented by pointers (e.g. the stress report which provided the load numbers for the evaluation). The ROC also has a digital analog which consists of the dense two page action summaries. These are searchable in the ROC work flow management tool for current jobs and in the STAIRS database for completed jobs.

I chose to review the records for the prior month, for only the short-haul support teams (as they had the highest number of jobs, and therefore the highest number of potential records to review). While this decision limited my access slightly (I was able to retrieve and review about 325 of the 500 jobs), the jobs described in the records overlapped with my first month of critical incident observation. Not only did this help to flesh out my understanding of these initial incidents but the immediacy of the jobs made it easier to talk to engineers about elements in the records I which did not understand.

The base of descriptive statistics which came out of this project, was a helpful foundation for comparison work later in the study. However, the most important element of the survey was understanding the role of the ROC as a pointer to information resources outside of the record. Often, these were details which had been “stripped out” before the record was stored (e.g. the original calculations that an engineer had found and modified to fit the current job).

• **Case development (lifecycles)** – Returning to Harper’s notion of information careers, in my final months at GTS-West (Phase III) I wanted to systematically trace a handful of action items from their inception (call placed by the airlines) to their resolution (response given and paperwork filed). At this point in the study, I had significant breadth of observation but wanted greater depth. The critical incidents I had collected had the limitation that they were snapshots in time and did not provide an understanding of the trajectory of the information artifact at the time of the breakdown. The sense of flow had been lost. I hoped to regain this with a more comprehensive capture of the entire

lifecycle of the artifact. Now, not only did I want to follow multiple items throughout their lifecycle, but I wanted to do it in as much detail as possible. I had been impressed with the level of detail captured by Mambrey, et. al. (1997) in their following of a speech draft through out a minister's office the German government, and desired to capture the same degree of specificity with this final survey.

• **BPR process flow diagrams and discussions** – Mid-study (phase II) I had the unique opportunity to be a participant observer for a 25 week business process re-engineering (BPR) team. This was a team of eight engineers charged with understanding current work practices and looking for ways, technical or organizational, to improve them. The team was focused on the working relationship between Structures and Stress and therefore culled members from those two groups and the periphery. There were three representatives from Structures and three from Stress (with care taken to have a representation from senior to junior members) and one each from Repair Design and Hydraulic and Mechanical. I was welcomed as a non-voting member, and in exchange agreed to some clerical duties (such as taking the minutes after some political contention with the team's secretary caused her to resign). In this capacity I was able to easily record field notes and review documents relevant to the debates.

This experience afforded an unparalleled opportunity to understand the different work cultures of the group and their collaboration, especially how they viewed the information flows and mediating roles of the information artifacts. I was able to capture most of the

debates surrounding these, and was able to have a number of informal interviews with the participants off-line to better understand their stances.

- **Standard metrics** – Lastly, I collected a variety of descriptive statistics throughout the year, from management’s call resolution metrics to simple samples of system usage.

While little of this is systematic enough to stand on its own, it proved valuable as background information and yet another point of triangulation in my analysis.

4.3.3 Understanding the Work

Harper’s third class of data collection involves gathering the most robust, valid data possible. Every data collection technique has inherent limitations and biases. A solid study relies on a triangulation of data gathered from multiple sources, in multiple ways. The more diverse data that are represented, the more balanced and less biased the findings will be.

Harper’s primary concern was getting the full story out of interviews, which often meant balancing a more formal interview with lots of informal input (e.g. the “breakfast clubs” at the IMF or water cooler discussions elsewhere). (Orr also found that some of his best data came from the quasi-social morning, breakfast planning meetings of the field service technicians.) In the same vein, I found that some of my best data regarding what was really going on came from frequent lunches with engineers. Away from the “hear all, see all” cubicle farm, and in an informal atmosphere, this is where people would really open up and “tell it like it is.” I routinely ate with small groups of engineers, and found the

interaction amongst the members to be as valuable as their direct interactions with me. In addition to lunches, I also tried to avail myself of every social opportunity possible, from morning breakfast hour to birthday celebrations to the annual departmental picnic. These informal interactions were a valuable balance to my routine of semi-structured and open interviews.

Harper's secondary concern was balancing the stories told in the interviews with the daily reality that was unfolding around him in his sites. The most common remedy was observation, ideally participant observation. Based on his experience studying air traffic controllers, he comments that

“in some settings it is certainly easy to observe, but downright impossible to participate. In ATC one can while away hours watching the work, but it would be extremely unwise actually to undertake the work (even if one were allowed to). One needs to observe and sometimes to supplement that with participation, if it is appropriate. It is certainly not essential.” (1998, p.85)

Over the course of the year, I picked up some of the basic principles involved with material strength analysis, but clearly never enough to ever be considered a true service engineer or analyst. I was involved with the ebb and flow of daily office experience deep enough and long enough to truly be able to view the world through their eyes. A colleague recently suggested that you know you have reached this point as a researcher when you can laugh at humorous occurrences with them, and know what it is about the situation that makes it so funny. Humor is indeed a subtle marker.

I had other opportunities which further developed my insight. Clearly the BPR team was a major window into the minds of the assembled engineers. I also found that the weekly

staff meetings (I alternated between Stress and Structures every other week) were another superb opportunity to gain insight. I came to learn not to focus as much on the information being presented, but instead on the audience's reaction to it. Lastly, the cases (described above) provided a rare opportunity to expand my coverage. I used the survey as a tool to gain extended interviews with almost every single engineer in both Stress and Structures. This is a strong point from which to triangulate. For example, with the short-haul teams, I acquired detailed explanations of every engineer's work practice firmly grounded in a contextual example.

In summary, I learned much about structural service engineering and stress analysis, though shy of actually becoming a service engineer or analyst. Nonetheless, I was able to triangulate among my varied data sources in order to achieve an understanding of not only the formal arrangements but the various informal and tacit arrangements that make the information work.

4.4 Data analysis techniques

Equally important with an understanding of data collection techniques is an understanding of the qualitative methods of analysis and theoretical constructs which will be utilized in interpreting this data.

The underlying theoretical, sociological perspective for the study is Symbolic Interactionism. With its origins in the Chicago School of Sociology in the 1930's, this is an approach which views human action as the result of complex interactions between

individuals and social systems. It favors a level of analysis at the individual or small group and elevates the examination of every day, routine action as the best means of understanding behavior in organizations.

At the basic level, data analysis will be approached from a Grounded Theory perspective (Strauss and Corbin, 1990). A method logically born of the symbolic interactionist movement, the underlying tenet of Grounded Theory is that deep understanding of social phenomena is best garnered from extensive real-world observation. Grounded Theory foregrounds this data and relies upon systematic coding to evolve hypotheses. This coding process is emergent (recurring patterns give rise to hypotheses) and iterative (an aggregate understanding is built upon multiple coding passes through the data). This Grounded Theory analysis was theoretically informed by Distributed Cognition (Hutchins, 1995; Halverson, 1995).

Having presented the methodological underpinnings of this study, discussed its general research design, and outlined the site selection criteria, attention will now shift toward the site itself. The following chapters will introduce GTS-West, describing its organization, environment, people, and routine work. This is all laying a foundation for the subsequent analytical chapters which examine reuse in this context.

CHAPTER 5: Organization, People, Systems, and Environment

5.0 Requisite background information

This chapter provides the relevant background information required to understand the coming analytical chapters. It presents an overview of the organizational structure, key roles, central information systems, and current operating environment. Emphasis is placed on the aspects which either constrain or facilitate reuse behavior in collaborative problem solving at GTS-West.

5.1 Organization

No factor has more of an impact on the work practices of the engineers at GTS-West than the web of organizational relations within which they operate. These will be discussed in detail, but at a high level, here.

5.1.1 The corporation: Global Aerospace

Global Aerospace is a Fortune 100TM corporation with \$58 billion in revenue in 1999.

Global has been an innovator on an international scale with respect to all aspects of aerospace industry, both military and civilian.

Global Aerospace has grown rapidly the past decade as a result of a series of mergers and acquisitions precipitated by the upheaval of the industry in the United States with the end

of the cold war. As a result Global is one of the largest corporate employers in North American, maintaining approximately 197,500 employees during the course of this study. The company is so vast that there is little employee awareness of other divisions outside of their geographic or functional area. Global's corporate culture is understandably in flux as it adapts to this growth, but one element has remained constant – bureaucracy. It is not surprising in such an engineering firm to find that its policies and procedures are decidedly labyrinthine.

5.1.2 The corporate division: Global Airframe

Global's founding division and still undisputed king of its revenue stream is Global Airframe (GA). GA is responsible for all aspects of the corporation dealing with the design, manufacture, and support of commercial aircraft.

The organizational structure of the GA corporate division mirrors its product. GA has only a handful of distinct aircraft models and almost all innovation results in specialized derivatives of these base models. Each resultant series is considered a member of that model family.¹⁰ All of the related organizational units (e.g. design, sales, manufacturing, support) are grouped by related model families and operate semi-autonomously. As a result, each grouping has a decidedly different character and is one of the noticeable cultural divisions within GA.

Interwoven with these group-based cultural distinctions are a set of disciplinary-based cultural distinctions. Each model family grouping has a web of major supporting sub organizations such as design engineering, modification engineering, production engineering, service engineering, and vast manufacturing groups. Each of these is comprised of a number of decidedly different communities-of-practice. Thus professional relationships within GA tend to form either along the warp of one's model family grouping or along the woof of one's community of practice.

5.1.3 The functional division: Technical Services

Global Technical Support (GTS) is the division within Global Airframe which provides technical support for the operators of Global aircraft (e.g. commercial airlines, air cargo companies). This support is required of all domestic airframe manufacturers by the Federal Aviation Administration (FAA). Although a mandatory service, technical support has been both a core asset for sales and is increasingly viewed as a potentially lucrative revenue stream of its own for GA.

¹⁰ This practice is standard within the industry. It is logical given the astronomical expense of aircraft design and requirements for proven reliability and safety. Well proven models, some of which have been in active service for over 50 years, form a solid foundation for future development and customer satisfaction. As an example, at industrial rival Boeing, a base model aircraft would be the 737, while the 737-800 series (an aircraft outfitted for longer range travel and higher passenger capacity through a stretched fuselage, modified wing, and improved avionics) would be a member of its model family.

There are GTS facilities distributed among the major model family groupings¹¹ and are located throughout the United States. Given their semi-autonomous organizational structure, each GTS facility is, not surprisingly, quite different from each other in organization, policy and procedure. While all three commercial GTS facilities were visited for this research project, this study focuses exclusively on GTS-West.

5.1.4 The Facility: Global Technical Services-West

The GTS-West facility is responsible for a particular grouping of non-Global aircraft model families, dating back to the early 1930s.¹² All of these aircraft models are post-production, having been previously manufactured by a merger-partner. GTS-West supports over 3,000 of these active, in-service aircraft.

As with all the other GTS facilities, the GTS-West office is located on site at the production facility for their grouping of aircraft model families. This location decision is deliberate, providing the service engineers with easy access to all of GA's expertise for their grouping, from original designers, to sales staff, to technical writers, to the team that rivets the nose cone fasteners. This proximity also allows GTS-West to draw its staffing heavily from those who have worked on particular aircraft models within another functional division (e.g. design, manufacturing).

¹¹ Three of these deal exclusively with commercial aircraft and are directly located within GA. The fourth is a hybrid facility dealing with support for military aircraft.

¹² The FAA mandates that as long as even a single plane remains in service, the entire model family must be supported.

5.1.5 GTS-West: Departmental overview

The GTS-West department consists of over 200 engineers and administrators divided among the following functional work groups:

Table 2: GTS-West Departmental Overview

Aircraft Service Areas	
<i>GTS-West Group</i>	<i>Responsibility</i>
Avionics & Electrical	All flight control and on-board electrical systems.
Hydraulics & Mechanical	All mechanical systems, most notably flight surface control and landing gear.
Propulsion & Environmental	Internal air flow systems and interface with engine units. (GTS-West does not support engines.)
Structures & Payloads	All structural elements, both internal and external.
Analytic Support Areas	
Stress	Performs mathematical modeling for repair durability.
Damage Tolerance Analysis	Advance stress analysis examining interactive effects between proposed and prior repairs.
Aging Aircraft	Designs special inspection programs for heritage aircraft.
Repair Design	In-house drafting team, local experts on parts.
General Customer Service Areas	
Airline Support	General customer service group, account managers reside here.
Warranty	Handles all warranty related issues for the operators.
Fleet Support	Tracks issues across all operators which affect all series aircraft within each model family.
Fleet Statistics	Collects standard operating statistics on each individual fuselage from operators. Maintains and distributes information internally as requested.
Flight Safety Investigations	Team dedicated to assisting investigations into incidents or accidents around the globe. Coordinates all related external requests for materials or expertise from GTS-West.

The simplified organizational chart provided below, provides an overview of these groups and their reporting structure.

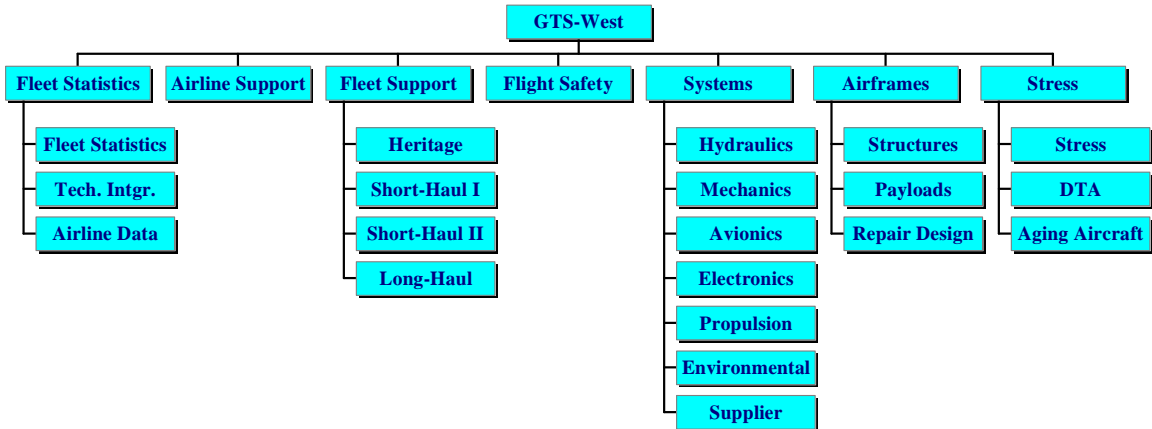


Figure 1: GTS-West organizational chart

While these basic groups have made up the core elements of the GTS-West department for many years, they have existed in subtly different configurations and are in noticeable contrast to the other GTS facilities. The two most important differences are that GTS-West is very centralized (the majority of the resources it needs are to be found within departmental boundaries) and politically powerful (historically the design engineering arm of airframe manufacturing companies are at the top of the political pyramid with support at the bottom, even below manufacturing. However, with this particular grouping of model families now out of production, support ranks much closer to the top.)

Airplanes are exceedingly complex pieces of machinery, with many interdependent systems. Thus, it is the exception, rather than the rule, that a service request can be resolved without consultation with at least two of the groups listed above. Almost all work requires inter-group collaboration of some degree. These collaborations can be loosely coupled (e.g. Electrical and Repair Design) or tightly coupled (e.g. Structures and

Stress). The most interesting venue for understanding the role of reuse in collaborative problem solving would come from one of these tight couplings, preferably one with a high volume of transactions. Thus the data collection for this study has concentrated on the relationship between the Structures and Payloads group and their primary analytic support group, Stress. These Airframes groups and their analytic support teams will be described briefly below.

5.1.6 GTS-West: Airframes

Structures and Payloads is a group of 31 service engineers, responsible for supporting all aspects of the airframe and all internal structural elements on both cargo and passenger aircraft. These engineers are managed as a single group, with a joint manager, but in day-to-day operations they operate quite independently under their own supervisors (“team leads”). The study focused primarily on the main Structures group because they fielded the most calls and their requests fit the standard prioritization pattern of the rest of the department. In contrast, Payloads was less representative, routinely focused on special long-term projects.

Indeed, Structures has the heaviest volume of service calls in the department, more than double the nearest group, Hydraulics and Mechanical. In 1999, they fielded approximately 12,000 actions and this number is climbing rapidly. (The increase has been 8-10% annually since 1993 and it is expected to be even higher now that all of their model families are post-production.)

There are a few characteristics of the Structures group which are important to highlight.

First, as service engineers, they are the front-line interface with the customer. They directly handle all structural technical support queries from the airlines. In this capacity their experience is quite similar to other common technical support environments. The days are long, grueling, and high-stress. The stream of incoming requests is unrelenting. Praise for a job well done is rare, while operator complaints are the norm. The primary motivations for the service engineers are also similar: a desire to work with people ('be where the action is'), a puzzle solving personality which is satisfied by this intellectually challenging work (to some degree this is akin to playing 'engineering Jeopardy'), being a valuable, if largely invisible, component of the success of a product ("without us, the planes don't fly"), and knowing that no two service requests will ever be exactly alike.

The service engineers know their customers well and are quite savvy at tailoring information to match their needs and abilities. Many of these service engineers have come up through the field support or manufacturing segments of Global Airframe and are accustomed to being close to both the customer and product.

Second, as with any good service group, Structures is very sensitive to customer feedback. Weekly staff meetings review operator feedback and discuss particularly problematic cases. Knowing that all operators are potential customers for future aircraft sales, the engineers are very attuned to pleasing the customer. Management encourages the engineers to not question the operators, and instead adhere to common support

mantras of ‘there is no stupid question’ and ‘the customer is always right.’ (While behavior is understandably anything but compliant internally, the service engineers are savvy with regard to face management, and the public image of the department among the operators is routinely high.)

Third, as a result of this customer-centered business model, Structures breeds a culture of efficiency and expediency. Everything is monitored and measured by management: timeliness, completeness of response, customer satisfaction, etc. These metrics are directly tied to each individual’s performance rating, salary increases, and bonuses.

Lastly, to the outside world, ‘the buck stops here’ – these service engineers have the final say on all repair recommendations and are ultimately held individually responsible for them. They are meant to be the sole contact the operator has with GTS-West, black boxing all other functions.

5.1.7 GTS-West: Analytic support teams

Many of these black boxed functions are provided internally by analytic support teams. These specialist teams are formally operated as services, with Structures contracting resources as needed on a case-by-case basis.¹³ In day-to-day practice though these

¹³ There are even accounting structures in place to do billable hours with departmental charge backs between the groups. Since service is provided to the operators free of charge, this is all meant for internal tracking. However, currently, such figures are rarely used, only surfacing to track grossly over budget projects or egregiously resource intensive technical support requests.

accounting divisions are artificial and service engineers routinely work hand-in-hand with the specialists.¹⁴

Structures routinely draws on four analytic support teams in dealing with their ordinary work. Each will be described below, with the primary team, Stress, receiving a detailed description and the other three a cursory overview.

5.1.7.1 Stress

Stress provides all of the advanced stress analysis for the airworthiness of repair actions generated by the operators and approved by Structures. These analyses are mathematical models of varying complexity which determine the impact the repair will have on the sustainable strength of the assembly and assist in predicting the repair's longevity.

Typical results of these models involve maximum load tolerances, expected lifetime of repairs, safety characteristics of repairs, and materials performance. Stress also initiates and coordinates the FAA approval process, via designated engineering representatives, for these repairs. Structures works closely with Stress for over 80% of their actions; their job simply could not be done without this collaboration.

In spite of their close collaboration, Stress engineering is an all together different community-of-practice than service engineering. Here are some of the major differences.

¹⁴ Although these functions remain opaque to the customer.

First, by deliberate organizational design, Stress engineers¹⁵ do not have contact with the operators.¹⁶ This objective detachment is one critical component of the system of checks and balances which yields the high-reliability of response required by aircraft repair. (Not having to please the customer allows Stress to be impartial in their assessments.)

However, this detachment is a perpetual source of friction between the two groups as Stress often has to work in a contextual vacuum, solely dependent upon the service engineer for the relevant details surrounding the current job. Also, not being customer-support oriented, they tend not to understand the job of the service engineers. Some view them as little more than technically savvy receptionists, who translate requests and results to the operators.

Second, the Stress team consists of “engineers’ engineers” – more abstract and theoretical than the hands-on service engineers. These sophisticated number-crunchers have traditionally been promoted internally into this department. In recent years, intense downsizing of the design and production units at this plant have enabled some of the best analysts from these groups to join Stress.

Third, most of their job consists of gathering information to build a case for a particular repair decision and then running through the requisite arcane analyses. The analysts stand in ultimate judgement on each repair – either it is safe or it is not, period. (Usually the

¹⁵ Stress engineers are fully credential engineers, however in a work environment where engineers are omnipresent, they are routinely referred to as Stress analysts. For the sake of clarity between service engineering and stress engineering, I will also use this term frequently throughout the dissertation which I hope at no point diminishes their engineering nature in the mind of the reader.

analysts are very helpful in explaining why they reached the conclusions that did, and frequently suggest how operators can best modify unapproved repairs for resubmission. Regardless though, they are often viewed as imperious and at times even capricious by the service engineers.)

Lastly, the interplay of these factors fosters a culture obsessed with quality. Stress engineers will proudly tell you that their calculations determine whether a plane flies or not. They take this responsibility very seriously and will work on a problem until they are convinced beyond reasonable doubt that the repair is suitable for strength. This passion for error-free evaluations comes at the cost of timelines (sometimes the calculations for particularly onerous problems can stretch over days), which clearly puts them at odds with Structure's response-time focused service engineers. Everyone in the department has safe air travel as their ultimate goal, but it is in this tension of tight conflict and collaboration that all the routine work is accomplished. The following three chapters will explore how this tension is played out in routine work.

Organizationally, Stress is equal in size to Structures/Payloads with 32 engineers. Stress has a single manager, but is subdivided into four teams, each with their own team lead. Three teams are arranged to mirror Structure's subdivisions (long-haul, short-haul, and heritage aircraft) and one exclusively serves the special analysis needs of the Hydraulic

¹⁶ This separation is strongly enforced. For example, Stress engineers are not permitted to initiate telephone calls with the operator and on more than one occasion I witnessed even well meaning perpetrators being read the riot act for their action.

and Mechanical group. While this study examined all four Stress teams, the emphasis was on the three that interact directly with Structures.

5.1.7.2 Other teams

Damage Tolerance Analysis (DTA) and Aging Aircraft are related specialist Stress teams. Created as part of the industry's response to the tragic 1988 Aloha Airlines metal fatigue induced explosive decompression incident, these teams are responsible for examining the inter-relatedness of repairs. DTA provides special analysis for repairs on older aircraft that are adjacent to regions that are structurally dependent on prior repairs. Aging Aircraft designs inspection programs for active, in-service aircraft that have passed their original design lifetime. Their work is complex, but has very long lead times and few time-critical requests.

Culturally these teams may be the only ones able to 'out-geek' the Stress engineers.

These three Stress-related groups have always been close, but mid-study the DTA/AA manager became the manager of the general Stress group, bringing them even closer.

They are occasional resources for the regular Stress engineers, helping with particularly challenging problems, and sometimes assist in more prosaic ways when there are staff shortages.

Composite specialists are another analytic support category. Modern aircraft design relies increasingly on composite structures. These are layered polymer components which are very sturdy, highly flexible, and have an extended lifetime. However, they are exceedingly

difficult to repair, requiring complex computer modeling for even the simplest of repairs. These specialists are often contracted from outside of Global and their skills are so critical to modern Stress work that they have been brought inside the department and seated physically adjacent to Stress.

Lastly, Repair Design is a support team which figures prominently in this research. Organizationally they are tied to the Airframes group and are primarily chartered to support Structures, but their services are employed by the entire department. Essentially they are an in-department drafting team. Informally, they harbor a wealth of local knowledge about parts and histories of special “one-off” repairs.

5.2 People

Having outlined the relevant organizational structures for GTS-West, we turn now to an description of the formal roles in the department as a means of examining those same structures though a different lens. The categorical cast of characters at GTS-West includes the following.

5.2.1 Director

The director’s official role is to oversee the activities of all groups within GTS-West. However, since the merger with Global, he has delegated this primary responsibility to his senior management team and instead focuses on managing the relations outside of GTS-West, most notably with the other GTS facilities, Global headquarters, and important operators. He views his contribution to the customer support mission of the

department as dealing efficiently with all of the high-level organizational issues, so that the work environment is stable and predictable, and that the team members can focus on their tasks uninterrupted.

Historically the directorship of the department has been a very stable position (the prior director had been there for over a decade). Given the new responsibilities as corporate liaison it is unclear if this will remain true in the coming years.¹⁷

5.2.2 Senior management team

The senior management team consists of a set of mid-level managers who all have had extensive experience throughout the plant. Those responsible for many teams, are understandably removed from day-to-day tasks; however, the managers for both Stress and Airframes are very hands on. Each has been with GTS-West for many years, and has earned the serious respect of their teams. The primary role of the senior managers is to serve as a conduit, raising concerns from the floor to be presented to the Director. The Stress senior manager is also the group's senior liaison with the FAA.

5.2.3 Group managers

The group managers are the equivalent of line managers who are responsible for managing their teams and handle routine day-to-day issues. They are responsible for maintaining metrics of their teams' performances and performing regular evaluations of

¹⁷ During the tenure of this study the directorship changed twice.

each staff member. While they are viewed primarily as human resource contacts by their engineers, they are frequently consulted for escalations on action items (though this usually tends to be for political problems with the operator, not technical concerns which tend to fall to senior engineers or team leads). They facilitate weekly staff meetings which involve a discussion of the prior week's unusual jobs and "flow down" information from the Director. They also help manage relations between GTS-West teams. For example, the group managers for Structure and for Stress now meet daily to discuss jobs which are in danger of being "waved off," not meeting their operator requested response deadlines, and to coordinate overtime schedules.

