

HL7 Context Management Specification

[Technology and Subject-Independent Component Architecture](#)

[Component Technology Mapping: ActiveX](#)

[Data Definition: Patient Subject](#)

[Data Definition: User Subject](#)

[User Interface: Microsoft Windows OS](#)

1
2
3
4
5
6
7
8
9
10
11

12
13
14
15
16
17
18
19
20
21

Health Level Seven Standard

Context Management Specification
Technology- and Subject-Independent Component Architecture
Version CM-1.0

DOCUMENT ID: HL7SIGVI_3_1_99
REVISION ID: March 17, 1999
FILE NAME: hl7_sigvi_arch_cm_1_0.doc
SUPERCEDES: n/a

Copyright 1999 Health Level Seven

1 Contents

2	1 INTRODUCTION	11
3	1.1 CLINICAL CONTEXT	11
4	1.2 LINKS AND SUBJECTS	11
5	1.3 ARCHITECTURE SUMMARY	13
6	1.4 READING THIS DOCUMENT	14
7	2 SCOPE AND OBJECTIVES.....	15
8	2.1 SPECIFICATION ORGANIZATION	15
9	2.2 ASSUMPTIONS/ASSERTIONS	16
10	2.3 CMA DESIGN CENTER	18
11	3 TECHNOLOGY NEUTRALITY	19
12	4 REQUIREMENTS AND CAPABILITIES.....	23
13	5 SYSTEM ARCHITECTURE.....	25
14	5.1 USE-MODEL	25
15	5.2 CONTEXT MANAGEMENT RESPONSIBILITY	34
16	5.3 CONTEXT CHANGE DETECTION	35
17	5.4 CONTEXT DATA REPRESENTATION.....	35
18	5.5 CONTEXT DATA ACCESS	36
19	5.6 CONTEXT DATA INTERPRETATION	37
20	5.6.1 Establishing the Meaning of Context Data Item Names	38
21	5.6.2 Establishing the Meaning of Context Data Item Values	39
22	5.6.3 Representing Context Subjects That Cannot Be Uniquely Identified	39
23	5.6.4 Context Subjects and Item Name Format.....	40
24	5.6.5 Standard Context Data Items.....	40
25	5.6.6 Non-Standard Context Data Items	41
26	5.6.7 Representing “Null” Item Values	42
27	5.6.8 Representing an Empty Context Subject	42
28	5.6.9 Case Sensitivity with Regard to Item Names and Item Values	42
29	6 COMPONENT MODEL	45
30	6.1 COMPONENT AND INTERFACE CONCEPTS.....	46
31	6.1.1 Interfaces and References.....	46
32	6.1.2 Interface Interrogation.....	46
33	6.1.3 Principal Interface	47
34	6.1.4 Interface Reference Registry	47
35	6.1.5 Interface Reference Management	47
36	7 PATIENT LINK THEORY OF OPERATION	49
37	7.1 PATIENT LINK COMPONENT ARCHITECTURE.....	49
38	7.2 PATIENT SUBJECT	50
39	7.3 PATIENT MAPPING AGENT.....	51
40	7.4 CONTEXT CHANGE TRANSACTIONS.....	51
41	7.5 JOINING THE COMMON CONTEXT SYSTEM.....	52
42	7.6 CONTEXT CHANGE TRANSACTIONS.....	53
43	7.7 TRANSACTIONAL CONSISTENCY.....	53
44	7.8 CONTEXT CHANGE NOTIFICATION PROCESS	54
45	7.9 LEAVING A COMMON CONTEXT SYSTEM	56

1	7.10	BEHAVIORAL DETAILS.....	56
2	7.10.1	<i>Application Behavior When it Cannot Cancel Context Changes</i>	56
3	7.10.2	<i>Application Behavior When it Does Not Understand Context Identifiers</i>	57
4	7.10.3	<i>Application Behavior with Regard to an Empty Context</i>	57
5	7.10.4	<i>Surveying Details</i>	57
6	7.11	COMMON CLINICAL CONTEXT USE MODEL	59
7	7.11.1	<i>Lifecycle of Common Context</i>	60
8	7.11.2	<i>Context Selection Change Use Case</i>	63
9	7.11.3	<i>Abnormal Termination of Common Context Use Case</i>	72
10	7.12	STAT ADMISSIONS.....	74
11	7.13	OPTIMIZATIONS	74
12	7.14	THE SIMPLEST APPLICATION.....	75
13	8	MAPPING AGENTS.....	77
14	8.1	ASSUMPTIONS AND ASSERTIONS	77
15	8.2	INTERFACES.....	78
16	8.3	THEORY OF OPERATION.....	79
17	8.3.1	<i>Initializing a Context System When a Mapping Agent is Present</i>	80
18	8.3.2	<i>Terminating a Context System When a Mapping Agent is Present</i>	81
19	8.3.3	<i>Distinguishing Between Mapping Agents and Context Participants</i>	82
20	8.3.4	<i>Mapping Agent Updates to Context Data</i>	83
21	8.3.5	<i>Conditions for Mapping Agent Invalidation of Context Changes</i>	83
22	8.3.6	<i>Treatment of Mapping Agent Invalidation of Context Changes</i>	85
23	8.3.7	<i>Mapping Null-Valued Identifiers</i>	86
24	8.3.8	<i>Initializing Mapping Agents</i>	87
25	8.3.9	<i>Handling Mapping Agent Failures</i>	88
26	8.4	MAPPING AGENT EFFECT ON APPLICATION SECURITY POLICIES	88
27	8.5	IDENTIFYING MAPPING AGENT IMPLEMENTATIONS.....	89
28	8.6	PERFORMANCE COSTS AND OPTIMIZATIONS	89
29	9	USER LINK THEORY OF OPERATION.....	91
30	9.1	USER LINK TERMS.....	92
31	9.2	DESKTOP ASSUMPTIONS	92
32	9.3	USER SUBJECT	92
33	9.4	USER AUTHENTICATION DATA IS NOT PART OF THE USER CONTEXT	93
34	9.5	USER LINK COMMON CONTEXT SYSTEM DESCRIPTION	94
35	9.5.1	<i>User Mapping Agent</i>	94
36	9.5.2	<i>Context Management Interfaces</i>	95
37	9.5.3	<i>Authentication Repository</i>	95
38	9.5.4	<i>Overall User Link Component Architecture</i>	96
39	9.6	USER LINK SIGN-ON PROCESS	97
40	9.7	DESIGNATING APPLICATIONS FOR USER AUTHENTICATION	97
41	9.8	SIGNING ON TO APPLICATIONS NOT DESIGNATED FOR AUTHENTICATING USERS	98
42	9.9	APPLICATION BEHAVIOR WHEN LAUNCHED	99
43	9.10	MULTIPLE CONTEXT SUBJECTS.....	99
44	9.10.1	<i>The Effect of Multiple Subjects on the Meaning of "Link"</i>	99
45	9.10.2	<i>Context Manager Support for Multiple Context Subjects</i>	100
46	9.10.3	<i>Effect of Multiple Subjects on Context Change Transaction</i>	101
47	9.10.4	<i>Context Manager Treatment of Multi-Subject Context Data</i>	102
48	9.10.5	<i>Effect of Multiple Subjects on Mapping Agents</i>	102
49	9.10.6	<i>Application Treatment of Multiple Subjects</i>	103
50	9.11	ACCESS CONTROL LISTS.....	103
51	9.12	EMPTY CONTEXTS	103
52	9.13	CHANGING USERS	103