The turnover for these positions is very high, with one engineer commenting that they are practically "revolving door" positions. During the duration of this study both of the group managers for Structures and Stress turned over. This transition resulted in multi-month periods when the teams did not have group managers. This had no measurable impact on the performance of the teams. While this may be surprising to an outsider, it made sense within GTS-West as the team leads handled all concerns for their team and merely reported directly to their senior manager. There was much discussion among the engineers of removing this position altogether in favor of strengthening the team lead position, described next, if it were not for the union and human resource issues involved.

5.2.4 Team leads

The team leads are the true local managers in all but name. (They are the highest level of union leadership in the department.) Global is explicit about classifying them as non-

management, however they perform all but the most formal responsibilities. Structures and Stress are currently configured so that there is a team lead responsible for each model family of aircraft. They are responsible for all routine work assignments and the quality of the work within their team. Almost without exception they are viewed as the most accomplished engineer, though not necessarily most senior engineer, within their team. In the past, these had been rotating positions, but recently they have been relatively stable, receiving wide-spread consensus and support even though the selection is made by the group manager, not by popular opinion. The role of the team lead, as the coordinator of his team's activities and gatekeeper for the teams' organizational memory systems, is one of the most critical roles in this work environment and will be discussed in detail in Chapter Eight.

5.2.5 Engineers

In general, the engineers in GTS-West are an accomplished crew. Most are survivors of the aerospace industry upheaval of the early 1990's. Demographically, they are well educated (most holding advanced degrees, including some doctorates within Stress), reasonably gender represented for an engineering position (12% female), and intensely multi-cultural, multi-ethnic, and multi-national. The senior engineers within both groups are viewed with respect and are leveraged as vital information resources, especially regarding information contacts around the plant outside of the department.

As noted in the organizational description above, the Stress team tends to have more design engineers with extensive academic training and lower employee turn-over.

Structures has engineers with extensive field experience (“wrench time”), many of whom have worked up through Global’s maintenance organizations,¹⁸ and have personalities well matched to technical support (many more admit to thoroughly enjoying their jobs than their counterparts in Stress.)

As an introduction to the working conditions of these engineers, photographs of standard workspaces for both Stress and Structures are provided on the following page.

¹⁸ During the course of this study there were two internal transfers into Structures, one from Stress and one from Repair Design. Another three came from elsewhere in Global to fill five vacancies.

5.3 Systems

The routine work of these individuals is mediated by a host of information systems. Some are deployed department-wide, such as their Microsoft Exchange based e-mail system, others are used by single individuals. With Structures and Stress there are also a significant amount of in-house or self-developed software tools whose use is spread by word of mouth.

Information systems services for the department are outsourced. Although the service provider has a significant staff located within the same building as the department, service complaints regarding slow service and limited support for older systems are common. After years of limited or no equipment improvement, the period of this study witnessed two complete group-wide machine upgrades for Stress and one for Structures.

There are two information systems though that dominate the activities of the service engineering and analytical support groups: GlobalCOM and the ROC system. Both are work flow management systems with striking points of similarity between them. However, for two years, and into the foreseeable future, both systems have been operated in parallel at great expense. Any time that equivalent, redundant systems are operated simultaneously, there is likely a set of interesting underlying issues. The early days of this research project sought to understand these issues, and the theme is raised again in Chapter Eleven. For now both systems will be briefly overviewed and compared.

5.3.1 GlobalCOM

GlobalCOM¹⁹ is a work flow management system which had been developed for other GTS offices. It has experienced widespread acceptance, and has expanded to include coordination between non-GTS members such as Global's field service reps. It was introduced to GTS-West shortly after the merger with Global. It is only available to service engineering teams, but has received tremendous acceptance among them.

Its foundational data structure is the action item, representing a repair request from an operator. The file for an action item has a very flexible structure allowing for the inclusion of all relevant digital materials pertaining to that job (e.g. electronic correspondence, faxes, photographs).²⁰ Its primary purpose is to coordinate the activities required to bring the action to completion. It has a peripheral purpose of maintaining a searchable history of all completed action items. The search facility is very limited (keyword only on a handful of fields), and is used by engineers at GTS-West to only recall jobs which they had previously completed.

Following is a screen shot from an incoming GlobalCOM request:

¹⁹ GlobalCOM has an Oracle back-end and is run on PCs in an emulated X-windows environment.

²⁰ One of the most important of the many innovative GlobalCOM features is the automated scanning and routing of requests. Given the variability of communication technologies that operators and repair stations have around the globe, the fax machine was settled upon as the lowest common denominator. Operators fax their requests into a central GlobalCOM telephone number, the faxes are digitized, and then routed (automatically if possible, by hand if not) to the appropriate service engineering contact (which at GTS-West is the team lead).

```

DATE: 05-Oct-99 02:14PM
PREPARED FOR: Ldt
PAGE: 1

View Message

Message Number:      Action File Name:      Stat:
-----
ZEO-LAX-00-00591F    ZEO-LAX-00-00591F      Open

Model: N-20

Subject: BRACKET - INBOARD FLAP - OUTBOARD FORWARD HINGE -
LIMITS/REPAIR -
8110#5

From: gcapu@gcapu-pop-02.global.com

/ATTN (SSM) P. HALLSWORTH AIRLINE SUPPORT MANAGER
/CC (PSRLAX04) K. TRAN FSR - LOS ANGELES

ZEO-LAX-00-00591F      05-OCT-1999
ATA 5752-00 MODEL N-20 26-OCT-1999 H
BRACKET - INBOARD FLAP - OUTBOARD FORWARD HINGE - LIMITS/REPAIR -
8110#5
REF /A/ 2 page fax /N/

Following message sent two P. Hallsworth with copy to K. Tran.

ZE have requested a Global DER review of the ref /a/ ZE generated
Procedure and provide 8110-3 coverage for this data.

Please ignore the airplane effectivity at the bottom of the ref
/a/ sheet. ZE is requesting effectivity coverage for the
following airplanes:

Q0942E-Q0944E
Q0946E-Q0949E

Action:
-----
Please review the ref /a/ ZE procedure and provide an 8110-3
form covering this data.

This request has been coordinated with P. Hallsworth.

GTS-GlobalCOM-12.1.9 TUE 05-OCT-99 21:51Z

```

Figure 4: Sample GlobalCOM request

5.3.2 ROC system

The Record-of-Conversation (ROC) system²¹ is a proprietary workflow management system developed in-house for tracking production problems, but co-opted and adapted by GTS-West to support service engineering. Until the introduction of GlobalCOM it served as the primary workflow management system for the department, but now only provides those functions for the analytical support teams. It remains the official record of a completed action item. Technical support for in-service aircraft is part of a regulated

²¹ The ROC workflow system is a proprietary program running in VM6 on an IBM mainframe. The historical archive is done in STAIRS, a legacy IBM database system famously studied in information science by Blair (1985;1996).

industry, in which an auditable legal paper trail is required for all actions. The ROC performs this function at GTS-West. (Further details on the aspects of regulation are covered in Chapter Six.) Its secondary service is as a rich searchable historical archive. While the system is only fifteen years old, it provides a database of nearly thirty years of historical data.

The ROC system itself is a hybrid digital and paper based system. The core element is the action item, the same as with GlobalCOM. This however exists in two forms. The first is a two page, highly condensed digital summary. This summary is what is passed through the work flow system and what is electronically archived. Accompanying this digital summary is the ROC packet – a coded file folder containing a print out of the summary and all relevant supporting documents (akin to the attachments in GlobalCOM, but not as freely defined). This also is physically passed around following the job assignment and is archived in vast filing cabinet bays.

Following are screenshots of a closed-out ROC digital summary:

<p>PRODUCT SUPPORT - RECORD OF CONVERSATION PAGE 1 OF 1 GTS 56-851 (Rev 5-94) LOG NO: 99-07-11-008 LDT02913</p> <p>PART NUMBER: 6845937-1 NEXT ASSEMBLY: 6845931-403</p> <p>SERIAL NUMBER: N/A ATA CHAPTER: 53-12-01</p> <p>DRAWING TITLE: LONGERON 20 -- STA 993 TO STA 1103</p> <p>PRIORITY: URGENT SERVICE ENGINEER: LINDA TILLMAN SUBJECT CLASS: INITIAL DEPT: L21 EXT: 17031 DATE SUBMITTED: 07/11/99 GROUP: CUSTOMER SERVICE REP DATE DUE: 07/11/99 SUPERVISOR: M. D. PIPPIN</p> <p>FAA/DER APPROVAL REQ: YES FORM 8110-3 DATE: 07/11/99</p> <p>MODEL: N-20 SERIES: 25 OPERATOR: TCA FUSELAGE NO: 2254 FSN: 58528 REGIS NO: JN-SLR FLT HRS: 36380 LANDINGS: 30777 CURR CERT GROSS WT: 130</p> <p>CORRES CONT: LONGERON 20 TELEX/MEMO: SLD-SNN-99-00053F</p> <p>PSE AFFECTED: NO PSE NUMBER: METHOD OF DETECTION:</p> <p>DISCREPANCY: ROC 99-06-13-003 PROVIDED FAA/DER APPROVAL TO SPLICE LONGERON 30 BETWEEN STA 993 AND STA 1050. THE REPAIR WAS ACCOMPLISHED PER ROC SKETCH 99-04-19-007 WITH REVISED DIMENSIONS PER ROC 99-05-08-014 (FUS 2153).</p> <p>TRANSCON AIR NOW INDICATES THAT THIS REPAIR HAS BEEN COMPLETED ON FUS 2254 EXCEPT THE VERTICAL WALL THICKNESS OF THE SPLICE HAD TO BE REDUCED FROM 0.250 INCH TO 0.187 INCH TO MAKE IT POSSIBLE TO INSTALL THE FASTENERS IN THE TROUGH. THIS REWORK GIVES THE TROUGH A WIDTH OF 0.566 INCH INSTEAD OF THE 0.44 INCH.</p> <p>*** AIRPLANE SCHEDULE FOR FLIGHT TEST FRIDAY - 12JUL99 ***</p> <p>SPECIFIC QUESTION TO BE ANSWERED: IS THE ACCOMPLISHED LONGERON 20 SPLICE REPAIR ON FUS 2254 ACCEPTABLE WITH THE VERTICAL WALL THICKNESS ON THE -5 SPLICE REDUCED TO 0.187 INCH FROM 0.250 INCH?</p> <p>PLEASE PROVIDE FAA/DER APPROVAL.</p> <p>ANSWER: YES.</p> <p>"GLOBAL AEROSPACE CORPORATION PROPRIETARY INFORMATION" "THIS PAGE IS NOT FOR DISTRIBUTION OUTSIDE THE COMPANY"</p>	<p>PRODUCT SUPPORT - RECORD OF CONVERSATION PAGE 2 OF 2 GTS 56-851 (Rev 5-94) LOG NO: 99-07-11-008 LDT02913</p> <p>THIS REPAIR SPECIFICATION, PER ROC 99-07-11-008, HAS BEEN SHOWN TO COMPLY WITH THE TYPE CERTIFICATION BASIS FOR THE AIRCRAFT AND IS FAA APPROVED.</p> <p>CAN ANSWER BE USED IN THE FUTURE?: NO</p> <p>ANALYST: JON C. LEW CHECK: OJB DATE: 07/11/99</p> <p>DESIGN COORDINATION: O. J. BROWN (DER)</p> <p>ENO.....: - CCN: EJ4Q83B1 JOB HOURS: 2</p> <p>REFERENCES: GTS-J6952 IV-2, TJ19875 P.13</p> <p>FOLDER NO: - LOG NO: 99-07-11-008</p> <p>"GLOBAL AEROSPACE CORPORATION PROPRIETARY INFORMATION" "THIS PAGE IS NOT FOR DISTRIBUTION OUTSIDE THE COMPANY"</p>
--	---

Figure 5: Sample ROC request

5.3.3 GlobalCOM / ROC system compared

The following chart provides a quick comparative overview of the two systems.

Table 3: Comparison of GlobalCOM and ROC Systems

GlobalCOM	ROC System
Workflow management (team based with personal workspace)	Workflow management (team based with personal workspace)
Internal / external communication	Internal communication only
Resource aggregation	Resource pointers
All digital record	Hybrid record (digital and paper)
Historical action item archive	Historical action item archive
Limited search capability	Flexible full-text search capability

Due to the costs incurred by the obsolete operating environment and hardware of the ROC system, and the company-wide standard of GlobalCOM, there have been plans from the very beginning to phase out the ROC system and cover its duties with GlobalCOM. This met tremendous resistance from within GTS-West, due to perceived alterations in the processes of reuse. Instead plans have been changed to the eventual integration of ROC functionality into GlobalCOM. In the meantime the department has the commitment from Global to operate both systems in unison.

5.4 Environment

The final section of requisite background knowledge for the study involves highlighting some key conditions of the operating environment with Structures and Stress during the period of data collection.

5.4.1 Productive culture of conflict

As presented in their group descriptions above, both Structures and Stress have different organizational mandates, goals, and world views. Often these are orthogonal and a source of significant friction between the individual teams and engineers. Individuals on both sides described this conflict as being beneficial for high reliability operations. It ensures that all positions are considered, and that when operating optimally, the customers' interests are perfectly balanced with Global's interests and the FAA's requirements for safety. When they are out of balance, things can degenerate quickly.

In an incident informally described as the "October/November war" the balance of power among the groups had shifted in favor of Structures who, in trying to meet new efficiency requirements placed on their teams, hounded Stress on issues of timeliness. On particularly heavy volume days during this period, the usual cordial informal relationship between the two groups evaporated and the formal, rarely invoked "service center" model was espoused. The later was so problematic,²² and the stress on the team leads who were the mediating points between the groups was so great, that there was immediate intervention by the senior management team. This included some organizational changes and some joint planning efforts, such as the business process reengineering effort. The short-term fixes sought to identify the issues, address the pressing problems, rebalance the system, and return it to working order. The long-term efforts sought to redesign the balance in order to address the growth problems described below.

²² This is akin to union workers "working to rule" instead of striking in protest.

5.4.2 Union issues

Another, more subtle tension during this period involved the local professional engineering union. It is important to note that all GTS-West engineers are unionized while management is not. (The position of team lead represents the highest level of supervision by a union member.) While other unions at this plant have had a notoriously raucous history, the engineer's union is relatively benign at GTS-West and the working relationship is deemed as very good by both sides.

This union/non-union distinction pervades all of the working relationships. Certain information flows stop at and certain interactions are carefully managed at this invisible union barrier. Normally this barrier is omnipresent, but backgrounded at GTS-West. However, two issues during the course of this study heightened the tension between the union and management.

The first was the selection of the replacement manager for Stress. There was concern among the engineers regarding the selection process and the long delay in filling the position. When asked about the lack of input, a senior manager described:

“Oh, no, that's a union issue. The way things work, union employees cannot have a say in non-union appointments. Period. The union has griped about this, 'why don't we have a say in these things.' They just don't, for a variety of reasons.”

In the end, the engineers universally affirmed the individual selected for the position as an exceptional choice and the tension was diffused.

The second case involved two other GTS facilities. At these facilities, the engineer's union, equivalent to the one at GTS-West, had its first ever extended strike. This 38-day strike was one of the largest white collar walkouts in United States history. These GTS facilities were quickly crippled in all but emergency functions. Tensions were heightened at GTS-West when it was unclear whether engineers would be temporarily relocated to the other GTS offices to provide coverage or whether teams would be forced to cover critical work which would be rerouted from these facilities to GTS-West.²³ In the end, the strike was resolved before either of these plans were implemented.

5.4.3 The current dilemma

In conclusion, there are a set of external forces which are straining the work practices at GTS-West, forcing reflection and eventual evolution. They are as follows.

- Rapidly climbing call volumes. As highlighted in the discussion of the Structures group, as more aircraft are out of production but still in-service the number of service calls will continue to rise. Also, the model families that GTS-West supports are renowned in the industry for their longevity. A sizeable minority of the current operating fleet have gone past their design life and are expected to continue to be operational for decades to come. While a great source of pride for Global, and value for the operators, the increasing burden on the technical support staff is significant. Another reason for the increase in call volume is the world-wide increase in air

²³ While there was little discussion of a sympathy strike at GTS-West, there was great personal discomfort about potentially having to “scab” activities which would hurt their fellow strikers at the other offices.

traffic, a trend that is not expected to reverse. The final reason is the growing trend among operators of outsourcing their maintenance to specialized repair companies, outfits which do not have the training and resources that Global provides the original purchasers of their aircraft.

- The resources at GTS-West are severely limited. There have been few capital improvements in the department for the last decade and they are enduring a multi-year hiring freeze. While they have not experienced the staff cutbacks demanded of other GTS facilities, GTS-West is already a desperately lean organization with little margin for error. There are no plans for increasing staff coverage in the foreseeable future.
- The quality of technical support has become a major point of comparison for potential buyers of new aircraft between Global and its international competitors. GTS-West is under intense scrutiny to improve the quality of their already legendary service (e.g. one goal is to ensure that customers are 99% satisfied with their response and 91% or more of their requests are resolved within their requested timeframe).
- Technical support has traditionally been a cost sink for Global, however, given strategic shifts in the commercial aviation industry, Global is reconsidering the role of this service. GTS-West is now under the directive to become a major source of revenue generation not drain.

In the next chapter we will examine in detail how all of these elements, the organizational structure, organizational roles, supporting systems and conditions of the operating environment work in concert to complete a routine repair request.

CHAPTER 6: Understanding Routine Work

6.0 Defining the work of GTS-West

The preceding chapter presented the organization, organizational roles, key information systems, and contemporary operating environment at GTS-West. This chapter will provide an introduction to the routine work of the department. It will focus specifically on the routine interactions between Structure and Stress involved with completing a technical service request. First, it will describe the customary work processes and constraints on those processes. Then, it will illustrate these with an example culled from the observed critical incidents.

6.1 Collaboration essentials

As already discussed in Chapter Five, Structures and Stress collaborate to some degree on over 80% of all incoming service requests. Through a series of redundant checks and balances embedded into the organizational structure and processes, and the result of divergent perspectives between the two groups, GTS-West is able to ensure a high degree of reliability, accuracy, and customer satisfaction in their problem resolutions. A closer examination of the collaboration required to accomplish this is warranted.

Organizational structure within GTS-West is geared toward inter-group collaboration. Three engineering teams within Stress are organizationally designed to align with the

equivalent groups within Structures.²⁴ These groupings are by aircraft model family: Short-Haul (intra-continental), Long-Haul (inter-continental), and Heritage (models out of production for approximately 20 years or more). The teams are staffed according to the number of active, in-service aircraft in their model family and a historical projection of their anticipated call volume. The Short-Haul dyad is the largest with 27 members (14 in Structures, 13 in Stress), servicing a fleet of almost 2,100 aircraft.²⁵ They process, on average, 21 action items per day. Long-Haul is the second largest dyad with twelve members (eight in Structures, four in Stress). Many of the long-haul aircraft are currently used as cargo freighters, which have different service needs than passenger aircraft. This helps account for the lopsided proportioning of engineers in favor of the service engineers. Lastly, the Heritage aircraft dyad is comprised of 8 members evenly split between Structures and Stress. The average action items processed per day for Long-Haul and Heritage are 12 and 6 respectively. These figures are represented on the illustrated blow up of the GTS-West organizational chart below.

²⁴ Recall that the fourth team is a specialist team dedicated to Hydraulic and Mechanical service requests.

²⁵ Given their size and call volume short-haul was the focus for most of the data collection in this study.

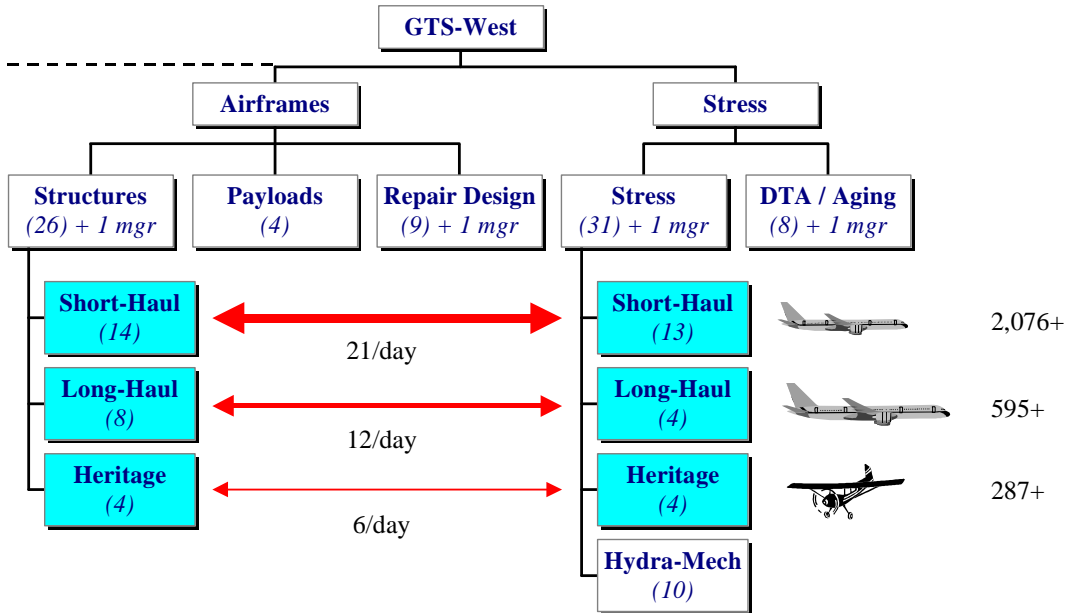


Figure 6: Collaborating teams

6.2 Service expectations, constraints, and abilities

The GTS-West staff takes pride in the quality of their work and their industry-wide reputation for service. Their work is evaluated by the following criteria:

- **Reliability** – That the response to the operator’s request is precise, definitive, correct, complete, and error-free. Not only are these high-reliability situations but they are also high-liability situations.
- **FAA compliance** – That the response to the operator’s request is in full compliance with all FAA regulations and that approval has been obtained directly from the FAA

for any non-standard repair (8110 or AMOC). In addition, a FAA approved record of the transaction, the ROC, has been completed and filed locally.

- **Response time** – That the operator’s request is resolved within the time period requested. If it is not possible to do so, the service engineer has notified the operator and secured their approval for an extension
- **Customer satisfaction** – That the request was correctly understood, that the response was appropriate, and that the operator understands how to implement the response (e.g. clear directions suited to their context).

In respect to key metrics covering these criteria, the Structures team is quite successful, routinely rated as the top support group at GTS-West.

Recall the current problems facing GTS-West presented at the end of Chapter Five. That Structures can achieve this degree of success with such limited staffing and resources is a significant accomplishment. They are able to attain this degree of efficiency due to three inter-related factors.

The first factor is the highly generalized knowledge of the workforce. Most technical support groups in the aerospace industry are comprised of teams of specialists (e.g. an expert in the repair of engine pylons for the N-20 series aircraft). In contrast, GTS-West has aggressively cross-trained their service engineers, expecting them to work outside

their immediate expertise. As a result most Structures engineers pride themselves on not only being able to repair any structural element in their *entire* model family, but also being able to address repairs in other model families as well.

The second factor is the service engineers' ability to leverage the information in their environment to augment their personal expertise. Engineers regularly comment that the ability to navigate the information space (which they define as the ability to find and use printed resources and human experts) is one of the most crucial job skills.

The third factor is fostering a culture of knowledge reuse. GTS-West maintains an extensive repository of high-quality, detailed cases over time. Being familiar with the value of reuse, both Structures and Stress engineers habitually consult the repository when working on a new job.

Again, these three factors are interrelated. The generalized expertise of the engineers allows them to be able to thoughtfully consult other experts, to make sense of a diverse literature, and meaningfully interpret a wide range of historical examples. In turn, that external information supports them when they are stretched beyond the bounds of their immediate expertise. It is a self-reinforcing cycle which has been used to great effect.

6.3 Classes of routine actions

Service requests originate from a host of potential sources: North American airlines, air cargo companies, foreign carriers, the military, independent repair shops, and even

private parties. Each have varying levels of engineering competence and available resources. As noted previously, the service engineers are highly adept at identifying the customer's profile and tailoring their response appropriately. There is a unifying characteristic to all requests though – their level of complexity.

The vast majority of technical problems arising in the routine operation of an airline fleet are resolved locally by the operator's maintenance crews using the structural repair manuals (SRM) provided by the manufacturer, Global. One engineer aptly described the SRM as a "cookbook for repairs." As with a cookbook, the recipes for repair have been distilled from a large set of possible repairs to address the most common requests, with clear, concise instructions, using proven repair techniques which have already been FAA approved. (In form the SRM is also roughly analogous to Frequently Asked Questions lists, FAQs.)

Only exceptional problems, or problems requiring special certification, which can not be handled by the maintenance crews independently are routed to Structures.²⁶ Common reasons for this escalation are as follows:²⁷

- There is no SRM repair for the given problem.

²⁶ In considering standard technical support environments, such as a telephone hotlines, the service provided by GTS-West is equivalent to "third tier" support.

- The SRM repair provided does not exactly match the given problem, and the operator suggests a reasonable deviation from the standard repair, requiring Global’s approval.
- The SRM repair provided does not exactly match the given problem, and the operator is requesting suggestions for alternate repairs.
- Regardless of SRM coverage, the FAA has a special directive for this specific damage and requires certification by both Global and the FAA for any repairs.

6.4 Prioritization

Returning to the constraints placed on the resolution of these routine requests, job prioritization is the primary stressor. It measurably alters work practices. For Structures and Stress there are three standard classes of prioritization, summarized in the table below. Each has its own character and set of constraints. The most pressing is “aircraft on ground” (AOG) which deals with aircraft in revenue operation and requires a same-day resolution.²⁸ The second class covers a range of priority situations which require next-day turnaround. These most often involve work stoppage crises at repair stations. The final class provides for the industry standard 3-5 business day response time. (By

²⁷ It is important to note here that while Global is bound to provide accurate information, they are not required to inspect the operators to ensure that they act as instructed. The ultimate responsibility for the safety of the aircraft rests firmly with the operator, not Global. Unwarranted deviance from GTS instructions though is viewed as negligent behavior.

²⁸ While it may appear logical to the reader that “airplane in air” issues should receive the highest priority, these emergency requests are not handled by GTS-West. The safe operation of aircraft is the sole responsibility of the airlines. In a mid-air crisis situation it is the operators’ foremost obligation to land the plane immediately. Diagnosis and repair activities are not to be performed in flight.

way of reference, Structures receives the highest percentage of AOG and high priority jobs of any support group in the department.)

Table 4: Prioritization

Priority:	AOG	High	Standard
Resolution:	Same day	Next day	3-5 days
Percentage:	30%	55%	15%

The introductory scenario, below, highlights an AOG action.

6.5 AOG scenario walk-through

Having covered the essentials of collaboration, service expectations, and request types (including prioritization level), it would now be beneficial to step through a typical job. Since no single example can present a holistic overview of the routine work practices, a stylized scenario is provided here. This is centered on an actual, extended critical incident which happened in the midst of resolving a routine AOG request. This incident is flesh-out with additional details from extensive observations of the resolution of other similar priority jobs. In addition, the scenario has been deliberately simplified, to highlight the major information flows, and condensed for readability. In actuality it would be continually interrupted with the breakdowns common to collaborative problem solving (missing files, unavailable people, miscommunications, distance, and timing).²⁹

²⁹ All other examples in this dissertation are presented unedited, directly from detailed case data or critical incidents.

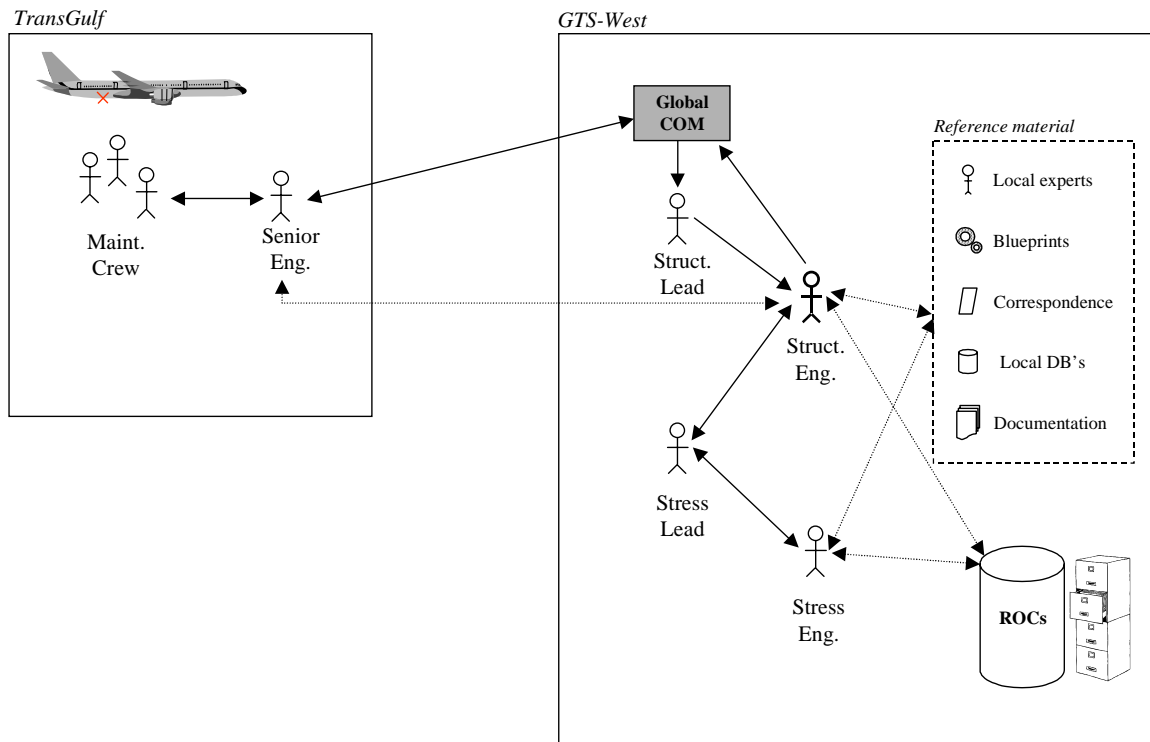


Figure 7: Information flows

Consider the following scenario: TransGulf Flight 471 to Houston is at the gate. Passengers have just finished boarding, the crew is ready, and the luggage is loaded. As the baggage conveyor is pulling away, the driver accidentally hits the accelerator, ramming the aircraft. A quick visual inspection by the co-pilot confirms a dented rear cargo door frame. Two hundred eighteen passengers wait patiently in their seats while the airline’s mechanics swarm around the cargo door taking measurements and photographs. The damage is serious enough to require guidance from GTS-West. This is a priority situation requiring immediate attention. “Airplane on ground!”

The chief engineer of the TransGulf maintenance crew, Takashi, calls the short-haul team lead in Structures, Tom, to let him know that an AOG request is en route and that TransGulf needs an rapid response. (Takashi and Tom have interacted many times over the years and are both comfortable with this initial circumvention of standard GTS submission protocol.) As Takashi is explaining the situation to Tom, another member of the maintenance crew has prepared the formal request – a handwritten document detailing the problem, a freehand sketch putting the measurements in context, two Polaroid photographs and a suggested temporary repair approach (enough to continue with this flight and complete a permanent when the aircraft returns to TransGulf’s hub airport in Pensacola). He faxes these to the dedicated GlobalCOM hotline. The digitized data is near instantaneously routed to Tom who is still on the phone with Takashi.

Within GlobalCOM Tom assigns the job to Nancy, a senior Structures engineer, based on content (he knows that she is a reasonable expert on cargo doors) and workload (he scans the system to confirm that she is not working another AOG at the moment). The TransGulf situation is now fully Nancy’s responsibility. She drops all other jobs to focus solely on resolving this action. The faxed incident details and images, routed to her by Tom, appear on her workstation within moments. She quickly reviews the documents, prints them out, and stamps the packet crimson: “AOG.” Pouring over blueprints, searching databases, consulting bound references and querying local experts, she must come up with an approved corrective action in a matter of minutes. The engineers at TransGulf and at GTS-West know what the problem is, but they need an immediate solution. The plane is still loaded and sitting at the gate.

Nancy knows that any repair they come up with will require Stress analysis for approval. The door jamb is a critical load bearing assembly. She rushes the AOG packet and blueprints over to Felix, the short-haul Stress team lead. “It’s an AOG, we need rapid turn around. I wanted to get this to you ASAP so you can assign it. Can you do it? Great, I’ll write the ROC up now.” Felix did have a Stress engineer, Jonathan, he could pull for the job. He hunches over Jonathan’s desk explaining the damage and prospective repair. Jonathan rushes off to find the stress reports for this model of aircraft, which contain the appropriate load data and standard calculations for the cargo door jamb and related assemblies, while Nancy returns to her desk. On her way back to Structures, she stops by Carl’s desk (the real cargo door guru of the group) to brief him on the problem and warn him that they may need to enlist his expertise. He jokes that this sounds like a great time to take a coffee break – anywhere but here.

Having assembled her team of local experts, all busy working the problem, Nancy calls Takashi for an update. While on the phone she begins to write up the ROC summary for Jonathan. Standard procedure is to write this on-line ROC summary first, and hand off the paper packet second. This job was reversed; with an AOG there just isn’t time. She hits “send” about the time that she hangs up the phone. Still in the ROC system, she switches over to the STAIRS database to search the historical record.

She finds one relevant case – “1991, that’s old.” She glances between her printout of the damage and repair details from TransGulf and this ROC record, comparing the two. It is

not a perfect match, but the depth and angle of the gouge are similar. She prints out the archival ROC summary and rushes it over to Jonathan. This record reminds Jonathan of other metal fatigue considerations with a dent this deep and in this configuration. He hurriedly works through a calculation on the back side of the packet. He taps his pencil on the final number for her to see. “No way, it’s not even close. Bad news for TransGulf, there can be no temporary repair for this damage. It’ll have to be replaced.”

Nancy heads back to her office to call Takashi and fax a formal response to his team via GlobalCOM. Jonathan writes up his rationale in his technical journal (TJ), completes the summary of the job in the on-line ROC file (including a pointer to his journal entry), stacks the packet and blueprint in the corner of his desk to file later, and returns to his interrupted job.

After forty minutes, Global’s answer is in: “Remove and replace.” The damaged plane is not airworthy. The passengers disembark, waiting in the terminal as gate agents scramble to cover their flights. TransGulf loses thousands of dollars an hour as the now out-of-service craft taxis toward the maintenance hangar. Nonetheless, a potential disaster has been averted.

6.6 Additional Complexity

There are many ways in which the standard AOG request just outlined in the scenario above becomes even more complex in even routine situations. There are a host of

players, all integral to the proper performance of exceptional airframe repair, which complicate things further. Three will be discussed here.

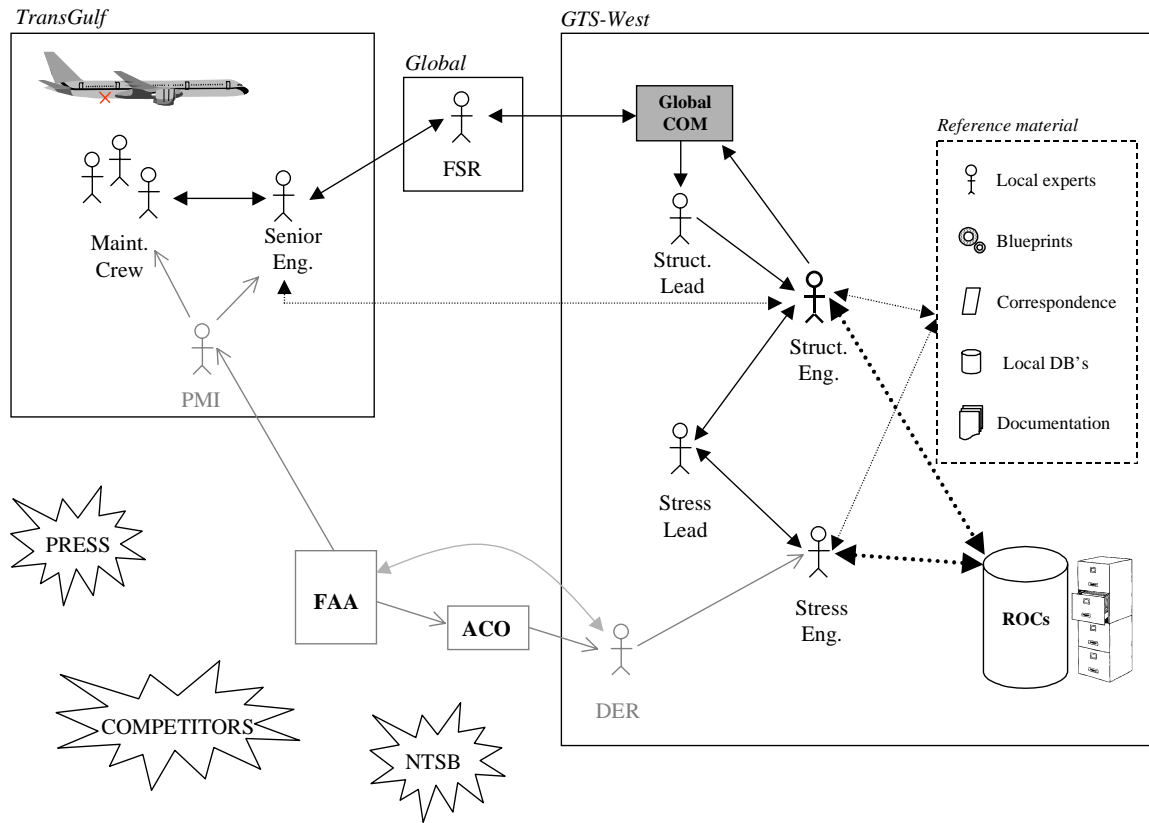


Figure 8: Information flows

6.6.1 FAA

Directly or indirectly, the FAA is involved in every request that passes through GTS-West. There are a tremendous number of generic routine maintenance and repair procedures which have been universally approved by the FAA, such as those provided in the SRM. The majority of these were negotiated with Global at the outset; when a new aircraft is designed and certified it must have complete maintenance and repair procedures. Other universally approved repairs, a significant amount, have evolved over

time, often in relation to legislation, operational incidents or improvements in repair techniques. An operator's maintenance crews are allowed to perform any repair contained in the SRM without requiring additional authorization (though many still call GTS-West for assistance). However, any repair which deviates from the SRM must, by FAA regulation, be routed through an approved FAA certification office. By far the most common venue for this is the manufacturer's service engineering department.

To accommodate this authorization procedure, the FAA has established "airline certification offices" (ACOs) with each American airframe manufacturer. These offices are part of the manufacturers' organization, yet are staffed with FAA representatives who are external to the organization. (The rationale for this bizarre self-policing arrangement is that aircraft design information is highly proprietary and cannot be stored outside of the organization, such as in a central FAA office in Washington, D.C.) One of the primary tasks of Global's ACO is to train, certify and supervise "designated engineering representatives" (DERs) within Global Airframe who can make preliminary FAA judgements. There are five DERs within GTS-West, all in the Stress group.

Repair requests which are minor deviations from the SRM (often cases which are subtle perturbations of a standard repair not covered in the SRM such as two stations to the left of what's covered) require an 8110-3 approval. DERs have the authority to grant this approval on the spot. Depending on the DER's availability (they are usually swamped) this can add significant time to the resolution of a request. In the AOG scenario above, it would have added at least half an hour to the resolution time.

Major repair requests substantially deviating from the SRM, especially requests done against assemblies declared in FAA airworthiness directives (ADs), focusing on trouble prone assemblies, must be routed through the regional FAA office for approval. The DERs coordinate this process of “alternative methods of compliance” (AMOC) approval. It may take anywhere from a few days to a few months.

On the operators’ side the FAA has “principle maintenance inspectors” (PMIs) to ensure reasonable compliance with GTS-West’s repair recommendations and all FAA guidelines. PMI’s are located at all major repair facilities, are semi-autonomous and very powerful – they have the final say on returning an aircraft to service. PMI’s are also fairly idiosyncratic and GTS-West engineers have learned that tailoring a repair to a particular PMI whenever possible will ensure optimal compliance and minimize follow-up requests.

6.6.2 Global Airframe

Other aspects of Global Airframe affect routine jobs such as the scenario described above in two fundamental ways. First, Global has “field service representatives” (FSRs) attached to most major clients (to be located at the discretion of the airline) and operating facilities (such as major airports). The FSR is meant to be a first line of defense, a gatekeeper, and general Global presence. They are trained to understand Global’s policies and procedures and are familiar with Global’s information systems. They work intimately with the repair staffs of the operators and are a superb filter for GTS-West, greatly reducing the number of incoming calls. However, when a problem does arise which

requires the intervention of GTS-West, the FSR must stay in the loop at all times, which is problematic, especially in time-critical AOG situations. In the scenario described above, additional mediation through the FSR would have added additional time and created additional opportunities for miscommunication.

6.6.3 Media, NTSB and public accountability

Air travel enjoys a high degree of public scrutiny. It is more visible and publicly accountable than any other transportation industry in America. Aircraft incidents and accidents are routinely covered by the press, to great effect. One classic example of this is the 1979 crash of a McDonnell Douglas DC-10 outside Chicago. Criticism in the press against the manufacturer was so harsh, and public opinion was resultantly so negative that the fledgling DC-10 program was nearly cancelled. (In the end crash investigators concluded that the primary cause of the accident was improper maintenance, not a design or manufacturing flaw.) All airframe manufacturers live in dread of similar coverage.

The GTS-West team is in a constant, wary state of readiness with regard to the press. They are poignantly aware of the potential damage of bad publicity. During the course of this study I had the unfortunate opportunity to be present for a number of fatal accidents involving aircraft supported by the department. GTS-West is the principle point of contact for all investigations into incidents and accidents. (Within the US these investigations are handled by the National Transportation Safety Board (NTSB).) While there is a small full-time investigation team within GTS-West, members from throughout

the department are routinely called upon to assist. All formal records, such as the ROC, are also subject to seizure and investigation.

The next four chapters will examine the role of the ROC in great detail, with Chapter Seven introducing its role in mediating the collaboration between Structures and Stress, Chapter Eight exploring this interaction in detail, Chapter Nine examining its role in managing inter-organizational boundaries, and Chapter Ten engaging the themes of the ROC in reuse.

CHAPTER 7: Boundary Objects in Collaborative Problem Solving

7.0 Boundary objects revisited

As described in the preceding chapter, information flow mappings are a valuable tool in understanding the routine work flows at GTS-West. While they are effective at foregrounding interaction patterns and are useful for comparing formal and informal communication, they have limited explanatory power. Information flow maps can highlight the channels that information courses through, but cannot speak to the transformation of the information itself. To move this analysis of reuse behavior at GTS-West deeper, the theoretic concept of boundary objects will be revisited and applied.

Boundary objects were introduced in Chapter Two when describing a class of CSCW research engaged with understanding the technical mediation of information use. As a quick reminder, boundary objects are information artifacts “which both inhabit several intersecting social worlds and satisfy the informational requirements of each of them. Boundary objects are objects which as plastic enough to adapt to local needs and the constraints of the several parties employing them, yet robust enough to maintain a common identity across sites” (Star and Griesemer, 1989, p. 393).

The information environment at GTS-West is rife with boundary object candidates. These range from such large grain concepts as the department itself (mediating meaning and action between Global and its customers) or an entire aircraft (the locus of discussion

over the appropriateness of a repair), to such fine grain concepts as a single measurement or a damage photograph. An examination of each, working in concert, will yield a richer understanding of the successful collaborative problem solving which occurs at GTS-West. However, given the bounds of this dissertation document, rather than superficially inventorying this vast collection of illustrative negotiation points, it would be best to focus on a single class of boundary object. This representative should be the one that is the most central to the routine daily work practice and ideally the most amenable to later system redesign.

7.1 The primacy of the ROC as boundary object

The Record-of-Conversation (ROC)³⁰ is clearly the most central, critical, and complex boundary object to be found in routine work at GTS-West. It mediates nearly every action item that is resolved between Stress and Structures. As will be described at length in the following chapters, the ROC also mediates among unexpected parties as well. In most cases, nearly the entire cast of characters presented in the TransGulf example from Chapter Six, negotiates through the ROC. This includes such individuals as the airline's repair technicians, Global field service representatives, and the FAA. Some negotiations are explicit, while others are implicit, shifting on a case-by-base basis. The boundary object approach assists in mapping out these myriad relationships which influence this

artifact over time. It highlights the varied perspectives that are brought to bear when interpreting, negotiating, or modifying a ROC instance.

Returning to the notion of granularity in boundary object analysis, there is significant fluidity in the concept. In that inaugural paper, both the state of California and a tagged bobcat femur bone acted as boundary objects. As noted above, the same spectrum is evident at GTS-West. In examining the granularity of the ROC, however, differences of kind come to the surface. Yes, it exists simultaneously as an ideal type in its abstraction and a standardized form in each instantiation, but it is also more than this.

In certain situations the ROC really functions as a meta-boundary object or boundary object repository, with its component parts as the true boundary objects. In the processes of creation and modification very little negotiation happens at the granularity level of the ROC itself. Instead it happens among its internal elements: arguing over a particular measurement, questioning the interpretations of particular individuals, or wrangling over alternative repair scenarios. In contrast, however, during the processes of reuse, especially recontextualization and evaluation, the negotiations across time do occur at the

³⁰ The ROC system was first introduced in Chapter Five and discussed within the context of work in Chapter Six. For the purposes of discussion in the remaining chapters of this dissertation the term “ROC” will refer to the abstraction of the complete action item record. This encompasses both the workflow and archival on-line entity, the two page summary, and the complete physical packet. Using the term in this fashion, to describe the entire amalgam, is the most common linguistic use by the GTS-West engineers. They only split it out into its requisite parts is when necessary (e.g. the computer system is down or the physical file is missing). The presentation here will conform to this, viewing the ROC as a conceptual whole, only breaking out its component parts when necessary for analysis.

granularity level of the ROC. The entire entity is evaluated for completeness, authority, and appropriate fit.³¹

In the remainder of this chapter the ROC will be introduced as a boundary object. First, an illustrative example from one of the cases will be presented, grounding the notion in an actual work situation. This example will then be unpacked to highlight some of the more interesting transformations. Lastly a boundary object diagramming technique will be presented which will be used in the following chapters of analysis.

7.2 ROC as boundary object illustrated

The general, high-level information flow for this example is nearly identical to that presented in the introductory example in Chapter Six. The action item followed the standard path through GlobalCOM, through team lead assignment, and the hand-off between Structures and Stress. Therefore those details, while present in the diagram, will not be repeated here. The key difference to note though between the examples is that the Chapter Six example was a high priority AOG case, while this was a low priority job with a 21-day completion time frame.

³¹ Considering that the primary research interest of this research is to better understand the behaviors of reuse, the ROC will be routinely referred to at this higher level of granularity. While the distinctions will be called out as appropriate, the reader should always bear in mind this dual nature of the ROC as both boundary object and meta-boundary object.

7.2.1 An example

Zip Express (ZE), a regional US carrier, has submitted a request regarding wear limits on the assembly below the pivot hole for the outbound forward hinge on the inboard flap for seven specific aircraft in its fleet. Once the wear limits were determined, ZE was seeking FAA certification for the numbers.

Cameron, the Structures engineer assigned to the job, spent six days preparing the case for the hand off to Stress. First, he had two cycles with the operator over the phone to try and clarify their request and elicit missing information. Then he collected key statistics about each of the seven aircraft, including the details in the ROC packet which was handed off to Anita, the assigned Stress analyst.

Due to its prioritization, Anita placed the job on the backburner until immediately prior to its due date. After researching prior cases, she found that the ZE's requested limits were actually within designed tolerances. There was no job here. Anita had Cameron call ZE back to verify this unusual request. Cameron was unable to contact the responsible party at ZE and the job was closed out.

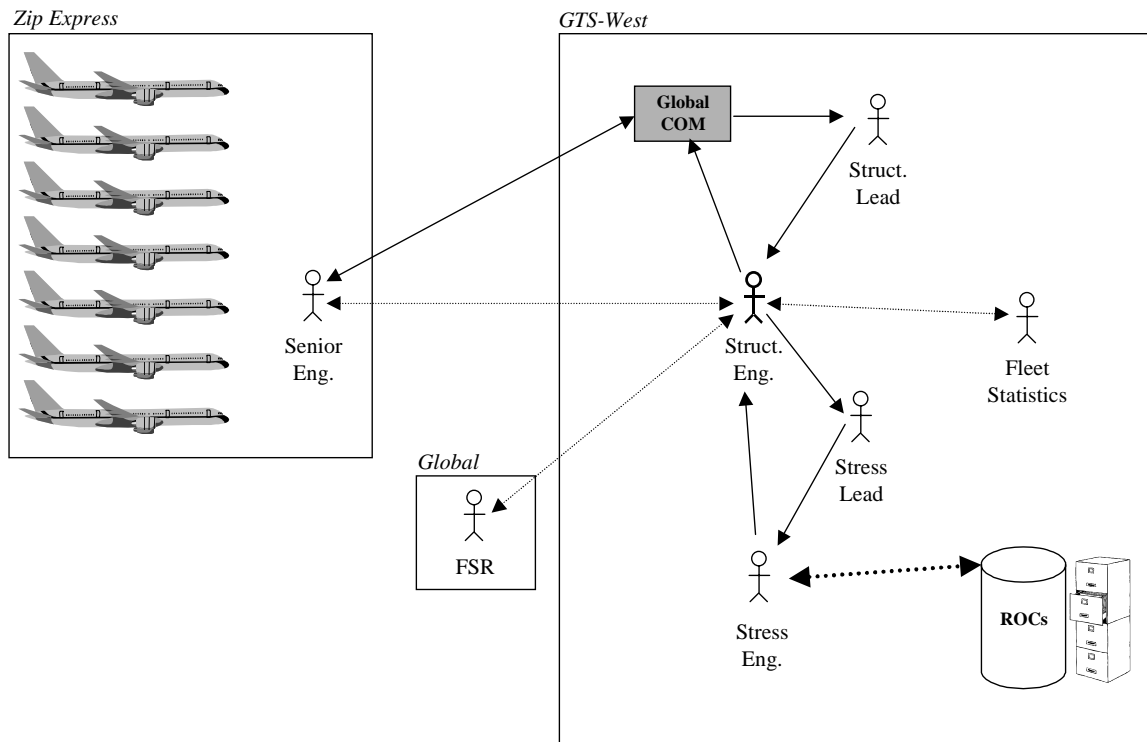


Figure 9: Information flows

7.2.2 An example unpacked

When the job first arrived Cameron realized that this would be a substantial task for Stress, because “we really don’t have numbers for this. Global doesn’t give [wear] limits [for this part].” The wear limit information, pertaining to acceptable tolerances beyond blueprint specification, is not provided in either the SRM or the Component Maintenance Manual (CMM). He explained that each prior case was evaluated individually and independently.