1	9.14	LOGGING-OFF AND APPLICATION TERMINATION	104
2	9.15	AUTOMATIC LOG-OFF.....	107
3	9.16	REAUTHENTICATION TIME-OUT	107
4	9.17	BUSY APPLICATIONS	108
5	9.18	CO-EXISTENCE WITH APPLICATIONS NOT USER LINK-ENABLED.....	109
6	9.19	POPULATING THE USER MAPPING AGENT	109
7	9.20	AUTHENTICATION REPOSITORY	110
8	9.20.1	<i>Repository Implementation Considerations</i>	<i>111</i>
9	9.20.2	<i>Populating the Repository</i>	<i>111</i>
10	10	CHAIN OF TRUST	113
11	10.1	USER CONTEXT CHANGE TRANSACTIONS AND THE CHAIN OF TRUST	113
12	10.2	CREATING THE CHAIN OF TRUST.....	113
13	10.2.1	<i>Object Infrastructures</i>	<i>114</i>
14	10.2.2	<i>Secure Communications Protocols</i>	<i>114</i>
15	10.2.3	<i>Security Building Blocks</i>	<i>115</i>
16	10.2.4	<i>Security Attacks On the Chain Of Trust.....</i>	<i>117</i>
17	10.2.5	<i>Chain of Trust Implementation Limitations.....</i>	<i>119</i>
18	10.3	DIGITAL SIGNATURES AND CMA COMPONENTS	120
19	10.3.1	<i>Public Key / Private Key Encryption as a Means for Generating Signatures</i>	<i>120</i>
20	10.3.2	<i>Incorporation of Signatures into the Context Management Architecture.....</i>	<i>122</i>
21	10.3.3	<i>Computing a Digital Signature.....</i>	<i>124</i>
22	10.3.4	<i>Public Key Distribution.....</i>	<i>125</i>
23	10.3.4.1	<i>Passcode Generation Requirements</i>	<i>127</i>
24	10.3.4.2	<i>Protecting Passcodes</i>	<i>128</i>
25	10.3.4.3	<i>Protecting Private Keys</i>	<i>129</i>
26	10.3.5	<i>System Configuration Requirements</i>	<i>129</i>
27	10.3.6	<i>Defending Against Replay Attacks.....</i>	<i>130</i>
28	10.4	TRUST RELATIONSHIPS	131
29	10.4.1	<i>Trust Between Applications and Context Manager</i>	<i>131</i>
30	10.4.2	<i>Trust Between Context Manager and User Mapping Agent.....</i>	<i>131</i>
31	10.4.3	<i>Trust Between Applications and Authentication Repository.....</i>	<i>132</i>
32	10.5	CHAIN OF TRUST INTERACTIONS	133
33	11	INTERFACE DEFINITIONS.....	137
34	11.1	INTERFACE DEFINITION LANGUAGE	137
35	11.1.1	<i>Interface Definition Body.....</i>	<i>138</i>
36	11.1.2	<i>Simple Data Types.....</i>	<i>139</i>
37	11.1.3	<i>Exception Declaration.....</i>	<i>140</i>
38	11.1.4	<i>Sequences.....</i>	<i>140</i>
39	11.1.5	<i>Interface References.....</i>	<i>141</i>
40	11.1.6	<i>Principal Interface</i>	<i>141</i>
41	11.1.7	<i>Qualifying Names.....</i>	<i>141</i>
42	11.2	INTERFACE IMPLEMENTATION ISSUES.....	142
43	11.2.1	<i>NotImplemented Exception.....</i>	<i>142</i>
44	11.2.2	<i>GeneralFailure Exception</i>	<i>142</i>
45	11.2.3	<i>Coupon Representation</i>	<i>142</i>
46	11.2.4	<i>Format for Application Names</i>	<i>142</i>
47	11.2.5	<i>Extraneous Context Items.....</i>	<i>143</i>
48	11.2.6	<i>Forcing the Termination of a Context Change Transaction</i>	<i>143</i>
49	11.2.7	<i>Character-Encoded Binary Data.....</i>	<i>145</i>
50	11.2.8	<i>Representing Message Authentication Codes, Signatures and Public Keys</i>	<i>146</i>
51	11.2.9	<i>Representing Basic Data Types as Strings.....</i>	<i>146</i>
52	11.2.10	<i>Pre-Defined Mapping Agent Coupons</i>	<i>147</i>

1	11.3	INTERFACES.....	149
2	11.3.1	<i>AuthenticationRepository (AR)</i>	149
3	11.3.1.1	Connect.....	149
4	11.3.1.2	Disconnect	150
5	11.3.1.3	SetAuthenticationData.....	150
6	11.3.1.4	DeleteAuthenticationData.....	151
7	11.3.1.5	GetAuthenticationData	152
8	11.3.2	<i>ContextData (CD)</i>	154
9	11.3.2.1	GetItemNames	154
10	11.3.2.2	DeleteItems	155
11	11.3.2.3	SetItemValues	156
12	11.3.2.4	GetItemValues	157
13	11.3.3	<i>ContextManager (CM)</i>	159
14	11.3.3.1	MostRecentContextCoupon	160
15	11.3.3.2	JoinCommonContext	160
16	11.3.3.3	LeaveCommonContext	161
17	11.3.3.4	StartContextChanges	161
18	11.3.3.5	EndContextChanges	162
19	11.3.3.6	UndoContextChanges.....	163
20	11.3.3.7	PublishChangesDecision	164
21	11.3.3.8	SuspendParticipation.....	164
22	11.3.3.9	ResumeParticipation.....	165
23	11.3.4	<i>ContextParticipant (CP)</i>	167
24	11.3.4.1	ContextChangesPending.....	167
25	11.3.4.2	ContextChangesAccepted	168
26	11.3.4.3	ContextChangesCanceled	168
27	11.3.4.4	CommonContextTerminated.....	169
28	11.3.4.5	Ping	169
29	11.3.5	<i>ImplementationInformation (II)</i>	170
30	11.3.5.1	ComponentName.....	170
31	11.3.5.2	RevMajorNum	170
32	11.3.5.3	RevMinorNum	170
33	11.3.5.4	PartNumber.....	170
34	11.3.5.5	Manufacturer.....	170
35	11.3.5.6	TargetOS.....	170
36	11.3.5.7	TargetOsRev	171
37	11.3.5.8	WhenInstalled	171
38	11.3.6	<i>MappingAgent (MA)</i>	172
39	11.3.6.1	ContextChangesPending.....	172
40	11.3.6.2	Ping	172
41	11.3.7	<i>SecureBinding (SB)</i>	174
42	11.3.7.1	InitiateBinding	174
43	11.3.7.2	FinalizeBinding.....	176
44	11.3.8	<i>SecureContextData (SD)</i>	178
45	11.3.8.1	GetItemNames	178
46	11.3.8.2	SetItemValues	178
47	11.3.8.3	GetItemValues	179
48	12	BACKWARDS COMPATIBILITY.....	181
49		APPENDIX: DIAGRAMMING CONVENTIONS	183
50		GLOSSARY	189
51			

1 **Figures**

2	Figure 1: Patient Linked Applications	12
3	Figure 2: Organization of HL7 Context Management Specification Documents	15
4	Figure 3: Overall Role of the CMA Specification.....	18
5	Figure 4: COM/Java/CORBA Interoperability	21
6	Figure 5: Patient Selection Change Use Case.....	28
7	Figure 6: Patient Context Automatically Changes within all Context Participant Applications	29
8	Figure 7: User Informed of Potential Data Loss and Cancels Context Change.....	30
9	Figure 8: User forces Application AAA to Become Out of Synchrony with other Context Participants	31
10	Figure 9: Context Participant Not Responding to Selection Change Request.....	32
11	Figure 10: User Accepts Consequences of going ahead with Patient Selection Change with all	
12	Applications	33
13	Figure 11: Patient Link Component Architecture.....	50
14	Figure 12: Patient Link Context Change Process	52
15	Figure 13: Common Clinical Context Use Model	59
16	Figure 14: Common Context Lifecycle Use Case	60
17	Figure 15: Context Selection Change Use Case	63
18	Figure 16: Abnormal Termination of Common Context Use Case	72
19	Figure 17: User Link Component Architecture	96
20	Figure 18: User Link Sign-On Process.....	97
21	Figure 19: User Subject Context Data Mapped for Different Applications.....	109
22	Figure 20: Signing A Message.....	122
23	Figure 21: Forming Signature Using Method Parameters	123

24

25 **Tables**

26	Table 1: User Link-Enabled Application Behavior for Termination and Log-Off	105
27	Table 2: Chain of Trust Attacks and Defenses	118
28	Table 3: Handling Transaction Instigator Failure	144
29	Table 4: Character Representations for Basic Data Types	147

1	Interaction Diagrams	
2	Interaction Diagram 1: Common Context Lifecycle.....	61
3	Interaction Diagram 2: Suspending/Resuming Context Participation.....	62
4	Interaction Diagram 3: All applications accept the changes	64
5	Interaction Diagram 4: An application conditionally accepts the changes; user decides to cancel	
6	changes	65
7	Interaction Diagram 5: An application does not respond to survey.....	66
8	Interaction Diagram 6: An application does not respond to change notification	67
9	Interaction Diagram 7: An application responds after context change transaction has completed.....	68
10	Interaction Diagram 8: A non-surveyed application participates in context change	69
11	Interaction Diagram 9: An application conditionally accepts the changes; user decides to accept	
12	consequences of change	70
13	Interaction Diagram 10: An application conditionally accepts the changes; user breaks link with	
14	common context	71
15	Interaction Diagram 11: Abnormal Termination of Common Context	73
16	Interaction Diagram 12: Simplest Application.....	76
17	Interaction Diagram 13: Context Change Transaction with Mapping Agent	82
18	Interaction Diagram 14: Mapping Agent <i>Invalidates</i> Context Change Transaction.....	87
19	Interaction Diagram 15: User Logs Off From One Application.....	106
20	Interaction Diagram 16: User Logs-Off From Desktop.....	106
21	Interaction Diagram 17: Populating Authentication Repository with User Authentication Data	134
22	Interaction Diagram 18: User Link Context Change Transaction.....	135
23		
24		

1

Preface

2

3

4

5

6

This document was prepared by Robert Seliger, Sentillion, Inc., on behalf of Health Level Seven's Special Interest Group for Visual Integration (formerly the Clinical Context Object Workgroup --- CCOW). Comments about the organization or wording of the document should be directed to the author (robs@sentillion.com). Comments about technical content should be directed to ccow@list.mc.duke.edu.