While it is the routine practice of service engineers to cut and paste the discrepancy from the GlobalCOM request into the newly created ROC with minor modification, Cameron

decided to do some advanced investigative work to ensure that the job was properly described and to assist Stress in its analysis. Using data gathered from Fleet Statistics, the customer service group within GTS-West responsible for maintaining data on each individual fuselage which is reported back from the operators, he built up profiles of the relevant statistics for each of the seven aircraft. His stated rationale for these actions was, “they have these seven planes, but they only provided the registration number. I needed more. See, I have this [mini-profile] for each, so I’ll have it later on. I won’t have to keep looking it up. [The] fuse numbers and all. These are converted freighters you know.” (The latter is a particularly interesting discovery that he made in this process. A reverse conversion is relatively rare with passenger aircraft and the information may prove useful later on.) In addition to this task, he also had to contact ZE twice to clarify certain elements and to collect information omitted from their original request.

The results of all of this communication and research Cameron kept written up on his “cover sheet.” This is a local memory device which he developed to help track his jobs. It is a self-designed form containing some of the basic information present on the ROC summary, along with a good six by eight inch white space for hand written notes. While his particular form is unique in Structures, the practice of maintaining a personal store of informal information on current jobs is not.

Cameron’s assumption about this being a daunting job for Stress was correct. Anita offered this as the reason why she delayed action on the job, “this is going to be big. A big job. It’s going to take days. I’m trying to get all the easy ones out of the way first.”

She appreciated the detailed information that was provided for each of the seven aircraft, but was unaware that Cameron had culled these details himself. While Cameron had anticipated these profiles to be the key information for Anita's job, Anita noted with exasperation that it was only a starting place, "first this, it's supposed to be just one part, but this has all these parts. You have to look at everything, and all that, for all these! Just finding all this stuff..." Anita realized that calculating the wear limits for this part depended on interactions with other local sub-assemblies on the inboard flap and would have to collect data on all of them, potentially for each aircraft, to make an appropriate assessment.

Knowing, as Cameron did, that these wear limits are traditionally decided on a case-by-case basis, Anita searched through the archival ROC system for similar cases that might give her a sense of the range of precedents.

"I searched. I was looking through the old records, jobs we've done before, all the old ROCs. I searched for the part to get the history. A lot of repairs referenced the CMM. They [operator] would say 'the limit is this, and we're this much over.' [They] were looking for approval."

With so many prior cases referencing the CMM for limit information, she researched it as well. While repair limits are not specified in the manual, the CMM does list design limits. ZE's request for each aircraft was within these limits.

At first, she was relieved with the discovery, "all seven [aircraft] were different, but they had this absolute. We had to verify [and that was it]." She did not have to do the extensive analysis for each of the seven aircraft. However, she soon became puzzled as to why the airline would have submitted such an obvious non-request. Clearly they have

their manuals and could have looked it up, and it most certainly did not require special FAA approval. She had Cameron call ZE back to ensure that GTS-West had correctly understood their request. “I have no idea why they sent this job in the first place, [so] that’s why we called.”

Cameron could not contact the ZE engineer who had initiated the request, “we called them and they never got back to us.” Undaunted, he called the Global field service rep that services the ZE repair facility, but this attempt was unfruitful as well, “he passed me off to someone else who passed me off to someone else. Not helpful... you’d think he’d know [whom to contact within ZE].” Cameron ultimately concluded that “there’s no one responsible for this [at ZE]” and closed the job out with a terse reply to read the manual and contact GTS-West again if there is a problem.

Cameron noted that this kind of pointless request is surprisingly common, “yeah, we get these all the time. They [operators] just submit. It’s [a repair already covered somewhere] but they ask for approval anyway... often because they just can’t find it [their manual and it’s easier to call].” Cameron stored his now densely scribbled coversheet in his desk, noting that he would not be surprised to see this job come around again. Meanwhile, the Stress team lead signed off on the ROC as resolved and submitted it to the archive.

7.2.3 Boundary object diagramming technique

To help visualize the various negotiations which utilize a boundary object such as the ROC, a diagramming technique was developed for this study. It will be used throughout

the remaining chapters, so will be briefly described here and illustrated with the preceding case-based example.

The particular instance of the boundary object, in this case ROC #000610049, is represented by the central, circular hub. Spokes protruding from the hub represent the various parties whose negotiations center on this boundary object instance. The length of each spoke represents the relative degree of influence on the boundary object instance that each participant has. This is roughly an approximate power weighting. For example, both Cameron and Anita had a high degree of influence on the ROC while others, such as the field service representative, the fleet statistics team, or the design engineers who authored the CMM, had less immediate influence. A dashed line indicates an indirect or peripheral influence (e.g. no data provided by this entity is directly present in the ROC). The resulting diagram for this example is:

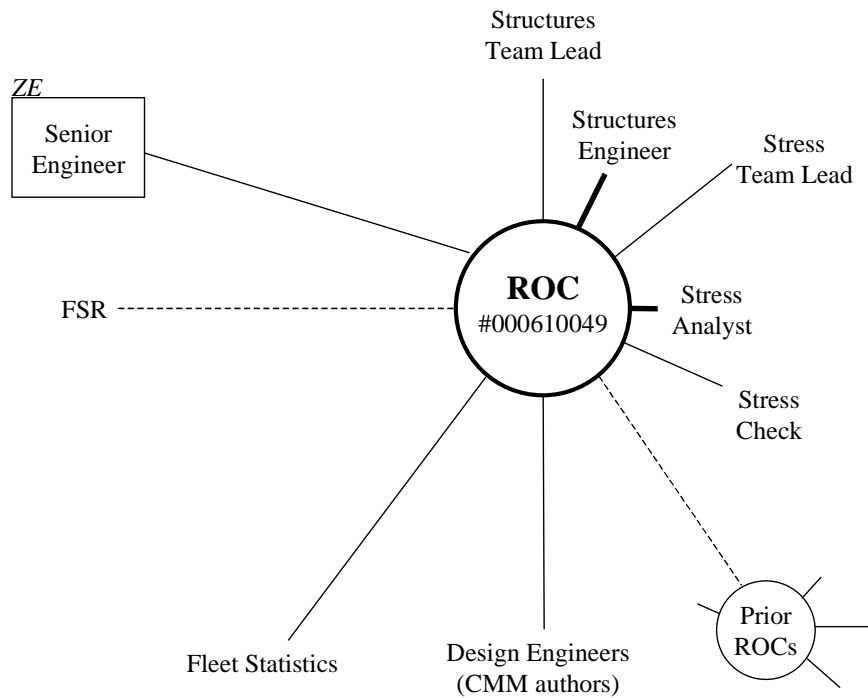


Figure 10: Boundary object diagram

Having introduced the notion of the ROC as a boundary object and illustrated it through a case-based example, the foundation is now prepared for more substantive analysis. The following three chapters will examine the ROC as boundary object concept in detail, focusing on the intra-organizational mediating effects, inter-organizational mediating effects, and its use across time.

CHAPTER 8: Coordination and Collaboration, Inter-group Boundary Spanning

8.0 Embeddedness of boundary objects in routine work practice

This chapter focuses on the manner in which the ROC, because it is a boundary object, mediates the coordination, collaboration, and negotiation between Stress and Structures. This will be critical when reusing the ROC or changing work processes.

The preceding chapter introduced the ROC as boundary object through an unpacking of a case involving a routine job action resolution. This chapter will deepen that analysis by concentrating on the ROC's role in negotiating completeness, facilitating multiple presentations of job resolutions, and supporting a single point of control within a set of unfolding, complex, multi-way negotiations.

8.1 Revisiting the nature of the Stress / Structures boundary

Recall from the discussions in Chapters Five and Six, that the engineers of Structures and Stress represent two distinct communities of practice. Their organizational cultural differences are substantial. In the business process re-engineering meetings, the two groups distilled the core distinction between themselves as follows. The service engineers in Structures are dedicated to their goal of providing first class customer support and as such they primarily prioritize their work according to "timeliness," which they define as getting an accurate response back to the customer within their requested timeframe. In

contrast, the engineers in Stress are dedicated to their mission of providing highly reliable stress analyses for submitted repairs, and as such they primarily prioritize based on “quality,” which they define as having as complete as possible information in order to perform the most thorough analysis to be able to offer the customer the best possible repair solution. Both groups are invested in the ultimate goals of aviation safety and customer satisfaction, but approach them through differing and often orthogonal means. This core tension permeates all of their interactions and will be a constant theme throughout this chapter.

In the face of these differences, Structures and Stress must routinely work together to resolve over eighty percent of all incoming job requests. Their collaborative problem solving activity is largely face-to-face, ephemeral, and unmediated by information technology. They frequently spend time at each other’s desks (mostly Structures consulting with Stress), talking through their various concerns, and negotiating to shared understandings.

The primary information artifact in, and often only permanent record of, these activities is the ROC. In essence, the ROC is the physical embodiment of an action item and ownership of that record represents the current responsibility and authority for that action. While most action item discussions wend their way through a diverse web of resources (e.g. blueprints, repair manuals, prior cases, FAA regulations), the ROC remains the constant focal point which participants return to again and again.

Every new stress analysis request by Structures is initiated with the creation of the two ROC components, the digital summary and the paper packet, and the submission of these to the appropriate Stress team lead. Every such request is finalized with a flurry of signatures on the packet and a reassignment of the on-line summary to the archive holding account.

The ROC is manifest throughout this process from creation to resolution. It is clutched in hand running down the hall. It is thumped on the desk. It languishes in the inbox. It is furiously annotated in a desktop debate. It is slipped into the brilliant red 'AOG' plastic sleeve. It is endlessly reshuffled for prioritization. It is passed from desk to desk. It is what fills a hundred filing cabinets lining the halls.

In the remainder of this chapter the ROC will be examined precisely for the properties which enable it to operate as the glue in these coordination and collaboration activities between Structures and Stress.

8.2 The ROC as foundational inter-group facilitation mechanism

The ROC is the central component in building and maintaining shared perspectives between the two groups. It provides a common ground for understanding about a current job, and augments this with a shared history. A request for Stress assistance on an action item is not official until a ROC has been generated. Surprisingly, this holds true informally as well. Beyond legal policy requirements, the ROC is the irreducibly minimal information artifact for a substantive interaction between the two teams. While

preliminary discussions regarding a job do happen with great regularity, the engineers universally (and rather quickly) reach a point where a shared artifact is required for further collaboration. Instead of grabbing for the proverbial napkin or scratch paper, the engineers will call a time-out, generate the basic ROC, and resume.

There are numerous exchanges which highlight this time-out behavior. In one instance, a Structures engineer introduces the problem to the Stress team lead. They have a brief, initial discussion regarding potential approaches to the analysis and repair. Satisfied that there is initial mutual understanding, and stymied by a few detailed questions, the structures engineer breaks it off and starts walking back to his desk, calling over his shoulder: “I’ll do you a ROC and sent it on over. They’re waiting for the answer. No, don’t [you] expect it in thirty seconds, but soon. I’ll do the ROC right now.”

Instances of formal Stress consultation on an action item not accompanied by a ROC are extremely rare. (There is a single case in this study where a stress analyst suggests that the structures engineer talk the operator through the issues using a shared photograph.) However, the tendency to err on the side of creating the ROC, does mean that there are many dangling requests in the workflow system which are incomplete, unanswered, duplicate, or simply useless.³² Even more frequent is the tendency not to include additions, corrections, or fixes to the ROC for reasons of the reliability of the record and deliberate, selective omission (each of which are discussed below under the issue of negotiated control).

By social convention, the ROCs for currently open jobs are always to remain publicly accessible. It is a norm in both groups that these are always left out on the desktop, in full view. (I did not encounter a single engineer in either group that placed current ROCs in their desk or personal filing cabinet.) This visibility provides for a fundamental awareness mechanism. A quick walk down the hall, surveying the desks, yields a wealth of information³³ – everything from the number of open jobs an engineer is working, to the complexity of those jobs, to their prioritization on that engineer’s schedule.

The ROC packet accessibility also facilitates team coverage. When an engineer is unexpectedly absent and one of her jobs is nearing its due date, the team lead can reassign the job electronically and the new engineer can collect the ROC packet from the absentee’s desk, reviewing personal notes, and surveying the resources they had assembled for the job.

The ROC also functions as the primary data exchange medium between the Structures and Stress. It is one of the few shared representational formats that both groups agree upon. This agreement is mostly a byproduct of historical inertia, as the data layout is actually not optimized for either group. However, the common fields and spatial layout of the ROC summary pages have a habituated familiarity for the engineers and provide a central synopsis of data translated from a wide variety of sources. The skeleton of the

³² Note though that these dangling records rarely achieve archival status. One of the tasks of the team lead is to evaluate each of these and filter appropriately.

³³ Most of the engineers for both groups engage in this informal scanning activity multiple times per day.

summary fields also provides for a minimum information threshold on any given job where omissions can be easily noted.

Lastly, the ROC plays a central role in the scheduling of time, resources, and engineers. As described in Chapter Five, the ROC system is a rudimentary workflow management tool. It supports the assignment, and reassignment, of jobs among the handful of engineers required to bring it to completion. It also provides basic awareness about the status of any given job. Has it already been assigned? To whom? Which stages in the four step process have been completed?

In understanding the role of the ROC as a mechanism for facilitating the daily collaborative acts between Structures and Stress, it is necessary to examine some key issues in micro-coordination which will be unfolded in the following sections. The first of these will be the ROC's role in representing the negotiations around the notion of completeness.

8.3 Completeness

The primary problem Stress has with any given ROC is its completeness, meaning that the record contains all of the relevant data, in the correct format, required to carry out the requested analyses. From their perspective the vast majority of ROCs are incomplete when assigned from Structures. One of the fundamental collaborative activities between the two groups is advancing the ROC to a stage of "complete enough" for the proper

analyses to be accomplished. There is no fixed definition for this state of agreed upon “complete enough”; it is a negotiated order (Strauss, 1993).

The root cause of ROC incompleteness almost always lies with the operator, who has failed to describe their request in sufficient detail or in the correct format. (The GTS-West / operator negotiation with regard to the completeness of requests will be discussed in detail in the following chapter.) Both Structures and Stress agree upon the culpability of the operator, however the resolution of this problem foregrounds the differences in organizational culture between the two groups. These differences must be successfully negotiated to each side’s satisfaction for the successful resolution of the job.

From the Structures standpoint, Stress analysts appear capricious, fickle, and obsessed with finessing the data instead of getting jobs out on time. The analysts will rarely settle for “good enough.” (Structures is also more of a puzzle-solving culture which finds challenge in limited information, not crisis as Stress posits.) In frustration one engineer commented, ““That’s the crunch. They [Structures] can work their asses off, be here 24 hours, but if they [Stress] don’t sign off they are screwed. What’s worse, the stress guy is probably just playing on the Internet.”

Soliciting the missing information from the operator often requires multiple call-backs from Structures. This greatly concerns the service engineers who are actively engaged in face management for the department. Calling operators days after their initial request to ask for more information makes the GTS-West look less competent. It leaves the operator

wondering why they were not asked this information up front, and why their job is apparently coming so late in the queue. These are both issues with which the operator will heartily harass the service engineer.

From the Stress perspective, Structures is so concerned about not questioning, second guessing, or disturbing the operator that they are willing to settle with sub-standard information. The average analyst does not understand the complexity of the service engineer's job, and one commented in frustration over the delay in a ROC creation, "they get the request, copy it verbatim, and hand it off."

The tension in these divergent perspectives on completeness was one of the primary issues in the business process re-engineering effort (which concerned decreasing call turn-around time). Stress complained vehemently of having to deal with incomplete or missing information on ROCs. Having complete information for every job, they insisted, would be the best way to maximize efficiency. Reducing the amount of time spent in this negotiation would allow for quicker turnaround.

Structures countered that Stress really does not know what they need, that not every element of data is important in every case, and that collecting all possible data for every job was simply too onerous of a task to impose on the operators. Stress must know what really matters for any given case prior to Structures eliciting that information. This eked out a reluctant admission from the Stress representatives that they really do not know everything that they will need until they actually begin an analysis. A compromise was

reached, which does not eliminate this negotiation, but seeks to reduce the number of iterations required to get the information that is needed.

Some high-level standards for facilitating this negotiation already do exist, which cover general information that should always be collected regardless of its intended use. A prior Stress manager had worked up a set of guidelines for all Structures engineers to follow for every job, listing what should and should not be present in every ROC. However, many engineers did not even know these existed, let alone followed them. There have also been attempts made at training the operators themselves (which are described in the next two chapters).

What is universally agreed upon is that each side should be timely in their requests for additional information. The official position on this can be found in the new employee operations manual for Stress analysts:

“Within a reasonable time after receiving a ROC, analyst checks it for completeness of data. The Cognizant Customer Service Engineer (CCSE) should be contacted immediately for getting the missing information. A ROC may be closed with the agreement of the CCSE, if the operator / repair shop is unable to provide the missing data in time.”

In general, Structures is very helpful in collecting any additional required information for Stress, despite their frustration and the awkwardness of initiating multiple call cycles with the operator. It is in their best interest that the jobs are processed as quickly as possible. Frequently they will try and predict what unusual information may be helpful for Stress and collect that in advance of creating the ROC.

One example of this is represented in the case from Chapter Seven. Recall that the operator, Zip Express, had a very general repair request which impacted a number of aircraft in their fleet. The Structures engineer, Cameron, followed up with them to make this general question more specific. Finding that this involved seven aircraft, he used the fuselage registration numbers to collect all of the background information on each individual aircraft before handing the job off to Stress. Finding all of these details in the ROC, the analyst was elated. Multiple cycles with ZE to elicit the details on each individual plane were avoided and a potentially onerous multi-day job was able to be completed in a matter of hours.

Likewise, Stress is also very willing to work with Structures engineers to help them understand some basic techniques to evaluate the quality of information required by Stress. In the following training example, a stress team lead helps train a relatively novice Structures engineer, even though there is a line of people waiting to talk to him.

Structures: “Just a heads up. This is what’s coming down the pipe next... This is the best I can get out of them [repair details]. Is it enough?”

Stress: “Show me the crack.”

The service engineer is unable to do so. By putting him on the spot to explain this simple yet critical detail the team lead helps him realize that the information is not even complete enough for himself. The team lead continues to discuss the importance of orientation (forward/aft, inboard/outboard) for the crack. Patiently he reviews all of these processes only to be rebuffed by the engineer who is nervously aware of the time.

Structures: “It’s not really clear, is it? Just say it’s not clear...”

The tension over completeness underpins the social control structures surrounding the ROC which will be discussed next.

8.4 Social control structures for the ROC

Given that the ROC is the central, and often only, information artifact which mediates between Structures and Stress for any given action, it is not surprising that tension of control is frequently paramount. At the root of this tension are three distinct information needs, the need for reliability, for customer presentation, and for historic preservation. Each will be discussed in turn.

8.4.1 Need for reliability

Recall that the “Record-of-Conversation” is just that, an accurate account of the resolution of every operator request serviced by the department. The FAA requires this legal paper trail to ensure that due process was followed and that the individuals responsible for each step in the process are clearly identified. As a result, modifications to this record must be carefully monitored.

By design once a ROC is released by one account and sent to another, it cannot be recalled or edited by the prior engineer. (The reader is probably familiar with similar frustrations with e-mail messages.) Considering the formal workflow this makes perfect sense. Jobs are only handed off when they have been completed by the party at hand. Legal responsible is cleanly transferred at that hand-off.

However, in reality action item resolution is a very dynamic process. For example, the Structures engineer usually has a stream of information coming in from the operator, rather than a discrete, one-time information dump. The ROC they create to represent this job is a snapshot of that information stream and what they believe to be the optimal time; the ROC will be complete enough to allow Stress to perform their analysis and close out the action by the requested due date. This is a fairly complex timing process for which service engineers are very adept and seem to master quickly. However, sometimes there are show-stoppers, e.g. the introduction of new information which may radically shift the understanding of the repair request. Depending on the degree of severity, the service engineer may close out the current ROC and start a new, replacement ROC, or request that the Stress team lead edit in a change, or simply pass off the information word of mouth to the analyst working the job. This decision is always made in consultation with the appropriate Stress team lead as they will know how it will be assigned (a more senior analyst might be okay with verbal instruction while for a more junior analyst it would be safer to ensure that it is included in the record), how the jobs are being evaluated (whether this is worth the performance ding of closing out an unresolved ROC), and also have a sense of how this may play out in future reuse.

This raises the issue at the core of the next control concern.

8.4.2 Need for control over archival process by Stress

Stress maintains that they need control over the ROC processes for the sake of historical reuse. The ability to control the processes of decontextualization for individual ROCs to

best fit their expected trajectory of future reuse is critical. As the most likely future audience for reusing cases will be Stress analysts, it is logical to allow them some degree of control over these processes. This control would be funneled through the organizational role of the team lead, described in great detail below.

8.4.3 Need for control over customer presentation by Structures

Structures has a very different, but no less important, need to control the ROC. There is measurable variability in the analytic services provided by Stress. (The methods used by Structures in managing this are described in the next major section.) In black-boxing these activities for the operator the service engineers must maintain the appearance of standardized request resolution processes.

The ROC information required to perform these simplification and face maintenance tasks for the customer are directly at odds with the information required by Stress to maintain a robust and relevant historical archive and the information needs for the ROC to serve as the formal regulatory account of these transactions. The tension produced by the confluence of these three factors appears insurmountable.

8.4.4 Resolution achieved through split views

GTS-West has found a method to relieve, if not completely resolve these tensions. The resolution preserves a consistent interpretation of the ROC and allows it to be maintained as a boundary object. This has been achieved via a set of organizational arrangements. With regard to the regulatory requirements, there are explicit locking mechanisms and

hand-off procedures built into the system. Record administration is divided by aircraft model family and is provided by a single point of contact, that model family's Stress team lead.

The resolution of the tension between the needs of Stress and Structures is more complex. Basically there are two different views presented of the same ROC boundary object.³⁴

Addressing the needs for maintaining a Stress-centric ROC archive, the ROC system is administered by Stress. Control has been institutionalized through the role of the team lead (which is examined in greater detail later in this chapter).

Stress team leads are the focal points, not only for the work flows represented in the system but for the content as well. In their role as aircraft model ROC account manager, they control more than the job assignment processes, they are also responsible for the ROC quality control elements: the check/verification process, the editing process, the formal close outs, and the batch archival process. These dimensions of control are critical for preserving the record for later reuse.

Addressing Structure's requirements for appropriately packaging the response for customer, they are given final editing privileges for ultimate response on an action item to the operator. The local GlobalCOM is primarily administered by Structures. As the

final response is provided to the operator via GlobalCOM, not the ROC system, it is during this process of transference that the service engineers make their modifications.³⁵

8.5 Managing cross-group expertise

One of the primary peeves of the service engineers is the variability of response they receive from Stress. As mentioned above, as customer support representatives they have an organizational requirement to present the appearance of a unified system to their customers. However, this conflicts directly with the nature of the analyses Stress undertakes. As with any complex engineering task, there is a degree of art balancing science. This allows for some measurable personal differences between analysts, not only in style and presentation but also content and approval. (In the most dire, and rare, circumstance one engineer may approve a given repair based on a particular set of stress calculations, while another would deem the same repair unsuitable for strength using a different approach.)

The service engineers' general frustration over this issue is aptly captured in this brief exchange. For staff load balancing reasons, a Stress team lead approaches the cognizant service engineer for a particular job and asks "Mind if I reassign this?" The engineer tentatively replies, tinged with sarcasm, "[the] answer is going to be the same whomever, right? Sure."

³⁴ It is difficult to assess at this point whether the ROC representing the actions is a single boundary object with two views or whether it is cloven into two boundary objects (one representing the account provided to the customer and one stored in the archive). While this has received extensive review by the author, it currently remains an open question.

As customer support representatives, Structures engineers desire the most generous response on a given job for their operator. To help ensure this, they often attempt to influence or control the ROC assignment process in order that their jobs are assigned to the Stress analyst that they deem is best suited for the task. Subtle ways they indicate their preference are to make suggestions to the Stress team lead or highlight an analyst's name on an included reference ROC. More overt ways are to engage the analyst directly to persuade them to request the job from their team lead or to try and commandeer the analyst directly for a particularly hot job.

This behavior is clearly organizationally problematic. Routing all jobs through a handful of exceptional analysts causes paralyzing bottlenecks. Jockeying for one's favorite analyst erodes the sense of team identity, promotes dissention, and limits the training potential for new employees. Most importantly it damages the sense of organizational autonomy for Stress. As a result, team leads actively discourage this behavior, though they acknowledge that in many cases, particularly high priority or continuing jobs, having a good match is indeed important for successful resolution.

Three of the most common motivations for service engineers to attempt to leverage particular analysts are described below.

³⁵ These changes are made almost exclusively without consultation with Stress. The cognizant Stress analyst is almost always unaware of the form of the final response.

8.5.1 Job continuation

One of the least problematic areas for Structure's requesting particular analysts is on continuing or follow-on jobs. When routed correctly through the team lead these requests are welcome and attempts are made to accommodate as best as possible given the current load and schedule among the analysts.

In the following example, as Structures engineer is addressing a particular Stress analyst directly. "Did you just answer a corrosion issue for EVI? I thought you did. You're the one. Well, they're back with some problems."

8.5.2 Perceived expertise

Service engineers selecting analysts based on perceived expertise is much more problematic. Clearly one cannot overly play favorites, so the service engineers frequently rationalize their requests with either common knowledge ('everyone knows that Philip is the expert on nose cones...'), referencing prior collaboration ('I worked on one like this with Philip last time...'), or substantiating it with ROC references ('I searched STAIRS and it looks as though Philip is the man for this...').³⁶

This behavior is illustrated in the following example. After preparing a ROC for a new job a Structures engineer set off to search the floor for the particular analyst he wants assigned to the job, Jim. He does not find Jim at his desk, but does run into him in the

³⁶ Although Stress discourages the behavior, the shared, archival ROC repository does allow for evidentiary comparison between analysts.

hall. He briefly describes the job and mentions that he thought that Jim had worked on a similar problem a few weeks back. [Even though he is holding a copy of that particular job in his hands, he is deliberately vague allowing Jim to recall the job on his own.] Jim remembers the job and is willing to work this new job. The engineer mentions that he has attached the old job as a reference ROC for the new job. Wary of this request being outside the normal channels, Jim asks if the engineer had submitted the job as “attn: Jim.” He mentions that he had not. Jim then instructs him that he must do so on the ROC for the job to be assigned to any particular analyst. He adds that it must be routed through the team lead, as he does all routing for the team, and that there was no guarantee that the assignment would be made.