1 Introduction

This document specifies the Health Level Seven Context Management Architecture (CMA). This architecture enables multiple applications to be automatically coordinated and synchronized in clinically meaningful ways at the point-of-use. The architecture specified in this document establishes the basis for bringing interoperability among healthcare applications to the point-of-use, such as the clinical desktop.

1.1 Clinical Context

Clinical context is state information that a user establishes and modifies while interacting with healthcare applications at the point-of-use (e.g., a clinical desktop). The context is common because it establishes parameters that should uniformly affect the behavior or operation of multiple healthcare applications. The context needs to be managed so that the user has a way of controlling it, and so that applications have a way of robustly coordinating their behavior as the context changes.

Examples of clinical context includes but are not limited to:

- The identity of a patient whose data the user wants to view or update via the applications.
- The identity of the user who wants to access the applications.
- A moment in time around which temporal data displays should be centered by the applications.
- A particular patient encounter that the user wants to review via the applications.

Healthcare application developers often implement a common clinical context capability for their own applications. However, there are currently no standards that enable independently-developed applications to share a common clinical context. Further, with the diversity of application programming technologies currently available, a common context solution should strive to be applicable to at least several of the dominant and emerging technologies.

1.2 Links and Subjects

The approach taken for the CMA is to define the architecture that enables applications to establish a single link based upon a set of clinical subjects of common interest. The applications automatically and cooperatively change their state whenever the user sets a new

value for one or more of these subjects. Two link subjects are defined as core to the CMA, and are therefore introduced in this document:

- Patient, which enables the user to select the patient of interest once from any application as the means to automatically “tune” all of the applications to the selected patient.
- User, which enables the user to securely logon once to any application as the means to automatically “tune” all of the applications to the user.

Applications that share the same common context are said to comprise a *common context system*. These applications have established and maintain a common context link. There is only one link, while there can be multiple subjects. However, in the vernacular that arose as the CMA was being developed, it became useful to refer specific link subjects. This has given rise to the terms such as *Patient Link* and *User Link*. An example of a set of Patient Linked applications is shown in Figure 1.

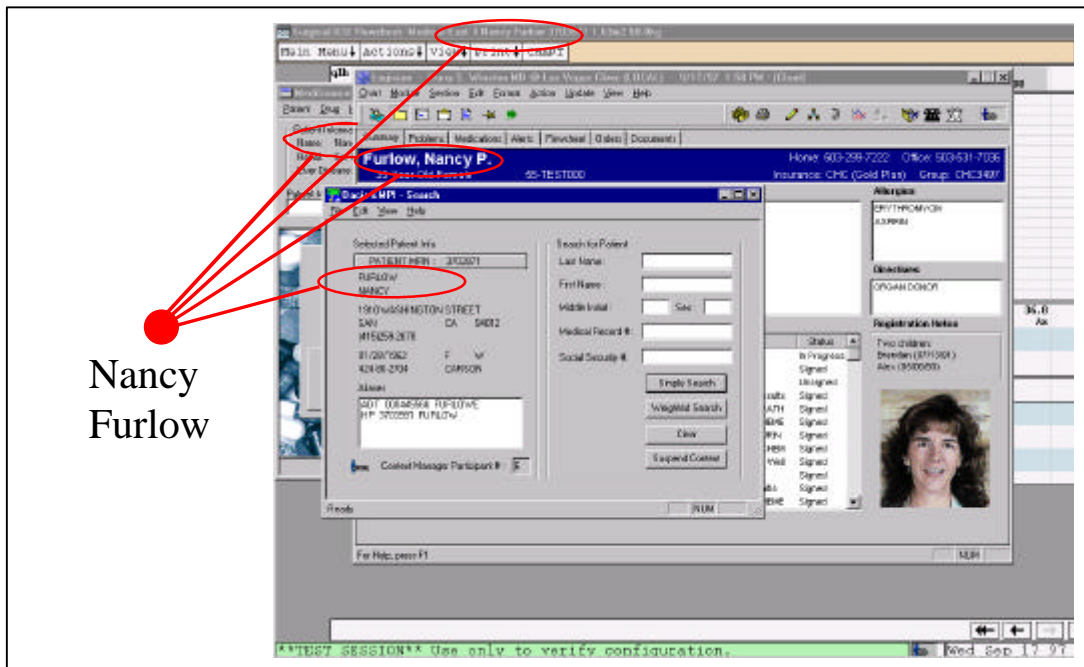


Figure 1: Patient Linked Applications

The architecture for Patient Link was developed prior to the extensions defined for User Link. In particular, User Link introduced substantial additional security-related capabilities. This specification presents a single consolidated view of the overall CMA.

The CMA enables additional subjects to be defined in a manner that does not require changes to the architecture. This capability is the basis for extensible standards-based context

management solutions that can evolve to address new requirements without requiring massive architecture or application implementation changes.

1.3 *Architecture Summary*

The CMA defines the interfaces between applications, known as context participants, and a coordinating component, known as the context manager. The CMA also defines the policies that govern the use of these interfaces and the interaction among and between CMA-compliant applications and components.

Applications that share a common context with each other, and the context manager that mediates the applications, are collectively referred to as a common context system.

Applications only need to interact with the context manager in order to participate in a common context system.

The data that defines the common clinical context for a common context system resides in the context manager. The data is organized as a set of name/value pairs that are grouped by context subject (e.g., patient, user, etc.).

When the user performs an application gesture that instructs the application to change the common clinical context (e.g., the user has selected a different patient), the application starts a context change transaction. Context items can be added or removed, or have their values changed, during a context change transaction. Only one transaction can be in progress at a time.

When the application that instigated the transaction has completed its changes to the context data, the context manager conducts a two-phase process to coordinate the propagation of the context changes to the other applications.

In the first phase, the context manager surveys the other applications to determine which ones can apply the new context, and which ones either cannot, or prefer not to. An application cannot apply the changes if it is blocked, for example if it is waiting for the user to enter data. An application might prefer not to apply the new context if, for example, doing so might cause the user to lose work-in-progress.

The context manager informs the instigating application of the survey results. If all of the applications are willing to apply the new context, then they are all instructed to do so. If at least one of the surveyed applications is blocked (“busy”) or prefers to keep the previous context, then the user is asked by the instigating application to decide how to proceed:

- The user can cancel the context change.

- The user can break the link between the instigating application and the other applications. The new context is then applied only to the instigating application, while the other applications remain linked together and tuned to the previous context.
- The user can apply the changes anyway (as long as there are no busy applications).

The context manager broadcasts the decision to all of the context participants to complete the second phase of the transaction. This approach ensures that the link among application is never broken unless the user has performed an explicit gesture instructing that the link be broken.

Mapping agents are an optional CMA components that provide an automatic means for adding data to the common context. The additional data augments the context such that all of the participant applications can “tune” to the same subject even when they do not necessarily have a common way to identify the subject. The specific job of a mapping agent is to map the context data set by the application that instigated a context change transaction to data that can be interpreted by the other context participant applications. A mapping agent only interacts with the context manager, so its existence is transparent to the applications.

Finally, for situations in which the secure conveyance of a context change is required, the “chain of trust” is defined. In the chain of trust, the applications and components in a context system use digital signatures to identify themselves in a manner that can be readily authenticated but not easily violated. The chain of trust allows only trusted applications and components to interact within a common context system.

1.4 Reading This Document

This document presents a comprehensive specification of the HL7 Context Management Architecture. The precision of the specification becomes increasingly more detailed as the document progresses. Several of the early chapters present concepts that underlie the architecture and lead the reader through the rationale for various architectural choices, while all of the chapters in this document include information that the reader should find pertinent to the explanation of the CMA.

However, Chapters 5 through 11 all contain normative content and as such should be regarded as the core of the CMA specification. In particular, Chapter 11, Interface Definitions, concludes the core specification with the complete set of CMA interface definitions, including methods and their argument signatures. These interfaces are ultimately the basis for the implementation of applications and components that conform to the CMA specification.

A compliant CMA application or component shall implement the relevant set of CMA interfaces exactly as specified. A compliant application or component implementation shall adhere to these interface definitions and to the policies specified throughout this document that govern the use and behavior of these interfaces.

2 Scope and Objectives

The HL7 Context Management Architecture (CMA) enables independently developed applications to share data that describes a common clinical context. This document emphasizes the policies, protocols, software interfaces, and responsibilities applications must implement and adhere to as participants in a shared context system.