8.5.3 Compliance with historical precedent

The final motivation for attempts at analyst selection deals with consistency. One of the most frequent complaints from operators to Structures on the content of their responses involves precedent (‘You let us do it this way a month ago, and now you say we can’t do it? Why?’)

It is the formal policy within Stress to be aware of all prior similar cases, and try as best as possible to maintain consistency. However, in practice this is very hard to accomplish. Few jobs ever match identically, there is simply too much variability in aircraft repair (especially when considering the diverse maintenance conditions of each individual aircraft). So analysts try to maintain consistency on a high level (e.g. process and heuristics), but acknowledge significant variability on a low level (e.g. repair details).

This is a distinction that is not at all obvious to the operator who, given limited information in their reply from service engineering, usually is only focused on the “approved” or “denied” status of their request.

8.6 Impact of ROC Visibility

The ROC system is a closed system whose contents are only visible to those within GTS-West. This has a number of benefits and disadvantages regarding both content visibility and process visibility.

The fact that the ROC, as a boundary object, is only visible by a subset of those between whom it mediates, and is only modifiable by an even smaller subset, instantiates a hierarchy of power over the repair request resolution process. Stress team leads have the highest level of influence, while the actual repair teams at the operator’s station have the lowest. (The role of the Stress team leads will be discussed in the next major section.)

A benefit of having the ROC system closed is that the information it contains can be slightly more honest and direct, knowing that the final response will be filtered by the service engineers in their final GlobalCOM reply. (The record still needs to be crisply professional though, recalling that this is the FAA mandated legal paper trail.) Informal information is indirectly evidenced by the routine peppering of the ROC with local knowledge and nomenclature (e.g. resources, individuals, acronyms). It is explicitly evident in the occasional still-in-process response. The following example highlights this behavior.

A Structures engineer is discussing a repair drawing with a Stress team lead. At issue is the appropriate fastener type for the repair – which class of aluminum or titanium is most appropriate for this doubler. The confusion is deepened by a discovered discrepancy between the SRM and the original design blueprint. The team lead agrees to investigate the problem, run the analyses, and finalize the decision. The Structures engineer thanks him, but stresses the time critical nature of this as he leaves. When the team lead suggests that he cannot formulate a definitive answer in time allotted, the Structures engineer suggests that he close out the ROC with his best formed opinion, explaining that “[you] don’t need it explicit. This is not direct for the airline.”

Regarding content visibility, both groups are satisfied with the current arrangement. (Although, other forces are at work attempting to increase the external visibility of the ROC system. These will be discussed in the following chapter.) However, disagreement arises between Structures and Stress with dealing with process visibility.

The introduction of GlobalCOM in Structures has revolutionized the visibility of content, communication, and workflow from their prior practices. After a period of initial discomfort, the engineers have found this new level of awareness to be extremely valuable. (While the boundaries of the GlobalCOM system extend far beyond the organizational boundaries of GTS-West, most of the workflow monitoring features are bounded at the group level. The ability of FSR’s or operators to monitor job assignments and progress would rankle nearly everyone in the department.)

As a result service engineers are staunch advocates for increased process visibility throughout the department, and see it as especially helpful in managing their relations with Stress. Stress is exceedingly reluctant to shift the current level of visibility offered with the ROC system. Their chief concern is in regards to autonomy. The work of the Stress analyst is idiosyncratic, and there is fear that opening this process up to those outside of their community of practice would lead to significant misunderstanding and mismanagement. Their fear is justified as the following exchange illustrates.

A senior Structure engineer approaches a Stress analyst with a follow-up job to one she had worked that morning. Already intensely overloaded, she rather coldly directs him to talk to the team lead. The team lead is visibly upset that the engineer had brought the ROC to the analyst directly and demands a reason for the behavior. Frustrated, the service engineer mentions that he knows there is some leeway in the Stress workload because of some recently completed jobs, “I heard that they have just finished, so they’re wrapping up. They did that in no time, so now they can certainly take this...” Even if this particular analyst could not accommodate him, he was confident that some of his work could be offloaded to another, now free, analyst. The team lead bristles at the assumption that the engineer would understand the schedule and prioritization of his team, takes the ROC from his hands, and places it firmly behind all others in his desktop assignment queue.

This exchange illustrates precisely why Stress does not want their process opened up any further. Undermining the control of the team lead, even by only adding increased

awareness, could unravel the fabric of the group. (However, observing the intense strain placed on the team leads of both groups, this deliberate bottleneck will not scale well in the future and the organizational structure will have to be re-evaluated at some point.)

8.7 The interdependence of the Stress Team Lead and the ROC

To fully understand how the ROC assists in the tug-of-war management of time and resources, it is also valuable to revisit the role of the team lead. Recall from Chapter Five, that within both Stress and Structures a senior engineer has been assigned to coordinate efforts of each aircraft model family-based team. While this is a non-management position, the team lead is responsible for nearly all moment-to-moment staffing decisions. They are also routinely viewed as the most knowledgeable individual in the group by their peers. They maintain a rich personal knowledge store which expands far beyond the technical details of their job to include extensive knowledge about the skills and expertise of the individuals on the team,³⁷ the expectations of the operators, and the inner workings of the group with which they are collaborating.³⁸ They have such extensive experience with the ebb and flow of this process that they are the most adept at making the best predictions about the future, essential for managing the allocation of resources in the moment.³⁹

³⁷ The team leads' abilities and behavior in this capacity are strikingly similar to those of the information concierges at MSC as described by McDonald and Ackerman (1998).

³⁸ In this regard the team leads function in a manner akin to Tushman et. al.'s boundary spanners individuals (1977;1980;1981).

³⁹ Only half-humorously a couple of engineers suggested that the best possible systems improvement in their group would be cloning their team lead.

8.7.1 Enactment of Perceived Authority

It is important to note that Stress team leads' organizational power does not stem from a formal position or entitlement, but instead it is derived from their management of the ROC, the fundamental boundary object stream between the two groups. There is a co-reinforcing structure between the ROC system and the Stress team leads centered on the enactment of authoritativeness (Weick, 1969).

The perceived authority of the ROC system is enhanced by the information gatekeeper role played by the Stress team leads. Through his final editing authority on the to-be-archived ROC, the team lead is the final redundant level of verification for the record. The multi-way negotiations present in the ROC as boundary object are reduced to a two way negotiation, in Stress' favor. The team lead's role as single point of contact for this multifaceted boundary object is very important for reliability.⁴⁰ The final approval by an known expert also enhances the recontextualizability of the ROC in later reuse. In addition to the archival ROC's validation, it has been appropriately filtered.

In turn, the team lead's administration of the ROC system provides them with unique awareness of every job being worked in their group. Their tight interaction with the current ROC workflow and historical archive enhances already extensive meta-expertise on the knowledge contained within their team, both contemporarily and historically. All of this increases their perceived authority within GTS-West.

The following are some examples of mutual reinforcement as observed in the routine operations of the Stress team leads.

8.7.2 Job assignment

When jobs are assigned to Stress from Structures, the ROCs are sent to a team account which the Stress team lead manages. Upon reviewing the request, the team lead surveys the expertise and current load of each of his team members to find the most appropriate fit and prioritization before making the assignment.

This is a complex process, with one team lead mentioning the tension that exists between best utilizing the skills of senior members while aptly training more junior members. For him the decision usually comes down to the prioritization of the job. If it is a difficult job and it is a rush, he will assign it to a senior engineer. However, he does routinely select both rush jobs and complex jobs, though not in combination, to assign to more junior members.

Scheduling of human resources, such as this, is one of the most divisive issues between the two groups. Clearly every Structures engineer would like to have the most experienced, most efficient analyst operating on his job. In diffusing this tension, everyone acknowledges that the mediation of the team leads is the organizational oil

⁴⁰ Many high-reliability organizations, such as the armed forces, have a similar hierarchical organizational structure with individuals operating as single points of responsibility.

which keeps the machine so well run. And the means by which the team lead accomplishes this is through controlling the assignment of ROCs in the workflow system.

To best understand how this behavior plays out it can be valuable to examine some situations in which things break down. The following are three such examples.

The first is a situation in which an attempt to circumvent the team lead backfires. A Structures engineer had previously worked on a job with a particular analyst that day and was quite impressed with the analysis and response given in her ROC. A similar job had arisen and the engineer sought out the same analyst, handing off the ROC packet to her. The analyst accepted, believing the job to be quick and easy. After getting the job underway she would coordinate with the Structures engineer to have her team lead assign the digital ROC to her. However, she soon discovered that the job was more complex and involved than she had expected. Since she had accepted the physical packet she felt obligated to carry it through to closure. (She could not informally reassign the job by handing off the packet to a colleague without team lead approval and assignment of the digital record.) Once the subterfuge came to light, the team lead was incensed that this underhanded exchange had occurred and that this new job was distracting the analyst from her other three higher priority jobs already assigned in her queue. The team lead removed the physical ROC packet to his desk, placed it in the appropriate place in his “to be assigned” queue. It was later assigned to another analyst even though the original analyst actually was best suited for the task.

Another routine complexity is dealing with absentee analysts. It is the job of the team lead to appropriately reassign jobs in such situations. The team lead does so by knowing a secondary password for every analysts account, logging in as them and reassigning the job back to the team lead account where he reassigns it to another analyst. In this particular instance, the job in question was a crisis. The operator was desperate for a reply. The analyst had changed her system password and neglected to update her team lead. (“The ROC is in her account. I’m supposed to know her password, but...”) After multiple attempts at gaining access, the account was automatically locked out. The job was delayed until he could contact the absent analyst.

Lastly, it is very rare that ROCs are not assigned in a timely fashion by the Stress team leads. However, there was at least one case where this had happened and the due date was missed. The team lead had the authority to re-order the prioritization within the group to ensure that this job was the highest priority the following morning.

8.7.3 Availability

The Stress team leads are always on-call. Each has been provided pagers and dial-up terminals to access the ROC system when coordinating after hours emergencies. On times when a given team lead is absent, a formal substitute is always provided. However, informally, the expectations for collaboration are lowered. Emphasis frequently is placed on resolution of jobs currently in the queue, versus new jobs. Interactions are more often between Structures engineers and the Stress manager or individual analysts. When crises arise however, tensions escalate without the strong, knowledgeable mediator in place.

8.7.4 Defender

The team lead also acts as a defender for his team against unrealistic expectations and spurious service requests. His control of the ROC assignment process allows him to call into question suspect ROCs received from Structures, prior to assignment to an analyst.

There are myriad examples of Stress team leads questioning the appropriateness of timing for particular jobs. (In fact the issue of “false AOGs” was one of the most contentious at the business process re-engineering sessions.) The following is just one example of the team lead confronting the originating service engineer: “Why is this AOG? It’s in D-Check! ... 29th? You let this go from Monday? Just getting it to us now?”

The Stress team leads also defend against unnecessary requests for FAA certification.

The DER’s at GTS-West are seriously impacted and their delay affects the time to completion of many jobs through the department. Most operators desire the FAA certification for all questionable repairs, just to be on the safe side. This overburdens the DER’s further and the Stress team leads actively question incoming jobs to avoid this.

The following exchange between a service engineer and a Stress team lead highlights one such defense:

Structures: “The clip to keel is not what’s critical but the bolts in the web. They want to repair to blueprint.”

Stress: “An 8110’s not necessary. They are doing it to blueprint. Why do you want this?”

Structures: “They think it is major, 85 rivets.” [a major repair requires FAA approval.]

Stress: “Who needs it? [8110] If it is just going back to the blueprint?
Makes it easier for us [not to do the 8110].”

In summary, the single point of control provided by the team lead via the ROC system, is able to simplify a complex multi-way negotiation into a more manageable two-way negotiation. In the end it is not really this simple. For example, they frequently need to integrate information from external expertise.

8.8 Integration of external expertise

Routinely the groups require specialized expertise outside of their purview to complete a job. The ROC becomes the means for integrating this information into the formal, historical record. The external information sources generally fit into three broad categories: within GTS-West, within Global Airframe and external to the corporation. Each will be discussed in turn, including some vignettes highlighting the role of the ROC in this integration.

8.8.1 Within GTS-West

Frequently the required expertise is available in-house, with the groups being able to leverage the knowledge of the different and diverse specializations within GTS-West. Both groups have routine collaborations with the other service engineering groups such as Avionics, Electrical, Environmental, Hydraulics, Mechanical, or Propulsion.

In one particular example, an auxiliary power unit door is being modified from one aircraft for temporary installation on another in an AOG situation where the door had

been lost in flight. The job fell to Structures since the door is a structural unit. The engineer is confident about the structural elements, but needs to be certain that the modified door does accommodate the electrical sub-systems, wiring, electromagnetic and heat tolerances. Appropriate expertise for this exists just down the hall in the electrical systems group. After performing the stress analysis the analyst prepares to hand off the job, mentioning that, “this is basically one half Systems and [one half ours].” With a normal priority job she would reassign the ROC to the appropriate person in Electrical. However, with an AOG she will handle this personally, collecting the information first hand and writing it into the ROC herself, “I’ll just walk down the hall and list him as a reference [on the ROC].” She mentions that the responsibility for this technically belongs to the Structures engineer, but she will run it over in the interest of time and because she knows that her team lead will ask her about the consultation before signing off anyway. She has worked with the expert in Electrical on jobs before, but does not know him personally. “I don’t know his name, but I know where he sits. I walked by his desk earlier and he wasn’t in. I was just going to write him a note to put on his chair [when he showed up].”

Frequent data gathering or coordination efforts also occur with the general customer service areas within GTS-West such as Fleet Statistics, Airline Support, Fleet Support, Flight Safety Investigations, and Warranty.

8.8.2 Within Global Airframe

Other times, the expertise search is flung further afield, but still within Global Airframe. Often they need to contact designers, lab technicians, flight simulation creators, and experts in specific material science sub-disciplines.

In the following example a Stress analyst consults with his team lead regarding a particular problem holding up the resolution of his ROC. The team lead advises he contact an outside expert,

“See if Jim is in, over in [building] 410, for the weld data. If he’s not in, we are stuck. It looks pretty simple, he should just be able to look at it if he is in.”

Even though the job was due that day, they decide that they will have to force a wave-off since there is no other expert available to consult.

8.8.3 Outside of Global

Occasionally, the expertise search extends outside of the organization. In the first example the team lead suggests to an analyst that he contact a parts supplier to Global to verify a figure critical to a current action item.

“This thing is going to drag on and on, so if you have an AOG or need help, just let us know. Call the company up that made the seal and see what they recommend for the valleys [grooves]. I think our numbers are a little high, not sure how we came up with .040 [“point oh-forty”]”

In the second example an analyst offers another analyst some fairly esoteric information regarding a vexing ROC he is working and offers to put them in contact with the external firm. The team lead, overhearing this discussion, agrees and encourages the analyst to do so.

“I was working with Aleutian Cargo back when they started flying to Russia. It turns out that at this time of the year the ice on the runways starts to crack... causing serious problems landing larger jets, chewing up the runway and damaging the gears. They decided to shift to 27's and 37's [competitor Boeing's models 727/737] which have a much narrower wheel base. They might have the info you need. Let me put them in touch with you.”

In this chapter the following has been examined, the various manners in which the ROC, acting as a boundary object, is deeply interwoven in the coordination, collaboration, and negotiation between Structures and Stress. With regard to notions of completeness, the ROC serves as the central negotiation point to reach a common ground, acceptable to both parties, on a case-by-case basis. We have seen how the tension surrounding the control of the ROC is relieved by supporting two different representations of the resolution of an action item: an internal view, used primarily for later reuse, controlled by Stress and an external view, presenting the results to the customer, controlled by Structures. In addition the enactment of authoritativeness over the ROC by the team leads provides a single point of control. This simplifies the complex multi-way negotiations mediated by the ROC into simple two way negotiation with a clearly dominant party. This assists in ensuring reliability and enhances later reusability.

In the next chapter we will examine the role the ROC plays in facilitating the collaborative activities across organizational boundaries. The increased complexity of these boundary crossings will highlight a different set of core issues.

CHAPTER 9: Negotiation and Translation, Interorganizational Boundary Spanning

9.0 Boundary: GTS-West / Operator

The ROC as boundary object, explained in Chapter Eight, with its crystallization of the coordination, collaboration, and negotiation acts between Stress and Structures, effectively hides the true complexity of these interactions. It is a snapshot of the myriad information and process streams surrounding any given job. Representative artifacts in those streams are folded into the ROC for later reuse by the engineers, and its entirety is ultimately evaluated, modified, and approved through a single point of control, the Stress team lead.

Much of the complexity hidden by the ROC, is exposed when coordination, collaboration, and negotiation are examined not within GTS-West, as previously presented, but outside of the organization. The term “boundary object” begins to break down here, with “object” hiding the complexity of the multifaceted information and process streams. The intent of this chapter is to expose those streams, as activities rather than as static information “objects.”

As discussed above, while Structures and Stress are two distinct communities of practice, they have much in common when compared with the cultures of those “at the other end” of an incoming repair request. The other organization may be an airline, a freight

company, a specialized repair center or even a private party. The individual may be a Global field service representative, a senior engineer at the operator, or a repair crew member with only a high school diploma belting out instructions from the tarmac during a typhoon.

The delicate dance of prioritization continues in these interactions, as between the two engineering groups. While safety is the paramount concern of all parties involved, almost all operators are profit-seeking entities and GTS-West is ultimately a customer service organization. The former desperately seeking to reduce repair down-time to increase revenue operation of the aircraft, and the latter seeking to appease the customer in the hopes of securing future sales.

Other communication complexities are problematic in these interactions as well: language differences, poor communication skills, different perspectives and the levels of abstraction required to discuss machinery as complex as an aircraft.

It is also important to remember that the majority of all repair issues are dealt with locally, at the operator or maintenance facility, without engaging GTS-West at all. (One cargo operator estimated their ratios at 80% internal only, and 20% with some degree of collaboration with GTS.) Therefore, only the most difficult and challenging jobs require interaction with GTS. By their very nature as jobs which cannot be solved locally, these jobs can be hard to characterize, describe and define.

9.1 Negotiating to completeness

As noted between Stress and Structures, the primary form of negotiation surrounding the ROC involves completeness – bringing the quality of the information represented in the record up to the standards of the group involved (most often that of the Stress Analyst). This tension is exacerbated in dealing with the operator. In almost all cases no GTS-West engineer will ever see, touch or test the aircraft in question. They rely solely upon the information provided by the operator. They must have a high degree of trust in the quality of this information in order to be confident in their suggested repair. However, unlike with Structures, the information provided by an operator is almost universally incomplete in some aspect. These omissions can be benign, such as a missing measurement, a unexpected dependence on a neighboring assembly, or they can be more suspect, such as deliberately withholding or amending information to speed repair time. In all cases, this incompleteness must be dealt with.

The iterative process of Structures and Stress building confidence in the data supplied by the operators is one of the primary external negotiation processes involving the ROC.

While there are hundreds of examples of this behavior observed in the study, an example from one of the cases is provided next to illustrate some of these behaviors.

9.2 An Example

Friday morning, a request comes in from an international repair facility servicing a TransCon Airlines (TCA) craft. The maintenance crew has discovered corrosion damage on a prior repair, an internal doubler. They would like approval to operate the airplane as-

is, but are suggesting a repair if that is unacceptable. If the corrosion has to be repaired they are requesting approval for the repair plan and FAA certification.

Reviewing the materials submitted by TCA, Rasim, the Structures engineer assigned to the task, is deeply frustrated. “There’s no part number, no location, [it’s] all too general... [and] look at this! Here’s the sketch they sent. It’s probably just an IRC copy. You can’t tell anything.” [The image is a high-level diagram of nearly the entire fuselage copied out of a manual. A hand drawn arrow pointing to the underbelly of the craft is supposed to highlight the damaged area.]

The supplied information is so sparse, Rasim is hesitant to offer even an educated guess. He cautions, “you need that to know exactly what the damage is so you can do the calculations.” He needs to contact the repair facility directly, however “it is like fourteen hours different, we’ve already lost a day.” Unable to contact them due to the time zone differential, he leave voice-mail and writes up his request and faxes it to the Field Service Representative at that airport.

On Monday, Rasim has not heard back from the repair facility, TCA or the FSR. He manages to contact the FSR, but cannot get the information that he needs out of him, “he’s a FSR, so it’s back and forth, back and forth, losing time. I couldn’t wait to spend even more time, it’s due today.” So, to meet his deadline Rasim decides to take matters into his own hands and make the best educated guess he can.

“I did the research. I had to search. I got the airplane master file list. I had to find the actual part number. I have to get the configuration for that

plane. In [this] case, there are many configurations, all for the same plane. I did [search STAIRS] and didn't get anything. There were no similar repairs. I had to pull out books though, the SRM. We have to compare their repair with the SRM."

After assembling a mountain of scratch work, Rasim was convinced that he had reached the most plausible explanation for the damage and requested repair. He passed his idea by the FSR. "I confirmed what I found with them. They checked. It was right." The ROC, which he had been working on in scratch form since Friday, had reached an appropriate level of "completeness" and was now ready to be formalized and sent to Stress for analysis.

In assembling the ROC packet for handoff, he included none of his scratch work but called attention to the results of his research ("I also highlight[ed] the details, material, thickness and all.") and included pointers to some of the most helpful reference materials. He also edited the original request to match the standard departmental abbreviations and omitted the fact that this corrosion had been found during a routine D-check.

An hour before the start of the work day on Tuesday, the team lead for Stress had reviewed the ROC and assigned it to Philip as high priority since it was due today. Philip began the job first thing. He was unaware of all the work Rasim had done in the prior days preparing the packet, commenting that "this was straight forward. It had all the information. The vendor sketch was really good [pointing to Rasim's annotations.] Shows you it all right here, site, damage, thickness and all. This was one of the better ones, I didn't have to use anything else [reference wise]. Some require a lot of research [but not this one]. Guess I was just lucky this morning. This wasn't bad. That really was

one of the easiest I have ever had. It had everything there, I didn't have to look for anything. Wish they could all be like that.”

He did however find an error in the double thickness measurement supplied by TCA and calculated the correct figure. “They claimed that the loss was .015, but I found that it is really .025, right? [he consults his scientific calculator] .023, it is really .023.” He makes a formal note of this discrepancy in his final reply (“... material loss is 0.023” not 0.015” as reported.”).

With this new figure, he is forced to reject their repair proposal but offers a similar substitution. “They want to leave it as is. That won't work. This is a high stress area of the door, down in the corner there... I replaced the amount of material lost [with] a standard 32, a .032 doubler. That will replace it. I recommended SRM 53-04, page 38, sheet 20.” He acknowledges that as a foreign repair station they may not have the materials on hand to match this new repair plan exactly, but so long as they can meet it or do better. The job is routed back through the team lead for check and verification, through the FAA DER for certification, and returned to Rasim for final reply.

While Philip was working the job, the repair facility finally returned Rasim's call regarding the missing details. “They [repair station] called back [after we had talked earlier this morning] and still didn't have the information. Can you believe it? They're just lazy. They didn't want to go and do it [verify the part and supply the additional information]. Isn't that vexing? They can be so lazy. My job is to get the exact

information for the exact repair. In order to support the customer, we spend more time doing the research. Often we get no support from the operator or rep, and they have it [the aircraft] right there! This is my daily life... [sighs]”

As frustrated as he was with the operator, he was ecstatic about the quick turn around on the part of Stress and the DER, “Sometimes it just works out. We really get to support the customer. I got it from them, 8110 and all, agreed and sent it. If I feel comfortable with it, I send it right away.”

The interactions evidenced in this case are typical, representing the majority of the day-to-day requests at GTS-West. Below are some important general observations about this case.

9.3 Opaqueness of routine work

Of course, the incompleteness of the information is exacerbated by the general opaqueness of the specific activities of the individuals to the others involved. The maintenance facility staff appears clueless throughout. The FSR does not understand the context for the repeated requests for additional information. In fact, his responses are infuriatingly unhelpful, spending most of his time thanking Rasim for the correspondence and nagging about the due date, without supplying any of the needed details. Rasim is uncertain of the DER’s 8110 approval process (“Sometimes you have to wait, like to four o’clock for [the DER] to see it and sign. That really slows it down. It’s beyond my control. I really don’t know how it all works [DER approval]. It can even be days

sometimes.”). Most importantly, the stress analyst, Philip is oblivious to the nightmare data collection ordeal Rasim had just completed.

This opaqueness of the different disciplines is a common source of frustration, friction and negotiation among the groups. In this particular case, the onus fell on the service engineer. More frequently it is on the Stress analyst (designing a new analysis from scratch or performing a complex set of load calculations can take hours or days). It can also be on the DER (especially if they require approval from an external FAA office) or even the operator (conditions at the repair site change, crews shift, more important problems are discovered). All of these processes are black boxed for the operators which are constantly cranking up the pressure, requiring more detailed work in shorter amounts of time. After one particularly frustrating case, a Stress Analyst commiserates with his Structures counterpart, “[You] bust your butt for an airline and they come back and complain... typical.”

9.4 Real and perceived operator abuses

The complaint that the operator is lazy or incompetent is almost universal among Structures engineers. There are clear informal rankings of the airlines and repair stations within GTS-West based on the quality of their submissions.⁴¹ Some organizations are held in very high regard, by only submitting the truly difficult questions, requesting reasonable response dates and always providing ample information. However, there is a

⁴¹ During my tenure at the facility, there were many occasions where engineers would hang “half-assed sketches” or ludicrous requests in the cubicle hallways as a humorous hall of shame.

common sentiment that the operator is always going to require more work of GTS-West than they ought.