A common context system is comprised of applications launched directly or indirectly by a particular clinical end-user, wherein the applications share the same context data. Also included in this system is a context management facility that enables applications to share the context data.

2.1 Specification Organization

It is beyond the scope of this document to provide all of the details that are needed in order to fully implement a conformant CMA system. The necessary additional details are covered in a series of companion specification documents. As illustrated in Figure 2, these documents are organized to facilitate the process of defining additional link subjects and to accelerate the process of realizing the CMA using any one of a variety of technologies.

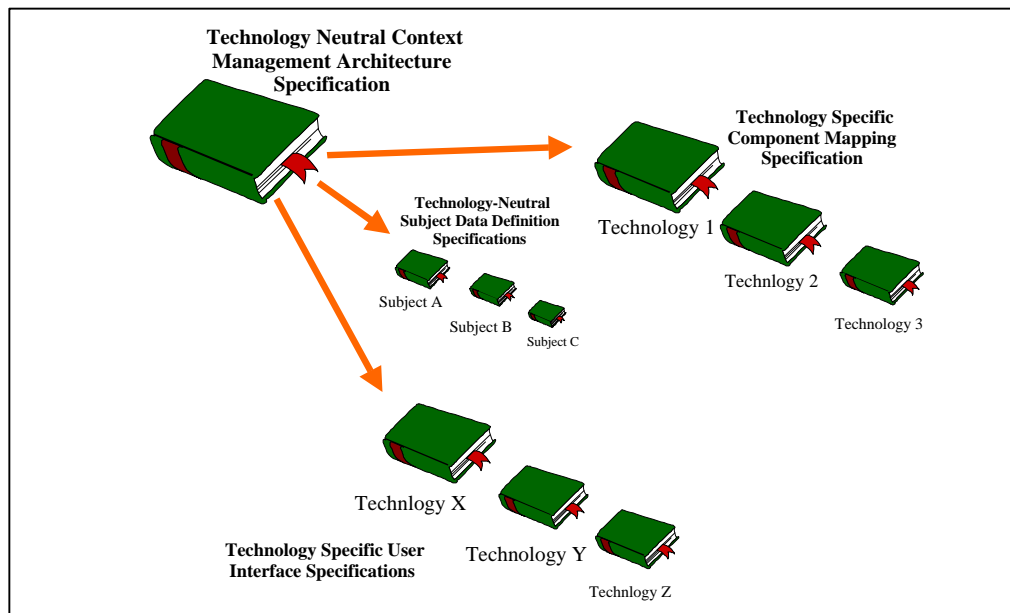


Figure 2: Organization of HL7 Context Management Specification Documents

The context management subjects and technologies that are of interest are determined by the HL7 constituency. There is an HL7 context management data definition specification

document for each of the standard link subjects. Each document defines the data elements that comprise a link subject. Concurrent with the publication of this document, the following documents have been developed:

Health Level-Seven Standard Context Management Specification,
Data Definition: Patient Subject, Version CM-1.0

Health Level-Seven Standard Context Management Specification,
Data Definition: User Subject, Version CM-1.0

There is an HL7 context management user interface specification document for each of the user interface technologies with which CMA-enabled applications can be implemented. Each document reflects the user interface requirements established in this document in terms of a technology-specific look-and-feel. Concurrent with the publication of this document, the following document has been developed:

Health Level-Seven Standard Context Management Specification,
User Interface: Microsoft Windows OS, Version CM-1.0

Finally, there is an HL7 context management component technology mapping specification document for each of the component technologies that can be used to implement the CMA. Each document provides the technology-specific details needed to implement CMA-compliant applications and the associated CMA components, as specified in this document. Concurrent with the publication of this document, the following document has been developed:

Health Level-Seven Standard Context Management Specification,
Component Technology Mapping: ActiveX, Version CM-1.0

2.2 Assumptions/Assertions

Key assertions and assumptions that were made during the course of developing the CMA are indicated below:

- The architecture does not intend to solve nor is it a substitute for solving the patient identification problem¹. However, the architecture does attempt to accommodate

¹ In general, patients cannot be reliably identified using their given name because given names are not necessarily unique. Identifiers can be assigned, but often a single person accumulates multiple patient identifiers over time. This is because the assigned identifiers are not universally unique, and generally only refer to a population of patients known to a particular healthcare institution, or known to a site within an institution. Government assigned identifiers, such as a social security number, may not be unique, or may change over time. In general, there is currently no simple and reliable way to identify the same patient across all possible systems that might contain data pertinent to the patient.

established means for achieving consistent interpretations of patient identification information.

- Architectural support for context data other than that which is used to identify patients is a non-objective to the extent it complicates the architecture. However, the architecture is currently applicable to a wide range of context data elements.
- Architectural support for distributed applications is a non-objective to the extent it complicates the architecture. However, the architecture is currently applicable to distributed as well as co-located applications.
- Context management is not a form of data interchange nor is it a substitute for data interchange. However, the common context might contain data that can also be obtained by an application through data interchange mechanisms such as those based upon HL7 (e.g., a patient's name or date of birth in addition to a patient identifier). When such data is provided, it is only as a means to simplify or optimize the sharing of common context.
- The context management facility is not visible to the clinical end-user. However, it might be visible to a systems integrator or systems administrator.
- The architecture is intended for use in clinical systems that are configured by an IT staff. Ad-hoc installation and configuration of a common context system by the clinical user is a non-objective to the extent it complicates the architecture.
- There is at most one context management facility per clinical desktop. However, applications shall work correctly with any facility implementation that conforms with the CMA specification. It is the decision of the IT staff as to which facility implementation is actually used by a clinical system.
- Implementation complexities will be shifted to the context management facility, as opposed to the applications, whenever this tactic is practical and reasonable. Minimizing the burden for the application developer is valued as an essential element for attracting the participation of the widest possible array of applications.
- It is assumed that the clinical data used by applications that share a common clinical context are appropriately synchronized (e.g., via back-end data interchange) to the degree necessary to ensure the consistent interpretation of the common context.
- It is assumed that any application that has been activated by the user can be used to set the user's common clinical context as long as the application conforms to the CMA specification. This enables multiple applications to provide context setting capabilities, which is convenient for the user.

- It is assumed that any application that does not understand or is otherwise unable or unwilling (e.g., for security reasons) to respond to a change in the common clinical context will ignore the change. However, any application that chooses to ignore a context change must clearly indicate its decision, for example by blanking its data display and/or minimizing itself.

2.3 CMA Design Center

The CMA specification is primarily aimed at enabling interoperability in the form of application control by the end user. Applications that interoperate in this manner appear to the user as visually integrated. This is because the user can see ways in which the applications interoperate.

This is in contrast to traditional healthcare standards, which have been primarily aimed at enabling interoperability in the form of data interchange between applications. Further, the design focus for the CMA specification is applications that have a means for interchanging clinical data. The overall role of the CMA specification is illustrated in Figure 3.

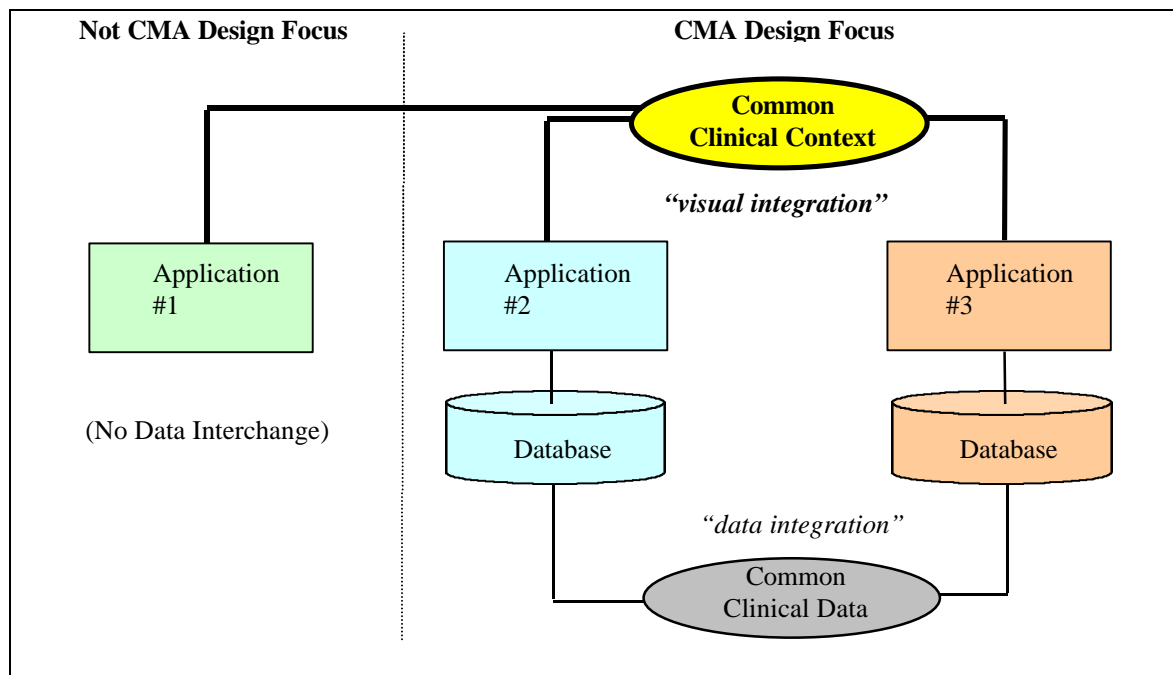


Figure 3: Overall Role of the CMA Specification

3 Technology Neutrality

As recently as one year ago, it would have sufficed to architect and implement a common clinical context solution that was targeted specifically for the Microsoft Window platforms. With the recent explosion of Web-based technologies, such as Java, this restriction is no longer practical. Fortunately, it is possible to architect a solution that is not predicated upon a specific technology. Specifically, in the architecture described in this document, the concept of technology neutrality is also applied.

The term “technology neutral” does not mean that any technology is applicable. Rather, it means that the common clinical context approach should work equally well with any one of a candidate set of relevant technologies.

The candidate technologies considered for this document are based upon market leadership:

- Inter-component communication: via Microsoft Automation through COM/DCOM; via any CORBA 2.0 compliant object request broker.
- Programming languages: any language that can be interfaced with Microsoft Automation and/or CORBA (e.g., VisualBasic®, C++, Java, MUMPS).
- Operating Systems: Windows 95®; Windows NT®; any platform that can host a Java virtual machine.

The primary reason that technology neutrality is practical is because all of these technologies have a lot in common, including:

- They are all based upon object-oriented principles.
- They are all embraced by Microsoft or are readily available on Microsoft platforms.

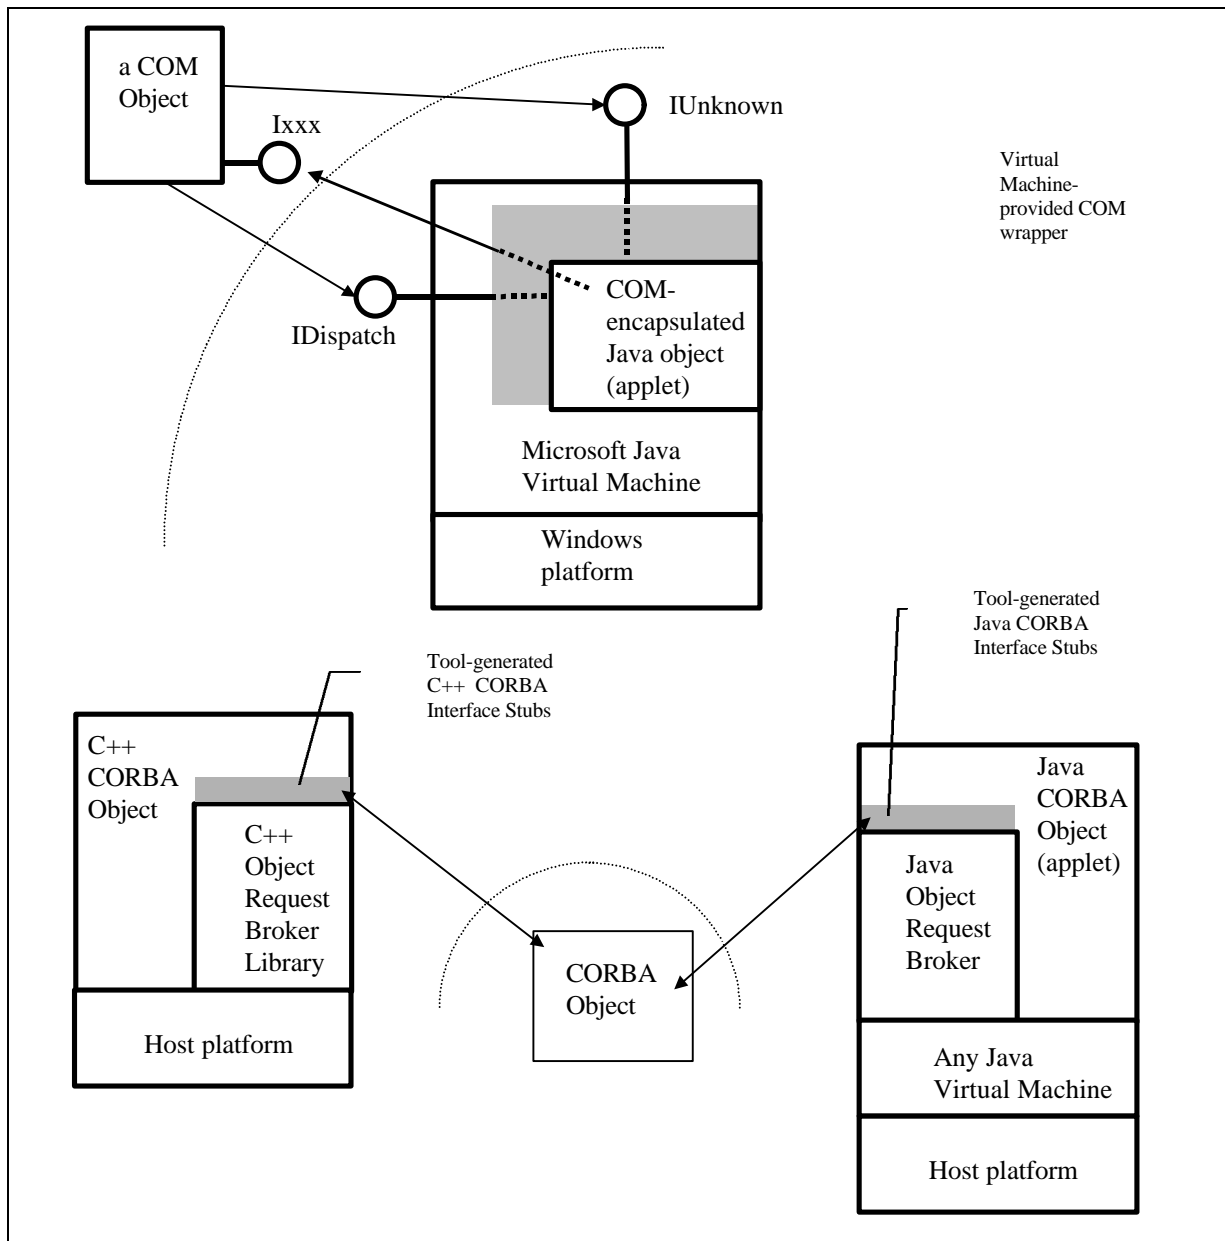
These two points have an interesting consequence: the technologies are compatible and interoperable. This makes it a lot easier to be technology neutral. For example:

- CORBA supports multiple programming languages. Support already exists for C, C++, Smalltalk, Java, and MUMPS. Objects implemented in any of these languages can transparently interoperate using CORBA.
- COM supports multiple programming languages. Support already exists for C++, VisualBasic, ObjectPascal, Java, and MUMPS. Objects implemented in any of these languages can transparently interoperate using COM.
- Most vendor’s CORBA object request brokers enable CORBA objects to transparently interoperate with COM objects.

- 1 • Microsoft's Java virtual machine enables Java objects (applets) to transparently
2 interoperate with COM objects.
- 3 • Java objects (applets) can transparently communicate with remote Java objects using
4 the Java Remote Method Invocation (RMI) mechanism.

5 Given the synergistic state of the dominant object technologies, the emphasis of this document
6 is on the structure of the common context system, the roles and responsibilities of the
7 components that comprise the system, the precise definition of the interfaces they need to
8 implement in order to be participants, the interactions between the components (via their
9 interfaces), and a host of architectural decisions that are intended to result in a robust,
10 practical, and useful common context solution.

11 Figure 4 illustrates a COM-encapsulated Java object that interoperates with other COM
12 objects, and C++ and Java CORBA objects that interoperates with other CORBA objects.



1 **Figure 4: COM/Java/CORBA Interoperability**

2

3

4 Requirements and Capabilities

The architecture described in this document is intended to serve as an extensible basis for future, more advanced, common clinical context capabilities. However, for now, an attempt will be made to focus on the immediate issue of developing a robust solution for sharing a common patient selection context.