Taken to an extreme, some small independent repair firms have come to rely on GTS-West for free engineering services. Having their repair designs unwittingly crafted by GTS through the iterative process of information gathering, question clarification, analysis and approval. The BPR team addressed this as an area of high concern considering the financial and resources drain it has on the department.

9.5 Information elicitation

The process of wrangling out the information required to complete the job is problematic and often unsuccessful. The pressures represented in this case are exacerbated when the job is a high priority, such as an AOG. A near hourly refrain from the service engineers in Structures is “[I’ve] called for more information. Will wait for them to call back.”

One means of facilitating the elicitation process is to employ a standardized form, which would be required of all operators making service requests. This is actually a practice, facilitated by GlobalCOM, used by other GTS facilities. The form had been experimented with here, deemed a failure, and abandoned. This resulted in a shift in tactics from requiring the operators to supply information in a particular format, toward enticing them to do so. One service engineer half comically described the sales pitch to the operators as, “having it in a more machine readable format, pre-populates many fields and assists in initial routing, reducing the time it takes to get a response. Sound beneficial to you?”

9.6 Coping with intractable operators and absent data

There are common approaches to dealing with absent information or intractable operators. Within Structures engineers will do as much research as they possibly can locally to fill in the gaps. Exceptional service engineers will do this automatically in advance of the hand off to Stress. Others economize, waiting for Stress to complain about specific missing pieces, before they start researching. Within Stress, their response is to be extra conservative with the repair. If precise measurements are not provided, they will fudge toward safety requiring a larger, thicker, stronger repair. (This frustrates operators though, as the best customer service gives the most minimal repair for required safety and strength).

The following dialog is typical of how this plays out between Structures and Stress. It involves an unhelpful operator requesting approval for a revenue ferry flight. (The aircraft is in between flights and they want to carry passengers until it is convenient to touch down, take it out of service, and do the repairs.) After many iterations with the operator trying to get the missing information, the Structures engineer is exhausted and he discusses the case with the Stress engineer. He is looking for a way to resolve this. Since GTS-West approval is required for the airline to make this flight, he asks for the worst case scenario that he can present to the airline.

Stress: “Well, the worst case scenario is complete failure, right?”

Structures: “They don’t need our approval for this ferry flight, however, I’m just tired of going back and forth with these guys.”

Stress: “My recommendation will clearly be to repair it before it goes into the air. Can they check on it?” [to procure the omitted details]

Structures: “No, there’s no one there who can.”

Stress: “What about their engineer? You have to install the external doubler on it before you do the ferry flight.” [if they are going to do the temporary repair, clearly they have someone standing by.]

Structures: “[He] only comes in at night. [I’m] trying to clarify the specifics and no one there can help.”

Stress: [sarcastically] “They can’t walk out and look at the airplane? We can’t give them approval based on this sparse information.”

Structures: “Field service rep wants just a ‘yes’ or a ‘no.’”

Stress: “Well, ‘no.’ That’s easy. We don’t even know specifically what’s wrong.”

Structures: “They can just blow off whatever we tell them and go ahead and fly it. The rep told me the airplane is still in service.”

Stress: “Whatever they want. This is their choice.”

They continue to discuss in detail what is missing and what they could research locally.

They agree that what they can do is not enough. Later on in the day the operator is complaining bitterly that the approval has not come in. They have changed their request to just flying the flight crew, no passengers.

Structures: “They feel that we have enough information right here to make a decision.”

Stress: “It’s a ferry flight only now, right?” [not revenue]

Structures: “Yeah. I must have had twelve phone calls on this so far this morning and it’s not even my thing. They are really pushing this...”

They decide that this is now a matter for management, and in a rare move, involve the senior service engineering management in this process.

In a second example of dealing with a problematic operator, a Structures service engineer is frustrated in his efforts to get complete details out of the operator's representative. The missing information is critical and he knows that the job cannot be handed off to Stress until the full details are in place, "The door has to fit perfectly – hinge, latch, everything."

Knowing that the contacts for this particular operator can be frequently clueless, he decides to step them through the visual inspection process over the phone himself. He will walk them through all of the relevant contact points. "They're usually pretty good, but there could be a gap this big [holds up index finger and thumb to indicate about an inch] and it is still miscommunicated."

If he is convinced that he has enough details for Stress to make a conservative repair, but not enough for their required degree of completeness, he will call out the differences in the final response to the operator. "If I'm uncertain, I'll put it in [the response]."

Basically he will be very explicit about the measurements for all of the contact surfaces.

If these measurements are all met, there is no way the door can be a misfit.

9.7 Due dates

Recall that due dates are imposed externally and rarely reflect the actual time required to accomplish the given task. The official position of the department is that the customer always knows best. They are the ones scheduling the aircraft and they know when they need it serviced by and back into operation. The problem with this philosophy is that the

airlines rarely build in slack (they used to). They have little institutional incentive for doing so.

All of this necessitates very complex, dynamic scheduling for coordinating the activities within and between Structures and Stress. This herculean task is the responsibility of the team leads within each of the model groups. In general the system works surprisingly well, however when the system is stressed, for example by a steady stream of AOG's one morning, it fractures quickly. One sample lament from a Structures engineer whose schedule was just upended "[the job] comes in on BOECOM with a due date of today, so I'm basically hosed." He cannot alter this due date, but he cannot complete the job on time, so it will have to be a 'wave-off' giving him a ding in his performance evaluation.

Stress' frequent suggestion is to interrogate the operator about their requested date or to simply request more time, but as customer service engineers, Structures is loathe to do this.

In summary, as one can see, incompleteness is an inherent aspect of the ROC. In no way can it show all of the complexities of the often complex flows of questions and answers from the operators to the GTS-West personnel. In Chapter Ten this incompleteness and the dynamic aspects of the ROC, and by extension other boundary objects, will be revisited to examine reuse.

CHAPTER 10: Recontextualization and Reuse, Boundary

Spanning Across Time

10.0 The notion of boundary objects across time

In Chapter Seven the theoretic concept of boundary objects was applied to explain the role the ROC plays in mediating the coordination, collaboration, and negotiation between Structures and Stress in daily work. As controlled by the Stress team leads, the ROC crystallizes a particular configuration of these interactions in time for the purpose of later reuse. Those processes of reuse will be addressed in detail in this chapter.

In Chapter Eight, it was shown that the traditional notion of boundary objects begins to break down when viewed in the dynamic stream of collaborative activities. This breakdown was most clearly obviated by interactions which occurred across organizational boundaries. The limitations of the boundary object notion to best represent these behaviors and processes are further expanded in this chapter. In discussing the problems of negotiation and representation between GTS-West and the operators, it was shown that more dynamic structures and views of the ROC were required.

In examining the ROC as a boundary object across time in this chapter, it will be shown that these breakdowns are even more severe. When dealing with a boundary as loosely defined as a temporal boundary, the other parties in the negotiation become less crisply defined (e.g. it may be a negotiation with individuals who are no longer available in the current context, or even a prior instantiation of oneself). As the contextual cues diffuse

and decay over time it becomes more difficult to assess the authority, accuracy and relevance of the information contained in the ROC.

However, the engineers in Structures and Stress are quite capable at leveraging this seemingly impoverished archive to meaningfully assist in the completion of their current jobs. This chapter will examine that ability. It will first introduce the general climate of reuse practice at GTS-West. Following, Ackerman's (1994;1996) notions of decontextualization and recontextualization will be revisited and briefly illustrated with data from observed critical incidents. These notions will then be fleshed out via a full length case-based example, which will reintroduce the complexities from Chapter Nine of negotiating across the organizational boundary with the operator.

10.1 Introducing the practice of reuse at GTS-West

Global aerospace has a corporate-wide policy on general knowledge management practices and a focus on repurposing⁴² efforts. The company is committed to fostering a stronger culture of reuse among its divisions. Due to the relatively recent corporate merger, the current environment at GTS-West remains noticeably different from the remainder of Global. In introducing the practices of reuse at GTS-West, the reader is reminded that these are unique to this site. (An obvious, engaging follow-on study would be to undertake a comparative study between the different GTS facilities.)

⁴² "Repurposing" is the common colloquial term for reuse among Global's engineers.

10.2 Culture of reuse

In an effort to maintain consistency with the operators of the aircraft models which GTS-West services, Global has maintained a decidedly laissez faire attitude with regard to the department. On key matters of service provision, GTS-West has not been forced into Global's standards and culture. This has allowed the majority of GTS-West's policies and practices to remain intact.⁴³ The most notable of these holdovers is their approach to reuse.

Records management policies between the two original corporations are distinctly different. Long influenced by government and military contracts, Global has maintained a records archival policy which mirrors these contracts even in commercial aircraft. With the exception of the most critical design documents, most are retained for seven years and destroyed. This policy yields an available knowledge base which is current, relevant, and accurate, yet sparse.

In contrast, GTS-West's original management exhibited much more of a pack rat mentality, maintaining warehouses of information dating back over eighty years. While the original company has merged multiple times since its founding, all merger partners have been distant landlords and have had little impact on this particular practice. This policy has yielded an available knowledge base which is rich and complex, yet nightmarish to navigate.

⁴³ This relaxed stance is not unusual given the semi-autonomous nature of the GTS facilities, as recalled from Chapter Five.

One of the Stress engineers had participated in a short-term exchange program, transferring to another GTS facility shortly after the merger. While he had an extensive commentary on the differences between the two offices, what surprised him most was this variance in record keeping practice. “[at the other GTS office it was] very different. For things like the stress reports, these [taps on the cover of one stacked on his desk], they’re not all together, not complete... not everything is there. Someone does something, writes it down, and just throws it away [motions toward his waste bin]. It’s not always kept. I find the books much better here. Now, that’s [the behavior at the other GTS office] not bad, they just do things differently.”

GTS-West has developed policies over the years which favor these rich local archives. The formal procedure for reuse within GTS-West is best captured by this reminder memo distributed by a senior manager to his reporting groups during the first phase of this study:

It has been brought to my attention that standard procedures are not always being followed when an analysis is being accomplished. Therefore, everyone needs to adhere to the following procedures when working on a ROC.

Check the history of the part requiring rework / repair. Use the STAIRS database and history files. If we have seen similar damage to a part, we may determine that service action is required. Also, if we have a prior repair, and this repair is still acceptable with today’s requirements, the same repair method should be followed to ensure consistency.

When completing an analysis, list a reference for the analysis on the ROC. This helps ensure traceability of the repair in the future. If a repair is signed off by comparison to prior analysis, the analysis for the prior work can be listed.

Note that the rationale for consulting the historical ROC archive is threefold: ensuring that the current analysis is in line with precedent, noticing any relevant patterns in repair history which may require resolution at a higher level (e.g. fleet-wide), and to repurpose prior work to assist in the resolution of the current job. This manager's memo then is an appropriate segue into a discussion of the different methods of reuse at GTS-West.

10.3 Methods of reuse

Within Stress there are four primary approaches to reusing archival ROC entries: precise matching, repurposing elements, precedent-setting and procedural. Given the high potential pay-off for a precise match, most Stress engineers perform an initial search in this manner for the majority of their jobs. As such exact matches are relatively rare, they routinely expand their search criteria to examine similar cases. Once these have been reviewed they weigh the costs and the benefits of recontextualizing the relevant archival record to assist in their current task. If they can gain a head start by adapting some material from the past, they will do it. However, if the work of recontextualization is too great, they will simply work their current job from scratch. The method of precedent-setting searches serves a different purpose and is of a different kind than these two. All four will be described in greater detail below.

10.3.1 Precise matching

In this situation there is a precise match in the ROC archive. After carefully confirming the match, this ROC can be used as a reference for the current repair – the analysis, repair and resulting approval can all be transferred to the new case.

A junior Stress analyst expressed excitement at the time saved by a precise match. He had been working through a lengthy but routine lug analysis, and realized that there may be a prior case which could shorten his work. He used STAIRS and found a match, “I just pulled it up, checked the analysis, and plugged my numbers in. Even made an Excel spreadsheet [to do this].” When asked about whether he saved the spreadsheet, he mentioned that he had, just in case it might come in useful later on. (This behavior is described in the procedural reuse section below). Yet only approximately one in fifty searches is successful, based on my observations.

Another example of precise matching is provided by the central case-based example of this chapter, in section 10.5 below.

10.3.2 Repurposing elements

This is a situation where there are not exact matches within the ROC archive, but there are a number of very similar historical cases (such as the similar dent damage from 1991 in the central example from Chapter Five). In this situation the engineer will likely either use the ROC as a starting place for his analysis, adapting relevant elements to give himself a head start, or use pointers from the ROC to find other relevant resources (such as a calculation template in a technical journal). Their questions are ‘what can I take away from this? What can I adapt?’

The following example captures an unsuccessful attempt at repurposing.

A military repair base is doing a run of repairs on the main landing gear assembly doors for a reconfigured aircraft model N-20. They are out of stock for parts and the delay in ordering is too long. They need the repair immediately and are asking – what are suitable substitutions? Even if they are temporary installations?

The most recent documentation for the doors, from 1982, is unclear with regard to substitutions. The service engineer comes over to the Stress team lead’s desk to “brainstorm together.” They both look at the problem and using their combined extensive personal experience they try to come up with options off the top of their head. They are not successful. The team lead suggests that they look at prior cases in STAIRS in order to try and find something that is a close fit or at least might give them a lead to some other resources. They start searching together.

The core problem here is that this is a unique repair request. It involves the hinge on the fuselage side *not* the door side of the assembly. They find a number of cases, but all deal with damage to the door side. None of these are close enough to support reasoning by analogy. Eventually they decide that there is no obvious substitute. They confirm this with each other (“you agree? Okay, we’re resolved.”) and the service engineer responds to the repair base that they will have to wait for the part.⁴⁴

⁴⁴ If there had been some relevant historical precedent from which to work, or if there had been more time available for doing some creative repair design engineering from scratch, the likelihood that the service engineer could have found a temporary replacement would have been much higher.

10.3.3 Precedent setting

This is the case where there need not be an exact item-for-item match for the entire ROC (e.g. the suggested repair could be different); however, some critical element in the ROC must match the current job exactly. Engineers use this method to ensure that their conclusions are in line with historical precedent. Their questions here are ‘what has been allowed before on this? Under what circumstances?’

In my observations, there were at least three motivations for utilizing precedent setting search, which I call *purposive*, *comparative*, and *incidental*. Each are described below.

Purposive precedent searching entails situations where the engineer needs to verify that his resolution is in line with prior cases before sending the final reply to the operator. In this example, an airline is in a fix. Its mildly damaged plane is marooned at an airport about a hundred miles from its central repair station. They are requesting approval for an unpressurized ferry flight to the repair base. The team lead from Structures, handling this high priority request directly, confers with his equivalent Stress team lead.

Structures: “Is it feasible to do it one time? Just a short distance...” [He is not pleading, but earnest.]

Stress: “Oh, okay. But before you do that let me check STAIRS and see if we have ever done one of these, a ‘one-time ferry flight with cracks.’ I’ll let you know.”

Comparative precedent searching is used less for validation and more for understanding the range of approved variances. This helps assess expectations, both from the operators and from fellow engineers within GTS-West. In the following example, a service

engineer from HyraMech searched to understand the range of approved variances in advance of submitting a high priority, but potentially problematic case to Stress.

“This is due tomorrow, not AOG but urgent. I looked through STAIRS and found a bunch like it, but none exact. I’ve seen ones that are less [in better condition than his case] that have been scrapped. Mine’s borderline, but I didn’t want to send it away ‘okay, go used,’ and have your guy come back and say, ‘shit, scrap it!’”

With *incidental* precedent searching, patterns of approval are revealed in the midst of other searching behavior. In the following example, a service engineer did not find either the precise match nor the repurposable match for which he had been hoping. However, he did notice that in reasonably similar skin repair cases 8110’s are not required and usually had not been granted.

Structures: “Here’s my AOG for the day. I searched the database and didn’t find anything like this. IndoAir, horizontal stabilizer, side load skin. We have no parts and it needs to be replaced. They had an earlier repair, but now found even more corrosion. [He unfolds the blueprint on the team lead’s desk as they talk through the details.]

Structures: “They want to make new skin. They only want to make it temporary, 135 flight hours, but I’ll doubt we’ll get the part in time. I’m going to have to write you guys a ROC on that, but just to give you a heads up that we haven’t done this before.”

Structures: “I don’t know if they are going to want an 8110. I hate when they do this – they say it is minor, but we realize later it is major and needs an 8110.”

Stress: “That’s the operator’s responsibility.”

Structures: “I know on the other skins we don’t [request an 8110], and this is a small piece and it’s temporary [even less logical to request an 8110], so it should be alright [without one]. I hate to do it to you but it seems as though every day I get hot ones.”

Stress: [questioning how hot this really is] “It’s still in C-check?”

Structures: “No, we found out later that it’s due out in three days [still AOG, a notch higher than work-stoppage]. I’ll be back with a ROC and all. Will type it up right now.”

In addition, the engineers are well aware of the precedent setting nature of ROC entries.

In one particular case, an operator is in a bind replacing a rudder. They are under time pressure, missing parts, and some equipment so would like approval from GTS-West to rebalance the rudder by calculation instead of physical testing as required by the SRM. This is generally allowed once in the lifetime of the part, but this would be the third time for this particular part. The service engineer and Stress team lead agonize over the decision. They trust the repair facility for this operator and would be willing to grant the exception except they are hesitant about setting a dangerous precedent in the ROC archive – “setting an example for other airlines, especially those that paint their rudders all the time. That builds up.” They search the archive for any other third time calculations and find none. In fact there is only a single case of granting an exception for a second time calculation, and that case included extensive argumentation from the operator (details they do not have for this case). They agree that they are going to have to reject this request for the moment. In the meantime the service engineer will check with some experts from Weight and Balance, a specialist department within design, about the grounds for granting exceptions in the future, though he is pessimistic that they will agree to anything beyond the SRM requirement.

10.3.4 Addendum: Procedural Reuse

There is a fourth method of reuse. It is less explicit and more abstract than these three. This is high-level, or procedural reuse, where patterns noticed across a large number of cases are generalized and formalized. The best example of this behavior in Stress is the window load calculation program and it warrants a few moments to explore in some detail. To start, the Stress engineer working this case explains the complexity of lightning strike damage and window repairs:

“When you’re flying and there’s lightning and it strikes, it usually hits the wings or the window area. This [case] hit the window area. When that happens they [operators] have to tell us where. They need to be precise. To tell us the exact location... The load changes from window to window. They’re all unique, so the location is important, very important. You need the depth. How deep did it go? When it [lightning] strikes, there is so much heat, it’s so hot that the metal fuses and loses properties [key material properties]. You need the range and the depth for this, oh, and the location, you must have the location. They [operator] want to know if it is okay to fly as is or if they need to fix... The window is so critical. If it cracks, it goes to the next window, they’re so close.”

As we have seen in the preceding chapters, eliciting this level of detail from an operator can be an excruciating task for the service engineer, although, in analyses like this completeness for these details is non-negotiable for Stress.

As no two windows are alike in load profile, and each location around the window is different, the resulting analysis can be quite complex. To ensure a high degree of reliability, window repairs have been traditionally very conservative. When operators began to complain (window repairs are very frequent), a senior Stress engineer in DTA noticed that the complexity had a regularity to it and leveraged this regularity to design an Excel spreadsheet to handle the special analyses for each sector, of each quadrant, of

every window on the aircraft. The spreadsheet is now used regularly throughout Stress.⁴⁵ It provides all of GTS-West's expected benefits of reuse: using prior cases to simplify current jobs, increasing the reliability of analyses, decreasing the variability in analytical approaches and improving customer satisfaction with more liberal repairs.

All of these four outlined processes of reuse are not error-free. In fact, difficulty in using these methods is the norm, not the exception. The next section reviews some of the common reuse problems faced by the engineers and the organizational work-arounds they have developed to cope.

10.4 Common breakdowns and resolutions

Recall from Ackerman (1994;1996) the processes of decontextualization (distilling information for storage) and recontextualization (reconstructing context for evaluation and use) in information reuse. Each will be discussed briefly here, framed in reference to the central scenario in Chapter Six and elaborated with breakdowns and resolutions common to everyday work.

10.4.1 Decontextualization

At GTS-West no action is considered closed and complete until it has been archived for future reuse. Thus at the conclusion of the job described in the scenario the resultant ROC, in both its digital and paper forms, would be prepared for archiving. In the process some elements of the interaction will be preserved, such as the original request and

⁴⁵ Print outs of the spreadsheet calculations are included in the physical ROC packet in lieu of the hand analysis or pointer to a technical journal.

details from TransGulf, as well as pointers to references such as the blueprint use, the relevant stress report and the entry in Jonathan’s technical journal. However, other information would be discarded, such as Jonathan’s initial calculations (prior to Nancy’s second visit) and the archival ROC which triggered the introduction of the metal fatigue standards which so dramatically resolved the request. The Stress team lead plays a critical role in this distilling process, as will be described in Chapter Seven. He works to ensure a high degree of precision in the archive and lends it a degree of authority by his personal review of every record.

Following is a table highlighting the most commonly observed difficulties encountered when these engineers decontextualize their actions for the ROC archive and their routine workarounds.

Table 5: Problems with Decontextualization

Breakdown	Work-around
The data for the repair request is incomplete.	Contact the operator to elicit the missing details.
The data in the repair request is suspect.	Be conservative in you evaluation and analyses.
Important information cannot be represented on the ROC summary.	Include in the paper packet or annotate the paper summary.
Background information or preliminary calculations which are important but could potentially confuse a future user with regard to the final analysis.	Include with notes in the referenced technical journal.
Important informal notes should not be included in formal record.	Store in personal filing space.

10.4.2 Recontextualization

Years later when a similar dented rear cargo door frame is encountered, a service engineer will likely retrieve the archival ROC describing this case. She will then have to

engage in the process of recontextualization for the record, rehydrating it in a sense, in order to evaluate its appropriateness for the task at hand. Some of the context required to do this will be available to her in her present environment, for example Tom and Nancy are still just a few cubicles down the hall and the cargo door blueprints remain identical. However, other requisite context has shifted over time, for example the SRM coverage for the door frame may have changed and the metal fatigue requirements may have been relaxed given new repair techniques. Still other requisite information is no longer available, for example Carl has retired and both Stress analysts are off-site for the week.

Following is a table highlighting the most commonly observed difficulties encountered when these engineers attempt to recontextualize an archival ROC. Note that the most common response is not listed on the table – find another example which is easier to recontextualize or abandon the reuse effort if working the job from scratch will be quicker. Given the time sensitive nature of their jobs, both Structures and Stress engineers are aware of the cost/benefit trade off for these activities.

Table 6: Problems with Recontextualization

Breakdown	Work-around
Listed contact or expert is unavailable.	Find nearest equivalent.
Listed resources are missing.	Search or find best approximate fit. ⁴⁶
Historical data is questionable.	Be conservative in evaluation and analysis.
Historical data is obsolete.	Trace the upgrade path to make it current.
Available, but non-local resource.	Request or travel to retrieve.

⁴⁶ One engineer highlighted the importance of referring to the original listed resources and frustration with the frequent inability to locate them, “we have to redo the analysis. If you can’t find the records then you have to redo the analysis.”

10.4.3 The special case of aging records

The most common problem with recontextualization is missing information (e.g. people have left the department, records are missing or misplaced). While the cause of this is usually recent (e.g. someone is out of the office for the day, or another engineer is referring to the manual you need), a persistent and pernicious problem is the aging of ROC resources. As mentioned above, the first response of an engineer is to ignore records that are too old. However, if the information is vital they are willing to invest time into updating the material (e.g. tracing the upgrade path for a historical calculation, taking into account new regulations and methods) or searching out a similar expert for one who is no longer available.

The first example highlights the engineer's preference for reusing more recent ROC's from the archive. A service engineer is discussing the case with a Stress analyst.

Structures: "Take a different intercostal and repair it to the SRM... You need to patch the lightning holes first."

Stress: "Yeah, a repair similar to this is acceptable in concept."

Structures: "They wanted an 8110 though. Acceptable to do minor deviation on a SRM? Vary the SRM and submit a final sketch for approval?" [Now turning to the details of the suggested variation.]

Structures: "This is similar to this one [an archival ROC, she places the pulled packet on his desk], but can't replace it there [points to the sketch] because of the extensive bundle of wires running through. That would be a major splice job."

Stress: "You've got this [one], but have you done a full ROC search?" [he is fishing for a better match than the one she has brought him.]

Structures: "Yeah. [she had done the full STAIRS search] They did this back in ninety [a better case], way back then, but I brought this from ninety-five as it's more recent. Looks the same in concept."

In the following example, a service engineer describes some of the problems with older records. In this particular case he is dealing with an operator's repair request which had uncovered a deviance from blueprint in the original manufacture of the aircraft. The role of precedence is important in a case like this if one could be found. He does manage to find one perfect match, but discards it due to age. In fact he views recontextualization as so problematic for this case that he does not even flag the archival ROC for the Stress analyst to even consider examining. He comments:

“I found a lot [of prior ROC cases], some things that could work. One ref sounded identical, unfortunately it was old, 1982. Records that old are very incomplete.⁴⁷ We'll probably have to do basic analysis unless Carlos can find something better. And some times he does. That's the problem with old repairs, old records – rough, vague and sketchy.”

Thus when an engineer describes the ROC archive as, “it is a real treasure, very powerful. It's ironic, what's so great about it is that it is so old,” the assumptions are that “old” is actually quite recent and that the truly old material still can be interpreted.