In a complete solution, at least the following issues need to be addressed:

- Extensibility - how can new context elements be easily added in the future?
- Coordination - how can applications be coordinated so that they respond to context setting changes in an orchestrated and manageable manner?
- Flexibility - how can applications and common context managers be structured so that they implement only the capabilities that they need?
- Performance - how can applications and common context managers be structured so that their temporal performance and utilization of computing resources is acceptable to the end-user?
- Localizability - how are internationalization issues addressed (e.g., local character sets, etc.)?
- Scalability - how is the performance of a common context system affected by the quantity of active applications?
- Applicability - how should context information be structured and managed so that application behaviors are useful to the end user?
- Usability - what are the policies that govern the use of a common context such that the resulting application behaviors are intuitive and reasonable?
- Verifiability - how will the correctness of independently developed common context implementations be verified?

Architectural approaches that address these issues are presented next.

5 System Architecture

At the most abstract level, the Context Management Architecture (CMA) provides a way for independent applications to share data that describe a common clinical context. However, the CMA must provide solutions for the following problems:

- What is the general use model for a common context, from the user's perspective?
- Where does the responsibility for context management reside?
- How are changes to context data detected by applications?
- How is context data organized and represented so that it can be uniformly understood by applications?
- How is context data accessed by applications?
- How is the meaning of context data consistently interpreted by applications?

Before drilling into the details of the complete CMA, this chapter presents approaches and associated trade-offs for the problems listed above.

5.1 Use-Model

There are many possible use-models for a common clinical context.

The extremes of application support for making context changes are represented by:

- Context changes can be performed only via a single, distinguished, application.
- Context changes can be performed via any application.

In the model chosen for the CMA, context changes can be performed via any application. This is because it is not reasonable to assume the universal existence of a distinguished application, and it is beyond the interests and scope of HL7 to specify one.

The extremes of application behavior when context changes are made are represented by:

- When the user changes the context, the changes are automatically communicated to all of the applications that share the context. Applications that are able and willing to apply the context changes do so immediately. Applications that are unable or unwilling to apply the context changes maintain their current context. It is assumed that the user can easily determine which context an application is using.

- When the user changes the context, the changes are automatically communicated to all of the applications that share the context. However, the context changes are only allowed if all of the applications are able and willing to apply the context changes immediately.

The model developed for the CMA is a hybrid of these two extremes that attempts to enable a high degree of automatic context management while also emphasizing clinical safety:

- The likelihood that applications can become uncoordinated with regard to a common clinical context is minimized.
- The circumstances that can prevent context changes from being automatically applied are expected to be infrequent.

The CMA model also respects the challenges of retrofitting common context capabilities into existing healthcare applications. Only modest assumptions about the capabilities of these applications and technology used to develop them are presumed. The CMA model is as follows:

- All or part of the common context can be set by the user from any application for which providing this capability is functionally relevant.
- When the user changes the context, the change is automatically communicated to all of the applications that share the context. The applications are expected to apply the new context in a clinically meaningful manner. In general, applications are also expected to apply the context changes immediately. Exceptions are described below.
- An application may choose to defer applying a context change until some time in the future. For example, an application that retrieves large medical image files (that require substantial processing) might choose to not retrieve images each time a different patient is selected as part of the clinical context. Instead, the application might wait for an explicit directive or gesture from the user before actually retrieving the image. An application that behaves in this manner must be sure that it does not show data for an earlier context. Blanking-out its data displays or minimizing itself are possible ways that this can be accomplished.
- An application for which a change in the context might result in the loss of work performed by the user can request that the user explicitly decide whether to proceed with the context change anyway, or to cancel the change. The solicitation of user input is performed by the application that is being used to change the context. The solicitation includes an identification of the application for which work might be lost and a description of the work that might be lost. An application that behaves in this manner is expected to be able to discard its work in progress and apply the context changes if instructed to do so. For example, a medication ordering application might

1 indicate that the inputs for a medication order that has not yet been completed by the
2 user will be lost if the context is changed to a different patient.

- 3 • When an application is unable to respond to a context change, perhaps because the
4 user left it waiting for user input, the user is asked to explicitly decide about how to
5 proceed. The solicitation of user input is performed by the application that is being
6 used to change the context. The solicitation includes the identification of the non-
7 responsive application and indicates that the application cannot respond to a context
8 change. For patient safety reasons, when there are applications that cannot respond to
9 the changes, context changes will not be automatically applied to the applications that
10 share a common context.
- 11 • When it is not desirable or possible for context changes to be automatically applied,
12 either because there are applications for which work might be lost, or there are busy
13 applications that cannot be notified about context changes, the user can explicitly
14 interact with these applications to correct the situation, and then apply the context
15 changes. For example, the user might complete or terminate a dialog that was left open
16 in order to enable an application to apply the context changes.
- 17 • When it is not desirable or possible for context changes to be automatically applied,
18 the user can also decide to apply the context change only to the application that is
19 being used to change the context. The decision to do this is typically in response to an
20 interruption during which the user needs to momentarily divert her attention to a
21 different context for a specific application. The application is, in effect, disconnected
22 from the common context, and must clearly indicate this fact to the user in a visual
23 manner. The application can be subsequently instructed by the user to reconnect and
24 apply the common context. The common context may have changed between the time
25 the application was disconnected and the time it is reconnected to the common context.

26 A high-level summary of the interactions between applications when a clinical patient context
27 is changed is illustrated below. Figure 5 illustrates the use case actors (i.e. external forces)
28 involved in a context change such as a patient selection. (The actors are the user plus
29 applications, all of which are represented in the Jacobson modeling technique as stick figures.)
30 Figure 6 through Figure 10 illustrate some possible instances of the Patient Selection Change
31 Use Case from the user's perspective. Not all possible instances of this use case are provided.

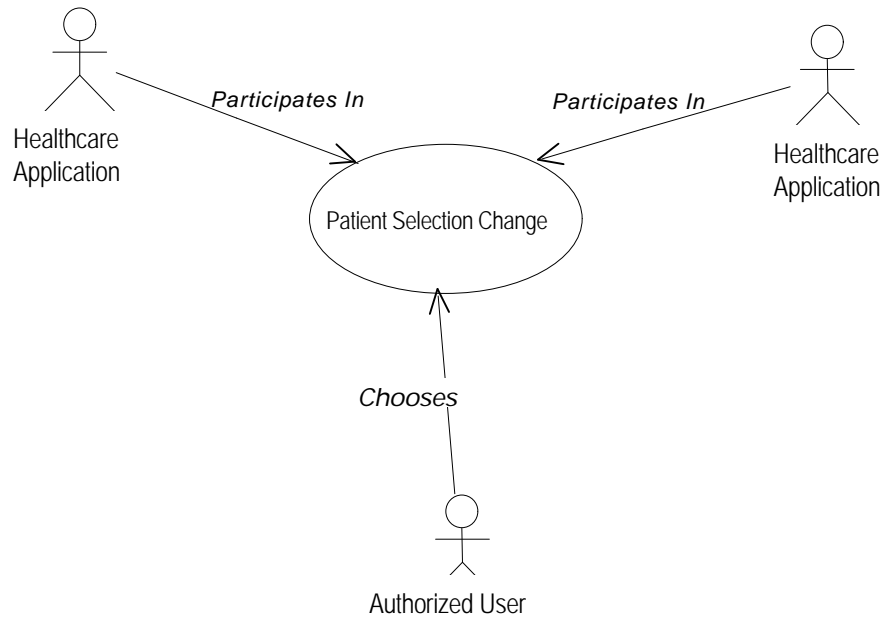


Figure 5: Patient Selection Change Use Case

The initial condition for each of the use case instances is that the currently selected patient is Jane Doe. In each instance, the user changes the common clinical context by selecting the patient Sam Smith. Some possible alternative outcomes follow:

- Figure 6 illustrates all applications reacting to the context change by changing their context to the patient “Sam Smith.”
- Figure 7 illustrates an application (Application DDD) conditionally accepting the context change and providing information describing work that could be lost if a context change occurs at this time. The user deciding to cancel the change is shown.
- Figure 8 illustrates a use case instance similar to Figure 7. However, the possible outcome of the user deciding to force a context change within Application AAA while the other applications remain with the original context is shown. This exemplifies Application AAA disconnecting from the common context system. Once disconnected, Application AAA’s context is no longer in synchrony with the other applications.
- Figure 9 illustrates healthcare application DDD not responding to a selection change request in a timely fashion. The user deciding to cancel the change is shown.
- Figure 10 illustrates the user being notified of potential data loss if selection change proceeds. The user accepting these consequences and proceeding with the change is shown.

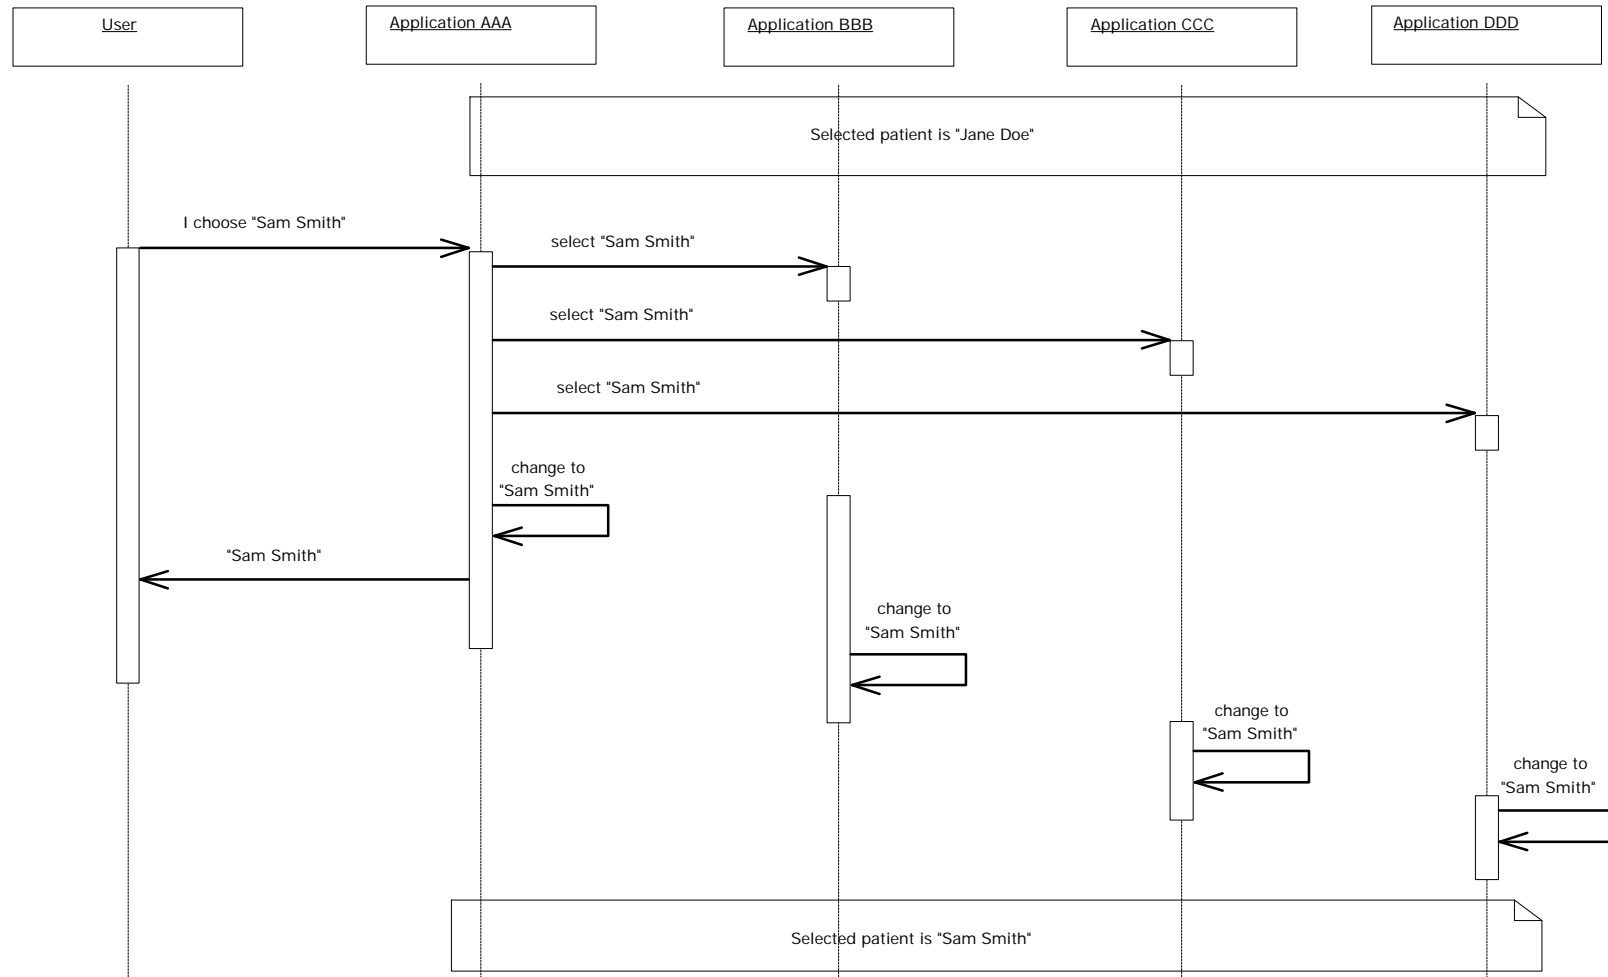
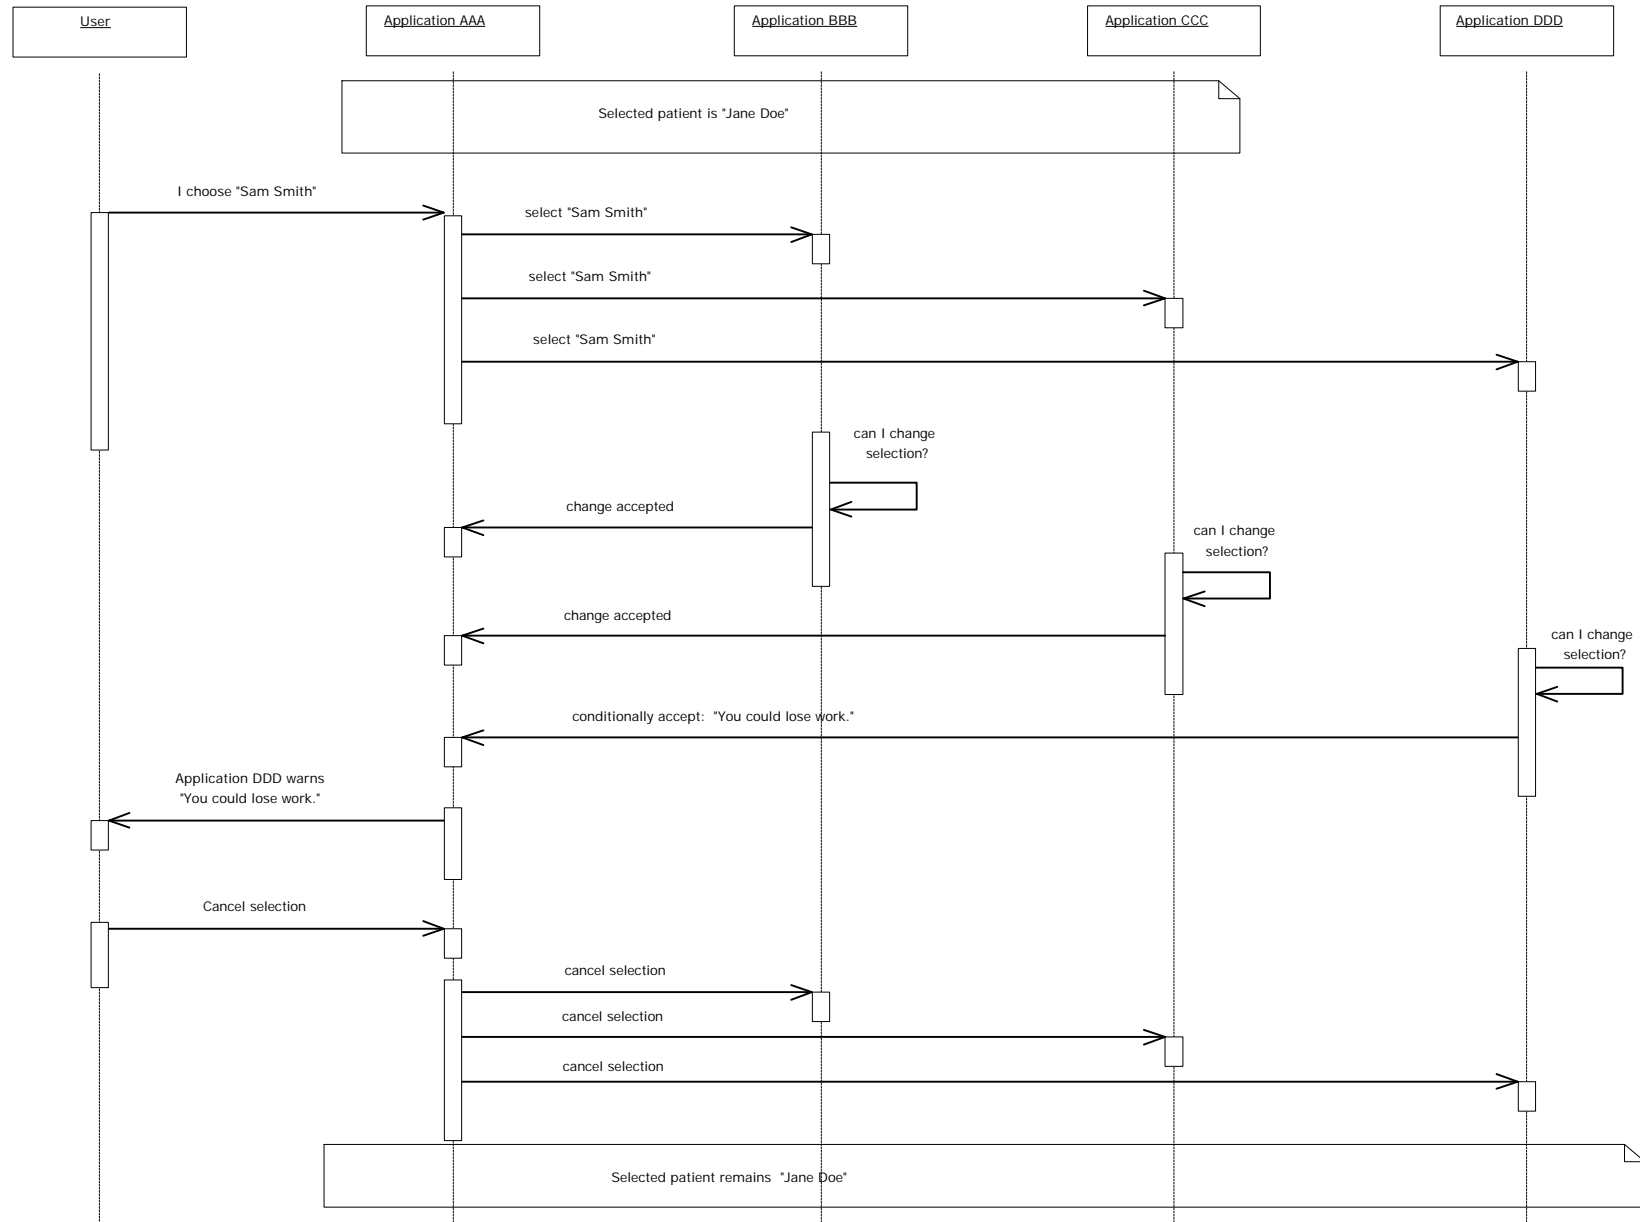