For many in the department, this ability to decipher these old records is one of the special skills of the “old guard” of senior engineers. Some have been affiliated with a particular aircraft model since the model's inception, and know where to search around the plant for arcane information sources or experts, while others who have had diverse careers at the plant have maintained useful, robust personal social networks.

⁴⁷ This is an extreme case of temporal decay. Prior to the introduction of the ROC system in 1985, two off-line systems handled similar functionality. As a requirement of abandoning these systems in favor of the ROC, the records from these two prior systems had to be incorporated into the new system. Given that the manual data entry task to accomplish this would be extensive, only a critical subset of fields were included.

Now an central example which highlights many of these reuse behaviors will be provided.

10.5 Example

Doris, a senior service engineer in Structures, was assigned an urgent new repair request from Pan-Eurasian Airlines (PEA). The extensive details of the problem and proposed repair were provided as attachments to the GlobalCOM action item she received from her supervisor.

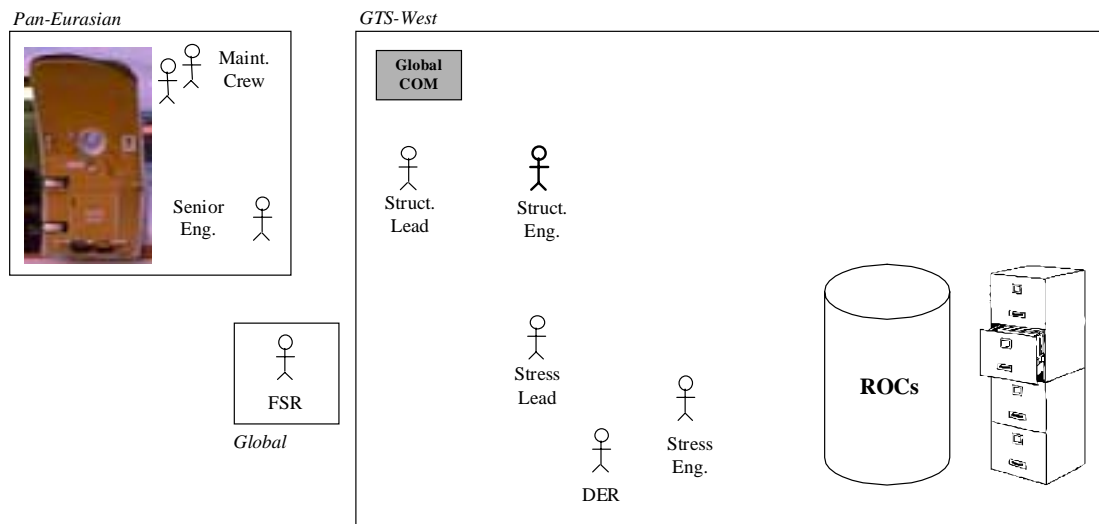


Figure 11: Information flows

During routine maintenance, PEA had found a crack in the service door pan of fuselage #13219. PEA engineers have experienced similar damage on other planes in their fleet. In fact, looking through their records, they discovered a GTS-West approved repair for a nearly identical crack signed off six months earlier. They include the correspondence identification number (CIN) of this prior job in their request for the new job. At best they

hope this will speed up the handling of their request. At worst it provides corroborating evidence for their current proposed repair.

Reviewing the materials, Doris realizes that this could be a complicated repair. “This is a problematic area. We’ve never developed an SRM [entry] for it before, so there’s not really a standard repair. This complicated area needs specific repairs.” From memory she believes GTS-West has seen one or two of these before, but “there’s not a big history of this.” This could be a lot of work and time is of the essence; at this stage in the process Doris has only 24 hours to get a reply back to PEA to be on time.

The CIN number provided by PEA is of tremendous assistance. (“The operator made reference that we had previously approved this... they tipped off that it had a history.”) It provides an advanced starting place for her search. From the CIN she finds the GlobalCOM entry for the prior job. Within this digital archive, she found a scanned copy of the 8110 which had the ROC number on it. She finds this ROC in STAIRS and prints out a copy to send to Stress.

With the provided materials, her basic research and the archival ROC the job is ready to be transferred to a Stress analyst. She writes up the current ROC request (on-line) and runs the paper packet over to the appropriate Stress supervisor.

The job is assigned to Sam, a senior Stress Engineer. The material in the packet is in order – sketches, measurements, and descriptions (“I really like them [PEA], they do a

very good job, at least in the description and the pictures, hopefully also in the repairs.”). The reference ROC is helpful, and he is not surprised PEA gave Doris a head start (“They’re good at that [record keeping].”).

It is Sam’s routine practice to always do a STAIRS search before starting a job. Even though he has the reference ROC given him by Doris, he searches anyway. There could be a better fit in the database. He finds three ROCs, all equally good, but this search triggers his memory – there was a recent job he did that could be even better. (“I did one last week that she [Doris] wasn’t aware that I did.”) This fourth ROC was not yet in STAIRS (ROCs are swept into the system, from current to archival status, on roughly a monthly basis.) Using his personal listing of recently completed jobs, he finds the ROC number for this job and pulls the paper file from cabinets. Reviewing the two jobs, he found that it was a perfect fit. “It was the exact same thing. I confirmed and that was it.” With this prior precedent, he was able to approve this repair in a matter of minutes.

With Sam’s answer in hand, Doris writes up the reply and gets the official GTS-West approval off to PEA first thing the next morning.

This approval process was complicated by the fact that PEA requested an 8110. As a foreign operator, Pan Eurasian is not regulated by the FAA and therefore is not required to request this approval on any of their repairs. “PEA didn’t used to request these, but they do now, on almost every job.” [can insert explanation of why here if needed. Benefit

to both PEA & GTS-West] It is departmental policy to comply with any such voluntary requests, so Doris places the request in her ROC.

Sam is good friends with the DER who would be reviewing this job (he sits at the desk next to Sam's). To help speed its approval, Sam mentions the prior approvals to the DER. ['this is an easy one']

Much of what transpired here never made it into the archival ROC for the PEA job. For example, of the four reference ROCs used in the evaluation, only the initial one (the one that Doris found with PEA's help) is listed. The fourth ROC, the 'exact fit' that Sam actually used for the approval, is not recorded.

Both Doris and Sam were surprised that all of the records on this repair were for "the same operator, same crack, same door, same location, same [model] plane." Clearly something was wrong at PEA, but they had neither the time nor the organizational directive to investigate this further. It just become social conversation, guessing what PEA could be doing wrong to cause this highly unusual damage. [None of this is recorded anywhere.]

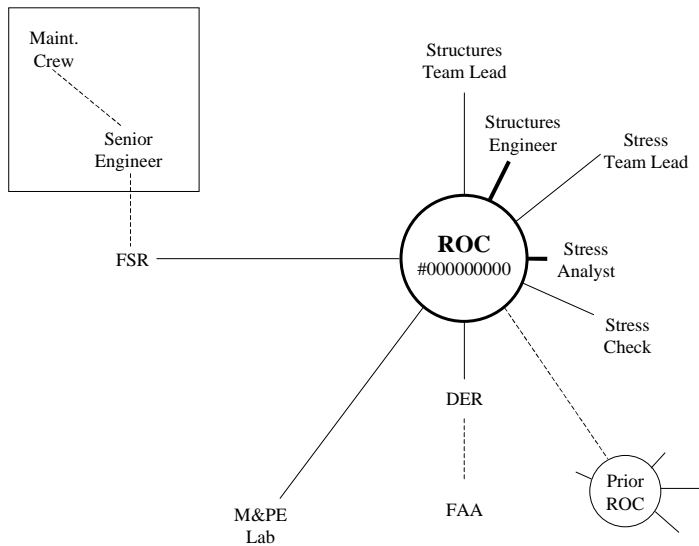


Figure 12: Boundary object diagram

The preceding example brings to light some further reasons why “ROC as boundary object” fails to provide meaningful theoretical coverage for all observed behavior. The two central reasons, each described in a section below, relate to the fact that the seemingly monolithic ROC archive is in fact surprisingly fluid across multiple boundaries, including some temporal, which are very difficult to reconcile with the theoretic notion of a boundary object.

When asked, most engineers at GTS-West view the organizational memory system represented by the ROC archive as a bounded and closed system – it is a database of prior cases that you can search (“It’s STAIRS.”). However, their behavior belies this surface understanding and reveals a very different belief in action. This is that the organizational memories captured by the ROC are actually fragmented throughout the department,

Global, and the customers themselves. Thus, the true process of recontextualization is the reassemblage of enough of these fragments to allow for successful reuse.

The two examples of this fragmented memory from the example above involve the role of personal archives and memory keys embedded in the environment.

10.4 Formal / informal archives

The role of informal archives, those private stashes of notes maintained by individual engineers, play a subtle but important role in augmenting the ROC system. Recall from Chapter Seven Cameron's proprietary personal coversheets on which all of his phone conversations with Zip Express were noted and all of his intermediate scratch work from assembling the details for each individual aircraft were recorded. Or in this example, the role Sam's personal memory played in locating the perfect record for reuse, which was not yet available in the system.

Sometimes the formal and informal cross and sometimes the informal is deliberately injected into the formal. The next example highlights the latter behavior.

A service engineer, wrapping up from a difficult and drawn-out action item resolution, is preparing his personal, informal file of the just completed job. He sorts through what he does not need and what he would like to save for the short term. (He does not tend to keep long-term personal records like some service engineers). A particular e-mail exchange he decides ought to be included in the formal, archival ROC packet. In this

private e-mail, a used part was identified which could be modified to fix the operator's pressing problem. The service engineer explains:

“No, this is from Jason, I think.[Jason is a team lead within Structures.] Jason has this incredible long term memory, like ‘oh fourteen years ago there was this crash and most of the fuselage was fine. It’s stored in this hangar at so and so airport. He knows the whole deal. People call him up all the time – ‘do you know if there’s a...? Do you know where it is?’ That’s what happened here.”

This e-mail exchange did not capture the details of the requisition for the used part, something reasonable one might expect to have included in the ROC. Instead it captured the details of the expert involved in the decision – lending credibility and authority to the unorthodox process (GTS-West is certainly not a used parts broker) and providing a critical component for anyone trying to recontextualize this case later on.

10.4.1 The Folder system

The other component to this intermingling of formal and informal archival systems is the Folder system. As mentioned, the primary focus of this study was on the Short-haul teams. However, both the Long-haul and Heritage groups maintain a team based organizational memory system in addition to the ROC archive – the Folder system.

Prior to the introduction of the ROC system it was the common practice of all three structural Stress teams to maintain folders. Folders are collections of all relevant information regarding each assembly on an aircraft. The information is sorted by assembly and covers information for that part across all aircraft repairs dealing with it across time. For example, a folder on a particular in-board flap joint may include detailed supervisor's notes from the shop floor regarding deviations during manufacture, incident

reports involving the part from the FAA, photographs from early repair jobs, personal notes and annotations from engineers, and, of course, all of the relevant ROC details. Understandably these folders are gigantic and they are housed in extensive file rooms. Their maintenance is prohibitively costly for the high-volume Short-haul team, and they had to abandon updating them over a decade ago. However, the lower-volume teams of Heritage, and to a lesser degree Long-haul, still update and use them regularly.

A Long-haul engineer explains their value in providing richer contextual cues than their equivalent archival ROC entries:

“You go to the folders and find more [than in STAIRS]. A report from Materials, or something from Design, or Propulsion. We’re not experts in everything. It’s important to get this from outside. You might be doing a repair and the airline wants this [physical] material. You check [the folders] and see another material is better. You talk with the Materials guys over there and find out why. It has better properties, longer life, you know, so you can explain it to the operator.”

Explaining what the folder provides that the ROC does not, in addition to the reports from other areas of GTS-West and the plant, he notes “that’s where you find the name. Now they may not be there any more [laughs], but you’ve found the right place [to look]. You can ask your question and someone probably knows.”

A junior Stress engineer from the Heritage team extols the value of the Folders for coming up to speed on his job. The requisite knowledge for working with heritage aircraft is daunting. (Interview questions are italicized and bracketed for context).

“It [the Folder system] is not only good for the job, it is good for me. I learn. It teaches me the history of things, which I need being new. [*Now you may be new, but you’re not **that** new. How long has it been?*] Two years, just over two years. [*That’s longer than some around here.*] “Oh, but not us [Heritage team]. All [are] older and really know it [their material] – fifteen, twenty years. I don’t have that, so I research everything. Always look it up [in Folders].”

10.5 Keys embedded in the environment

The preceding section examined fragments of the organizational memory captured by the ROC but fragmented outside of the ROC archive. This section explores two more extreme situations of fragmentation where the memory exists outside of the organizational boundaries of GTS-West.

10.5.1 Correspondence identifiers

Keys into GTS-West’s organizational memory systems are sometimes embedded, deliberately or inadvertently, by service engineers in the responses returned to operators. For example, in the case above, a faxed document to Pan-Eurasian Airlines by GTS-West in the resolution of a prior job, contained a correspondence identification number⁴⁸ which was of significant assistance to Doris in searching the ROC archive for relevant case candidates for reuse.

Operators frequently try to make sense of fragments, such as the CIN number here, and leverage them in the future. In some situations, such as the example above, this is helpful and welcomed. In others, such as referencing as authoritative “mystical” internal GTS-

⁴⁸ This number belonged to a proprietary system used by GTS-West prior to the introduction of GlobalCOM. GlobalCOM has a similar action item identification number which is more explicitly presented to the operator on all correspondence.

West documents of which they have no knowledge, is very problematic and borderline illegal. One senior service engineer describes the problem as follows:

“[We] don’t give out ROC numbers any more. They [operators] abuse it. [uses mock voice to imitate an operator’s repair technician] ‘Doing repair per ROC blah, blah, blah’ [now visibly quite upset] They have no authority to do that. Each case is unique. When they use over and over ‘against the ROC so-and-so’ they put their own spin on it. We wrote this addressing corrosion and send it out. Three years later they have a crack. ‘We’ll throw a ROC repair on there!’ Burn holes, puncture holes, whatever. It is a completely different set of circumstances! [*How do you even find out that they’re doing this?*] You only catch it with an 8110 request. We assume nothing. Question everything – the operator, fuse, part, measures, everything. It’s a chain of detective work all the way through.”

10.5.2 Serialization

Finally, another very interesting example of externally embedded memory keys is that of the serialized part. Here the key is present in the very physical artifact itself. Serialization is standard in areas such as HydraMech where the components and assemblies are independent of individual aircraft (e.g. landing gear) and interchangeable. However, serial numbers are also used in Structures to track exceptional parts. These are parts, often “one of” manufactures, that were allowed to be installed under special circumstances with explicit restrictions. Knowing that once parts are out in the marketplace Global has no control over them, each of these exceptional parts is etched with the ROC number corresponding to its genesis. The number is meaningless to the operator, but deeply meaningful to GTS-West. Instructions are given to contact the manufacturer and give this number, which GTS will recognize as a ROC and pull the full history of the part, including all of the constraints surrounding its exceptional approval.

A service engineer who had just serialized a part in the resolution to a particularly complex job resolution described the role of the serialized part in this case:

“That part was reworked so extensively, it’s not even close to the configuration. Off the airplane [this specific aircraft], no, is it much of... you know, [pause] [*A bit of a monster?*] Exactly. It is not like anything we have seen before. We didn’t know what to call it, so we had them serialize it by ROC number. So the next time some guy calls in, ‘we’ve got a weird part number here, oh, oh, dash, oh, eight, dash, ... [00-08-...] then we know, oh, that’s a ROC number. We can go look it up and get the whole story. It’s a real Heinz 57!”

He also described the importance of serialized ROC number etching in promoting safety.

The part was designed to be a short-term repair. Once the airline repairs the part with a permanent replacement, what is the likely fate of this unique creation?

“They’ll probably sell it to some unsuspecting third party. There are so many problems with spares. There simply are no new parts. We don’t have any of these doors in stock. It’s a high heat area, so that’s difficult. There’s not a lot you can do, so here we have it, really thick skin, non blueprint stiffeners, and all.”

This chapter has examined the complexities of interacting with boundary objects across time. It engaged the problems that occur with both the processes of decontextualization and recontextualization, pointing to some additional limitations of the boundary object concept. The following chapter will address design considerations which address the collection of issues regarding the use of boundary objects in reuse that have been raised in this and the preceding three chapters.

CHAPTER 11: Implications for Design

11.0 Introduction

While the primary goal of this research has been to better understand the routine behaviors of reuse in collaborative problem solving situations, there has been an underlying interest in informing the design of socio-technical systems to support this behavior. This chapter will review some of the implications for design which are suggested by the results of this study. As will be mentioned in the final chapter, the design implications are one of the most fruitful avenues for future research.

Much was found which supports some of the design considerations presented in earlier work such as: the value of free text search (Blair, 1985;1996), the importance of fostering a culture of collaboration (Nonaka and Takeuchi, 1995; Davenport and Prusak, 1998), the necessary support for heterogeneous resources, and the integration of both technical and social mediation (Ackerman and Halverson, 1998;1999;2000).

11.1 Designing for negotiation points

In projects whose goal it is to support, augment, or improve existing work practices, this research shows that it is extremely valuable to examine boundary points in an organization's information flows. These can provide key insights for critical points in the software design lifecycle.

Initially a boundary point analysis highlights the richest confluence of contextual information in the flow. Information artifacts at this juncture, boundary objects, are the most valuable for understanding current state and preserving for future reuse. This is the tightest intersection point among multiple participants and focused on the key elements of translation amongst them.

Stakeholder analyses and process mapping are now routine activities during the requirements gathering phase. An examination of boundary points extends this understanding by highlighting the interaction of these two. This yields an awareness of the routine negotiation processes, common breakdowns, and standard repairs.

During the design process the boundary object concept accentuates the process constraints and performance expectations for varying information pathways. This can help assess the appropriate prioritization. Also, an understanding of these constraints can point to opportunities for iterative real world evaluation – key relationships for prototyping, beta testing, etc.

Lastly, an understanding of boundary objects is valuable in deployment and evaluation. These negotiation points are likely to be the greatest points of friction and resistance in the system. Paying special attention here can really pay off. The understanding of the intersection point of different communities of practice also helps to understand adoption trends (who is benefiting the most, using the most, and why).

11.2 Designing for recontextualization

As Ackerman and Halverson (1998;1999;2000) noted, the problems of recontextualization are the core problems of reuse in collaborative problem solving. This study validates that finding and suggests measures to improve the processes of recontextualization which will be discussed in this section. Each of these will be presented within the specific context of this field site.

Prior to this discussion, it is important to note that the goal of this IT support is to augment personal abilities, not to automate the reuse process. These processes of interpretation, evaluation, and recontextualization are the result of a complex interplay between tacit and explicit knowledge distributed cognitively throughout a team of experts. While routine diagnostic activities have been shown to benefit from limited automation exercises, such as medical expert systems, there is little at GTS-West that can be deemed routine and repeatable (with the possible exception of skin corrosion blendouts). Each case is subtly but meaningfully different, and an exact match with a rule or historical record is almost non-existent. These processes then of reasoning by analogy are often driven by intuition – tacit knowledge supported by experience over time, which can only be socially mediated. What then is the role of IT in supporting these endeavors?

Primarily, IT provides the malleable substrate for storage and retrieval. However, it can also offer much more, while still remaining shy of outright automation. These contributions are through enhancing the decontextualization and recontextualization

processes by means of the lightweight preservation of context. Four of these approaches, borne out of the analyses of the work practices at GTS-West, are presented here.

11.2.1 ROC Completeness

As mentioned in Chapter Eight, one of the primary tensions between Structures and Stress involves the issue of completeness. Completeness means different things to the two different teams and the resulting conflict is routinely resolved on a case-by-case basis through the boundary object of the ROC. While attempts have been made at standardizing and routinizing this process (e.g. the completeness guidelines provided by Stress), little has been as effective as the actual format of the data fields in the ROC itself – the lowest common information denominator between the two groups.

Enforcing ROC completeness through the code-level means would be highly problematic. For example, this could prohibit a ROC from being sent from Structures to Stress without having all fields correctly completed. As shown in the data, this is a troublesome design decision. There are multiple situations where omissions are deliberate and warranted (e.g. airline refuses to provide data), and GTS-West has work arounds to ensure that these jobs can be completed regardless (e.g. extra conservative repairs). To prohibit the processing of these jobs would drastically affect these critical work arounds.

However, even a less draconian system-based completeness enforcement program is problematic. There is the tension between the informal and formal representations of information and thereby at GTS-West. Since the ROC is a formal record, requiring the

inclusion of informal information is inappropriate. For example, the stipulation that all due dates be correct, precludes the situation where Intra-Asian Island Air's jobs need to be accelerated due to the fact that the Structures manager is currently visiting with IAIA officials in Indonesia. Another example would be when working figures are "off the record" and need to be delivered in person. Recall the tension that exists between reliability's need for 'absolutely complete and accurate information' and the liability risks entailed by this.

The solution to this is the inclusion of automatic checking systems or agent-based critics (Fischer, et al., 1992). These can notice omissions, irregularities, inconsistencies, or errors and flag them for review by the Structures engineer prior to sending the job to Stress. If not resolved, these flags can remain visible for review by the Stress team lead or analyst, who can use them as prompts for discussion with the Structures engineer as needed. These critics can not only provide automated verification (e.g. ensure that measurements are within bounds), but also "intelligently" suggest corrections or alternatives in place.

The improved correctness and completeness of the ROC would not only enhance the workability of the contemporary ROC but also the processes of retrieval and reuse later on. Improving the quality and reliability of the information in the archival ROC is of tremendous assistance in recontextualization.

An aside to this discussion involves search. Recall that the archival ROC system is a STAIRS database utilizing a free-text search engine. A routine problem in retrieval is finding the appropriate records. A degree of standardization could be introduced into the ROC creation and modification phases assistance by providing pre-population of fields or auto-complete mechanisms. As with the critics mentioned above these should support the engineer's expertise in ROC creation, not supplant it. Flexibility must remain. However, given the degree of important information lost in the processes of decontextualization currently due to ambiguous and sporadically enforced group standards, intense time pressures, and notable system limitations, the payoff from such improvements would be substantial.

11.2.2 Resource coverage

As noted, once a reasonable match is found in the ROC archive, a tremendous amount of time is spent trying to reconstruct the prior case and adapt it to the task at hand. As seen, much of this effort involves physically locating resources (e.g. technical journals, stress reports and blueprints), certifying their correctness, and evaluating their authority. The time spent on these tasks for any given action item is frequently greater than that spent on the numerical analysis itself.

The fact that resources in the department's information space are rapidly becoming digitized (e.g. the stress report project), should be leveraged by embedding hyperlinks for cited reference material into the ROC. Reducing the frustrating complexity and length of

the recontextualization process by this means would not only increase the quality and frequency of reuse, but also further reduce call turn-around time.⁴⁹

Also, as already begun in the stress report digitization project, experts within the teams could be identified to coordinate and maintain the hyperlinked resource bases within their specialization. As this project is the first of its kind in the department, there is little hard evidence for the tradeoffs between the benefits of these easily accessible resources and the costs of maintenance.

11.2.3 Similar cases

Given the problems with recontextualization, such as incomplete searches, missing resources, and poorly aged ROC files, engineers do not trust a single hit and usually try to cull multiple prior examples to examine. As outlined in many of the cases previously discussed, engineers routinely examine a subset of matching cases (usually 3-5) to find the most appropriate fit, if it exists.

As the most common search criteria are usually known at the moment a new call arrives; automatically data mining the ROC archive from these would be easy. This would provide the engineer with a higher quality sub-set of matching cases to evaluate in reduced time.

⁴⁹ While few engineers enjoy this maddening daily document hunt, there are elements of worth in the process – namely the mental break from analysis or call handling and the forced interaction with colleagues. However, the former could be addressed by minor work schedule modifications and the latter

Given that Stress routinely re-does the ROC archive searches performed by Structures to verify that a more appropriate prior case was not missed or to more expertly refine their search, unobtrusively providing pointers to the cases returned from the data mining exercise and evaluated by Structures would jumpstart Stress' evaluation process by providing more robust contextual details for their search rationale. As this process is a critical redundancy check required in high-reliability, high-liability work, it is a process that should not be removed or replaced, only provided additional resources such as this to ease the cognitive burden for the analyst.

11.2.4 Awareness of knowledge/expertise network

Less vexing than the routine resource hunt, but no less critical is the search for appropriate expertise. While employee turnover in Stress is relatively low, it is higher in Structures, and throughout the department and other departments relied upon for consultation. One does not have to search far back into the ROC archive, only a couple of years, before personal contacts grow stale. In most cases engineers ask more senior engineers for expertise recommendations (e.g. 'looks like Sally must have been the expert on nose cone corrosion, who does this now?'). While this process should not be circumvented (often the consultation may yield a further refining of the problem or valuable instruction in local history), it could be better supported. This would allow the simple, most direct cases to be handled directly without consultation, and for those that

need not be removed entirely only reduced. Once relevant resources had been found there may still be the need to interact with others in correctly interpreting and adapting them.

require consultation, better information will be provided, both of which would reduce the load on the local senior engineers.

One way of accomplishing this would be to modify the ROC archive system to better represent the expertise network of the organization. In a manner similar to Answer Garden (Ackerman, 1993), the system would have links to the relevant experts for each case. As it is now, these would most likely be the originating engineers and analysts. However, these links could be updated automatically to match the ever-shifting network of expertise. For example, clicking through on a contact link for an original engineer who has since left the company, may yield suggestions with contact information for three other likely engineers with similar expertise regarding that particular case or the details it contains. These contact pointers could even assist in navigating the network through an iterative series of recommendations and escalations akin to that proposed with Expertise Recommender (McDonald and Ackerman, 1998;2000). McDonald and Ackerman's research also suggests manners in which such a recommender could automatically leverage existing digital information sources to reduce the burden of the maintenance required to keep the system current.

11.3 Designing for inter-organizational boundary spanning

As discussed in Chapter Nine, there can be significant value in extending organizational memory systems beyond the boundaries of the organization. Maintaining a semi-permeable membrane, so that only certain types of information can flow and ownership remains with the original organization, allows for leveraging the reuse behaviors of the

other organizations. For example, recall the case from Chapter Ten where the savvy senior engineer from PEA was able to work with the local FSR to uncover a correspondence identification number which was the key to reuse in that situation.

There are clear concerns about extending the organizational memory too far outside of the organization such as losing ownership, control, and interpretive power. There are issues of intellectual property, reliability (who outside of the organization can verify the accuracy of the interpretation of a prior case?) and liability (what if an operator misuses GTS's information).

One solution to handling formal information has long been in place – the Structural Repair Manual, and as mentioned, there are efforts underway to increase its scope. But what of coverage for more informal or case-by-case information, as represented by the ROC archive for example? A design option could be to provide designated contacts within the customer organizations with keys back into the system. These individuals would have some characteristics of Tushman's boundary spanning individuals (1977;1980;1981) and of Orr's key operators (1986), notably that they would have some knowledge on how to "talk the talk" with service engineering. . This contact would have enough information to significantly ease the reuse efforts back at GTS-West, while not possessing enough information to allow the customer to continue on independently.

11.4 GlobalCOM integration revisited

Recall in Chapter Five that one of the motivating factors for this research was examining the impending integration of the ROC system into the GlobalCOM system. Both Stress and Structures engineers had successfully lobbied to delay this integration as long as possible to better gain an understanding of its implications. While many of the technical systems-level concerns are clear, such as providing better search capability, a lengthier archive, alterations to key forms, and modifications to the represented workflow, many of the social issues are not. As this change is practically inexorable, it is valuable to understand what may be altered.

A move to a future GlobalCOM instantiation of the ROC system will shift control over GTS-West's archival system from local administration to the next higher organizational layer, the general GTS offices. While this will certainly provide improved technical support and visibility, the change could also be catastrophic if not handled correctly. The most important effect of such a move would be the reallocation of political power away from Stress, though not necessarily toward Structures, and significantly reduce the gatekeeper role played by the team leads. While the ROC is the negotiating point for multiple parties, Stress' ownership of the key ROC processes (e.g. editing, archiving) give them the final authority. This is appropriate as the information contained in the ROC is the most critical for Stress' information needs. Stress engineers are also the most likely audience for future reuse, so the filtering and modification done by the team leads helps to ensure the appropriateness of the archival record, reducing the required effort for later recontextualization and reuse. GlobalCOM's data model is flexible enough to handle any

special needs that Stress may have, such as including additional information in the record, however its workflow model does not support the structures required to ensure the high degree of reliability of this information -- it is currently too loose and too general.

There are similar roles in both systems, but subtle variations which make big differences. For example, both systems have a notion of a supervisor or team lead, who accepts incoming jobs and assigns them to the most appropriate engineer. However, the GlobalCOM model does not support any of the editing or archival features mentioned above, or the feedback loops for verification so critical to the redundant check systems required by Stress. The value of that information would change as a result. While the recommendation and assignment of a job to a particular engineer by a team lead could still signify that this was the most appropriate engineer for the job, no longer would it carry the understanding that this work had been verified by that team lead as correct. One level would be removed.

CHAPTER 12: Conclusions

12.0 Reuse at GTS-West

The story of Global Airframe's Technical Services West (GTS-West) division is a story of remarkable success, borne in no small part by innovative knowledge management practices. GTS-West has managed to leverage its pre-existent organizational culture of knowledge sharing and archival dependency to address major current challenges (e.g. increasing call load with decreasing resources). These practices were forged by crisis and necessity. Not only has GTS-West survived the onslaught of these new challenges, but they have been able to out perform other GTS offices, with only a fraction of their resources and infrastructure.

12.1 Review of key findings

While this study has uncovered much that is valuable to the continued maturation of the knowledge management efforts at GTS-West it also speaks to a set of more universal concerns. The central theoretical findings of the study are summarized in the following sections.

12.1.1 Integrating boundary objects for knowledge management

The selection of appropriate information artifacts to preserve for later reuse is a core knowledge management concern. One needs to retain critical information elements and

contextual content, without provoking information overload for the employees. The successful identification of these artifacts from the usually vast information workspace is a known challenge. One heuristic offered by this study is the value in identifying primary work artifacts which exist at negotiation points and which mediate collaboration among diverse participants – boundary objects.

The selection of an artifact already in the existing work flow is efficient as no additional preparation is required. It is grounded, ensuring that there is good preservation of context. It is also information rich, in that it can represent the perspectives of multiple participants simultaneously.

In addition, it can be more valuable to preserve such an artifact itself (e.g. the actual work product) than a post-hoc rationalization (e.g. after-action report). The work artifact best reflects the formal and informal information flows that were required to complete the job, important details commonly stripped out of later reports. The work artifact also contains more robust pointers into the knowledge network of the organization than a summary written by any individual party. Lastly, it preserves elements of the negotiation without suffering the implicit biases of any particular participant.

As was presented in Chapters Seven and Eight, GTS-West has been able to integrate an existing information artifact in their routine work flows, the ROC, to form the foundation for their successful knowledge base and reuse practices.

12.1.2 Leveraging boundary objects to enhance compliance

Another primary concern with knowledge management systems is compliance, ensuring enough routine participation by employees to achieve critical mass and cross the utility threshold for the system. This study has shown that by leveraging boundary objects which are primary work artifacts, participation is near universal. As shown in Chapters Seven through Ten the ROC mediates practically all collaboration between Structures and Stress. Preparation for reuse, often an onerous burden in systems requiring the preparation of additional reflective materials, is practically a peripheral concern at GTS-West. With only minor editorial control by a gatekeeper, the Stress team leads, the archival artifact is assured to be of reasonably high quality and authority with minimal effort on the part of the individual engineers.

Also, by maintaining a culture dependent on reuse, the GTS-West engineers frequently experience the value of reuse on a personal level. This is a strong motivator for compliance. (If you might be the next benefactor of an archived ROC, you are more apt to be attentive when assisting in decontextualizing it.)

12.1.3 Gatekeepers

This study highlights the importance of key individuals who represent intersection points of information flows, work flows and knowledge/expertise networks within the organization. This analysis supports current notions of the roles of gatekeepers (Allen, 1997), boundary spanners (Tushman, et al. 1977;1980;1981) and expertise concierges (McDonald and Ackerman, 1998;2000), while expanding them to examine the

individual's role in creating, maintaining and interpreting archival material. Chapter Eight examined in detail one such class of individuals at GTS-West, the Stress team leads, highlighting the processes of enactment of perceived authority and reliability which occur between these individuals and the ROC. This enactment process was critical for further reuse of the ROC; here, the gatekeeper role functioned for these ROCs over time as well as in the moment.

12.1.4 Methods of enhancing recontextualization

The better "preserved" an archival boundary object is, the more facile the process of adapting it for later reuse (recontextualization). This does not mean overloading the artifact with more data or more comprehensive coverage. Instead it means expertly distilling the information which best represents its path through the knowledge network. In a way this is preserving the information fingerprints, instead of the content itself. Chapter Nine showed how this behavior is exemplified by the economical yet powerful ROC summary.

This finding holds the greatest promise for improving the performance of knowledge management systems through enhancements to the processes of decontextualization. First, any given ROC needs to be "complete," meaning that it contains essential information in an agreed upon format. Automated checking of ROCs would improve the baseline of the archive. Second, while distillation is valuable (not everything should be preserved), addressing unknown future reuse scenarios makes this process difficult. Coping with lost context is the central problem in recontextualization. The ROC system would benefit

from augmentation with support for lightweight preservation of context. Detailed recommendations were presented in Chapter Eleven.

12.1 Extreme organizational stressors revisited

Other factors unique to GTS-West come to bear on these general findings, of which two will be briefly discussed here: reliability and liability. While valuable in highlighting the behaviors under investigation in this study, they are a potential impediment to the generalizability of the findings.

12.1.1 High-Reliability

Recall that GTS-West is a classic example of a high-reliability organization. These are organizations with no tolerance for error, where even the slightest mistake can have catastrophic consequences. As a result of their high-stakes operation, GTS-West has multiple layers of redundancy designed into their procedures. These layers are clearly represented in the ROC, providing assurance of the credibility of any given record. The engineers regularly assume that once a ROC is archived, there is a high probability that the material it contains is error-free, having been checked by multiple individuals and passed through the final editorial review by the team leads. (Problems in recontextualization stem from incorrect interpretations based on unclear or omitted information, not erroneous information.)

It is clear that the assurance of such high-quality, high-relevance archival material is a tremendous boon to any knowledge management effort. Could this be one of the leading

factors in the successful reuse practices at GTS-West? Obviously, the answer is yes. However, the degree to which this is the case warrants further investigation. In particular studying sites with a more degraded knowledge base, but with similar organizational and cultural characteristics would be of great utility.

12.1.2 High-Liability

Liability is also a key concern as GTS-West operates in a regulated environment. Remember that the ROC serves as the FAA-auditable centerpiece, the formal “legal paper trail” required to be kept for every job. Personal responsibility is explicit at every step in the routine resolution of repair requests, and this tracking is facilitated by the ROC system. This has two primary effects.

The first is that the ROC contains a degree of explicit personal information (e.g. pointers to the experts consulted), far higher than most equivalent information artifacts at other organizations. While most boundary objects do contain pointers to the individuals involved in the negotiations, as in Mambrey and Robinson’s study of the speech draft (1997), they do not exhibit the deliberateness, weight and authority that they do with the ROC. This is tremendously advantageous for the processes of recontextualization and later reuse.

The second is that the legal nature of the ROC system limits the range of informal information which can be integrated within the ROC, and thus archived for future reuse. The organizational work around for this limitation is that the ROC does provide pointers

to information resources outside of the system which do capture some informal details, such as technical journals, personal filing systems and memory, but this is weak. It would be valuable to examine other knowledge management systems which capture more informal information in order to gauge the severity of the limitations, if any, in the ROC system.

12.2 Beyond boundary objects

As presented in Chapter Three, the boundary object concept appeared to be the most appropriate analytic construct for the examination of reuse behavior at GTS-West. This assumption was borne out well in Chapters Seven and Eight, which foregrounded the negotiations surrounding the ROC in routine work. However, limitations in the boundary object concept began to surface in Chapters Nine and Ten.

12.2.1 Representing the embedded nature of boundary objects

In Chapter Nine it was shown that the “objectness” of boundary objects limited what could be represented. Clearly the design and use of the ROC system supported this object-centric view of work practices. However, where there were breakdowns with the ROC in reuse, the concept of boundary objects faltered as well. The information flows surrounding the negotiation, coordination and collaboration between the operators, Structures, Stress and other relevant participants in any given repair were dynamic and complex. Yet the ROC was relatively static and simplified by comparison. Most difficulties in recontextualization trace their roots to this fact. (It is important to note that

many of the benefits of this system share the same root cause – an authoritative yet simplified snapshot of negotiation resolved.)

Regardless of how the ROC system develops in the future to handle this, the boundary object concept clearly needs to be expanded to better represent the embeddedness of these objects in complex and dynamic streams.

The closest development in the literature, Kim's notion of boundary instances, does not address this problem. He completely abandons the object nature of these boundary points for certain circumstances (e.g. engaging negotiations that are not yet represented by physical information artifacts). This is a meaningful extension, but it is not relevant to the observed behavior at GTS-West where clearly the negotiation points are instantiated at some level in documents such as the ROC, but are merely too static and simplified to be a reasonable analytic representation. Further theoretic development is required to address this limitation in the boundary object concept.

12.2.2 Representing the trajectory of use for boundary objects

The engagement of temporal issues in Chapter Ten uncovered additional limitations in the boundary object concept. As contextual cues regarding the central negotiations diffuse and decay over time, one would expect that the boundary object would rapidly lose utility. This anticipated cycle was exhibited at GTS-West; the older the ROC, the less likely that it would be reused to assist a contemporary problem. However, the degree to which old ROCs were able to be reused was astounding. The boundary object concept

has no explanation for how this could occur – once the details of the negotiation are lost, they are gone. How then does GTS-West accomplish this? They do so by maintaining a sense of trajectory for the boundary object, both embedded socially in the group and technically in the ROC. There are a set of expectations regarding future reuse which guide the processes of decontextualization. Relevant resource pointers placed in the document, information stored outside of the system, and personal memories, especially by team leads, all play a critical role in recontextualization.

The concept of boundary objects would benefit from a deeper examination of temporal issues, most notably this understanding of expected trajectory for the information artifact which influences every aspect of its use from creation to archive to later reuse.

REFERENCES

- Ackerman, Mark S. (1993). "Answer Garden: A Tool for Growing Organizational Memory," *Dissertation*, Massachusetts Institute of Technology.
- Ackerman, Mark S. (1994). "Definitional and Contextual Issues in Organizational and Group Memories," *Proceedings of the Twenty-seventh IEEE Hawaii International Conference of System Sciences*, pp. 191-200.
- Ackerman, Mark S. (1996). "Definitional and Contextual Issues in Organizational and Group Memories," *Information Technology and People* (9:1), pp. 10-24.
- Ackerman, Mark S., and Halverson, Christine. (1998). "Considering an Organization's Memory," *Proceedings of the ACM Conference on Computer-Supported Cooperative Work*, pp. 39-48.
- Ackerman, Mark S., and Halverson, Christine. (1999). "Organizational Memory: Processes, Boundary Objects, and Trajectories," *Proceedings of the IEEE Hawaii International Conference of System Sciences*.
- Ackerman, Mark S., and Halverson, Christine. (2000). "Re-examining Organizational Memory," *Communications of the ACM* (43:1), pp. 59-64.
- Ackerman, Mark S. (2000). "The Intellectual Challenge of CSCW: The Gap Between Social Requirements and Technical Feasibility," *Human-Computer Interaction* (5:2/3), pp. 181-203.
- Allen, Thomas J. (1977). *Managing the Flow of Technology*, Cambridge, MA: MIT Press.
- Auramaki, Esa, Mike Robinson, Ann Aaltonen, Mikko Kovalainen, Arja Liinamaa and Taina Tuuna-Väiskä. (1996). "Paperwork at 78kph," *Proceedings of the ACM Conference on Computer-Supported Cooperative Work*, pp. 370-379.
- Bannon, Liam, and Kari Kuutti. (1996). "Shifting Perspectives on Organizational Memory: From Storage to Active Remembering," *Proceedings of the Twenty-ninth IEEE Hawaii International Conference of System Sciences*, pp. 156-167.
- Bechky, Beth A. (1999). "Crossing Occupational Boundaries: Communication and Learning on a Production Floor," *Dissertation*, Stanford University.
- Bently, R., J. A. Hughes, D. Randall, T. Rodden, P. Sawyer, D. Shapiro and I. Sommerville. (1992). "Ethnographically-informed Systems Design for Air Traffic Control," *Proceedings of the ACM Conference on Computer-Supported Cooperative Work*, pp. 123-129.

- Berg, Marc and Geoffrey Bowker. (1997). "The Multiple Bodies of the Medical Record: Toward a Sociology of an Artifact," *The Sociological Quarterly* (38:3), pp. 513-537.
- Blair, David C., and Maron, M. E. (1985). "An Evaluation of Retrieval Effectiveness for a Full-Text Document-Retrieval System," *Communications of the ACM* (28:3), pp. 289-299.
- Blair, David C. (1996). "STAIRS Redux: Thoughts on the STAIRS Evaluation, Ten Years After," *Journal of the American Society for Information Science* (47:1), pp. 4-22.
- Bowker, Geoffrey C., et. al., eds. (1997). *Social Science, Technical Systems and Cooperative Work: Beyond the Great Divide*, Hillsdale, NJ: Lawrence Erlbaum.
- Bowker, Geoffrey C. and Susan Leigh Star. (1999). *Sorting Things Out: Classification and Its Consequences*, Cambridge, MA: MIT Press.
- Carstensen, Peter H. and Ulrika Snis. (2000). "On Knowledge Management: A Field Study." In Schartz, David G., Monica Divitini and Terje Brasethvik. *Internet-Based Organizational Memory and Knowledge Management*, Hershey, PA: Idea Group Publishing, pp. 170-199.
- Cicourel, A. V. (1990). "The Integration of Distributed Knowledge in Collaborative Medical Diagnosis," in Galegher, Jolene, Kraut, Robert E. and Egido, Carmen (eds.) *Intellectual Teamwork: Social and Technological Foundations of Cooperative Work*, Hillsdale, NJ: Lawrence Erlbaum, pp. 221-242.
- Davenport, Thomas H. and Prusak, Laurence. (1998). *Working Knowledge: How Organizations Manage What They Know*, Boston: Harvard Business School Press.
- Dourish, Paul, Victoria Bellotti, Wendy Mackay and Chao-Ying Ma. (1993). "Information and Context: Lessons from a Study of Two Shared Information Systems," *Proceedings of the ACM Conference on Organizational Computing Systems*, pp. 42-51.
- Emerson, Robert M., Rachel I. Fretz and Linda L. Shaw. (1995). *Writing Ethnographic Fieldnotes*. Chicago: The University of Chicago Press.
- Fahey, Liam and Laurence Prusak. (1998). "The Eleven Deadliest Sins of Knowledge Management," *California Management Review* (40:3), pp. 265-276.
- Fischer, G., Girgensohn, A., Nakakoji, K., Redmiles, D. (1992). "Supporting Software Designers with Integrated Domain-Oriented Design Environments," *IEEE Transactions on Software Engineering, Special Issue on Knowledge Representation and Reasoning in Software Engineering* (18:6), pp. 511-522.
- Halverson, Christine A. (1995). "Inside the Cognitive Workplace: New Technology and Air Traffic Control," *Dissertation*, University of California, San Diego.

- Harper, Richard H. R. (1998). *Inside the IMF: An Ethnography of Documents, Technology and Organisational Action*, New York: Academic Press.
- Heath, Christian and Paul Luff. (1991). "Collaborative Activity and Technological Design: Task Coordination in London Underground Control Rooms," *Proceedings of the Second European Conference on Computer-Supported Cooperative Work*, pp. 65-80.
- Heath, Christian and Paul Luff. (1992). "Collaboration and Control: Crisis Management and Multimedia Technology in London Underground Line Control Rooms," *Computer Supported Cooperative Work* (1), pp. 69-94.
- Henderson, Kathryn. (1991). "On Line and on Paper: Visual Representations, Visual Culture, and Computer Graphics in Design Engineering," *Dissertation*, University of California, San Diego.
- Henderson, Kathryn. (1991). "Flexible Sketches and Inflexible Data Bases: Visual Communication, Conscriptio Devices, and Boundary Objects in Design Engineering," *Science, Technology and Human Values* (16:4), pp. 448-473.
- Henderson, Kathryn. (1991). "Introduction: Social Studies of Technical Work at the Crossroads," *Science, Technology and Human Values* (16:2), pp. 131-139.
- Henderson, Kathryn. (1998). "The Role of Material Objects in the Design Process: A Comparison of Two Design Cultures and How They Contend with Automation," *Science, Technology and Human Values* (23:2), pp. 139-174.
- Henderson, Kathryn. (1998). *On Line and on Paper: Visual Representations, Visual Culture, and Computer Graphics in Design Engineering*. Cambridge, MA: MIT Press.
- Hutchins, Edwin L. (1995). *Cognition in the Wild*, Cambridge, MA: MIT Press.
- Grudin, Jonathan. (1988). "Why CSCW Applications Fail: Problems in the Design and Evaluation of Organizational Interfaces," *Proceedings of the ACM Conference on Computer Supported-Cooperative Work*, pp. 85-93.
- Grudin, Jonathan. (1994). "Groupware and Social Dynamics Eight Challenges for Developers," *Communications of the ACM* (37:1), pp. 92-105.
- Katzenberg, Barbara, Fred Pickard and John McDermott. (1996). "Computer Support for Clinical Practice: Embedding and Evolving Protocols of Care," *Proceedings of the ACM Conference on Computer-Supported Cooperative Work*, pp. 364-369.
- Kim, Jeffrey Y. (2000). "Knowledge Management in DRAM Fabrication." *Dissertation*, University of California, Irvine.
- Kovalainen, Mikko, Mike Robinson and Esa Auramaki. (1998). "Diaries at Work," *Proceedings of the ACM Conference on Computer-Supported Cooperative Work*, pp. 49-58.

- Lofland, John and Lyn H. Lofland. (1995). *Analyzing Social Settings: A Guide to Qualitative Observation and Analysis (Third Edition)*. New York: Wadsworth Publishing Company.
- Mackay, Wendy E. (1999). "Is Paper Safer? The Role of Paper Flight Strips in Air Traffic Control," *ACM Transactions on Computer-Human Interaction* (6:4), pp. 311-340.
- Mambrey, Peter and Robinson, Mike. (1997). "Understanding the Role of Documents in a Hierarchical Flow of Work," *Proceedings of the ACM Conference on Supporting Group Work*, pp. 119-127.
- Markus, M. Lynne. (1983). "Power, Politics, and MIS Implementation," *Communications of the ACM* (26:6), pp. 430-444.
- McDonald, David W. and Ackerman, Mark S. (1998). "Just Talk to Me: A Field Study of Expertise Location," *Proceedings of the ACM Conference on Computer-Supported Cooperative Work*, pp. 315-324.
- McDonald, David W. and Ackerman, Mark S. (2000). "Expertise Recommender: A Flexible Recommendation System and Architecture," *Proceedings of the ACM Conference on Computer-Supported Cooperative Work*, pp. 231-240.
- McDonald, David W. (2000). "Supporting Nuance in Groupware Design: Moving from Naturalistic Expertise Location to Expertise Recommendation." *Dissertation*, University of California, Irvine.
- Nonaka, Ikujiro and Hirotaka Takeuchi. (1995). *The Knowledge-Creating Company: How Japanese Companies Create the Dynamics of Innovation*. Oxford University Press.
- O'Day, Vicki L., Daniel G. Bobrow and Mark Shirley. (1996). "The Social-technical Design Circle," *Proceedings of the ACM Conference on Computer Supported Cooperative Work*, pp. 160-169.
- Orr, Julian E. (1996). *Talking About Machines: an Ethnography of a Modern Job*, Ithaca, NY: Cornell University Press.
- Orlikowski, Wanda J. (1992). "Learning from NOTES: Organizational Issues in Groupware Implementation," *Proceedings of the ACM Conference on Computer Supported-Cooperative Work*, pp. 362-369.
- Pentland, Brian. (1992). "Organizing Moves in Software Support Hotlines," *Administrative Science Quarterly* (37:4), pp. 527-548.
- Prusak, Laurence, ed. (1997). *Knowledge in Organizations*. Boston: Butterworth-Heinemann.
- Radding, Alan. (1998). *Knowledge Management: Succeeding in the Information-based Global Economy*. Charleston, SC: Computer Technology Research Corp.

Robinson, Mike, Mikko Kovalainen, Mikko, and Esa Auramaki. (2000). "Diary as Dialogue in Papermill Process Control," *Communications of the ACM* (43:1), pp. 65-70.

Schmidt, Kjeld and Carla Simone. (1996). "Coordination Mechanisms: Towards a Conceptual Foundation of CSCW Systems Design," *Computer Supported Cooperative Work* (5), pp. 155-200.

Star, Susan Leigh. (1989). "The Structure of Ill-Structured Solutions: Boundary Objects and Heterogeneous Distributed Problem Solving." In Gasser, Les and Michael N. Huhns. *Distributed Artificial Intelligence – Volume II*. Morgan Kaufmann Publishers, Inc.

Star, Susan Leigh and Griesemer, James R. (1989). "Institutional Ecology, 'Translations' and Boundary Objects: Amateurs and Professionals in Berkeley's Museum of Vertebrate Zoology, 1907-39," *Social Studies of Science* (19:3), pp. 387-420.

Strauss, Anselm L. and Corbin, Juliet. (1990). *Basics of Qualitative Research: Grounded Theory Procedures and Techniques*, Newbury Park, CA: Sage Publications.

Strauss, Anselm L. (1993). *Continual Permutations of Action*. New York: Aldine de Gruyter.

Suchman, Lucy A. (1983). "Office Procedure as Practical Action: Models of Work and System Design," *ACM Transactions on Office Information Systems* (1:4), pp. 320-328.

Suchman, Lucy A. (1987). *Plans and Situated Actions: The Problem of Human-Machine Interaction*. Cambridge University Press.

Tushman, Michael L. (1977). "Special Boundary Roles in the Innovation Process," *Administrative Science Quarterly* (22), pp. 587-605.

Tushman, Michael L. and Ralph Katz. (1980). "External Communication and Project Performance: An Investigation into the Role of Gatekeepers," *Management Science* (26:11), pp. 1071-1085.

Tushman, Michael L. and Thomas J. Scanlan. (1981). "Boundary Spanning Individuals: Their Role in Information Transfer and Their Antecedents," *Academy of Management Journal* (24:2), pp. 289-305.

VanMaanen, John. (1988). *Tales of the Field: On Writing Ethnography*. Chicago: The University of Chicago Press.

Walsh, James P. and Gerardo Rivera Ungson. (1991). "Organizational Memory," *Academy of Management Review* (16:1), pp. 57-91.

Weick, Karl E. (1969). *The Social Psychology of Organizing*. Reading, MA: Addison-Wesley.

Weick, Karl E. and Roberts, Karlene H. (1993). "Cognitive Mind in Organizations: Heedful Interrelating on Flight Decks," *Administrative Science Quarterly* (38), pp. 357-381.

Wenger, Etienne. (1998). *Communities of Practice: Learning, Meaning and Identity*. Cambridge University Press.

Wolcott, Harry F. (1990). *Writing Up Qualitative Research*. Newbury Park, CA: Sage Publications.