Figure 6: Patient Context Automatically Changes within all Context Participant Applications

Context Management Specification, Technology and Subject-Independent Component Architecture

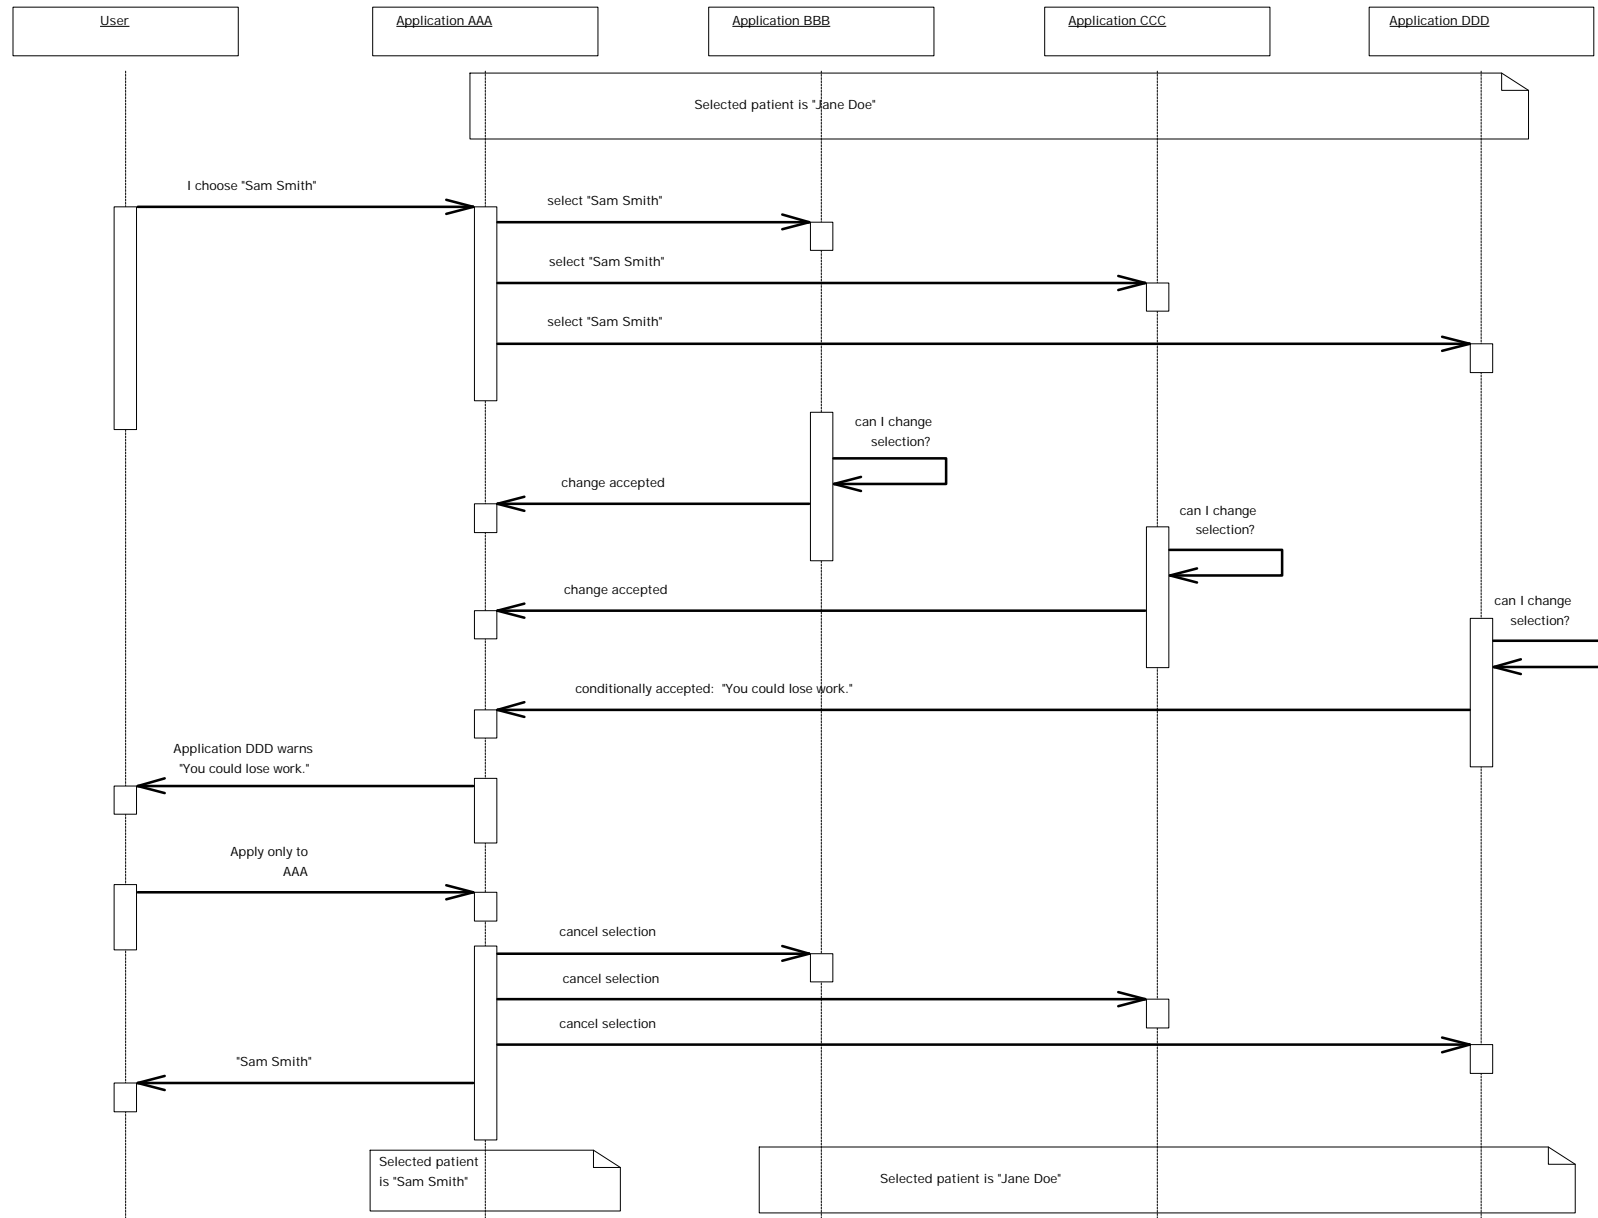


1

2

Figure 7: User Informed of Potential Data Loss and Cancels Context Change

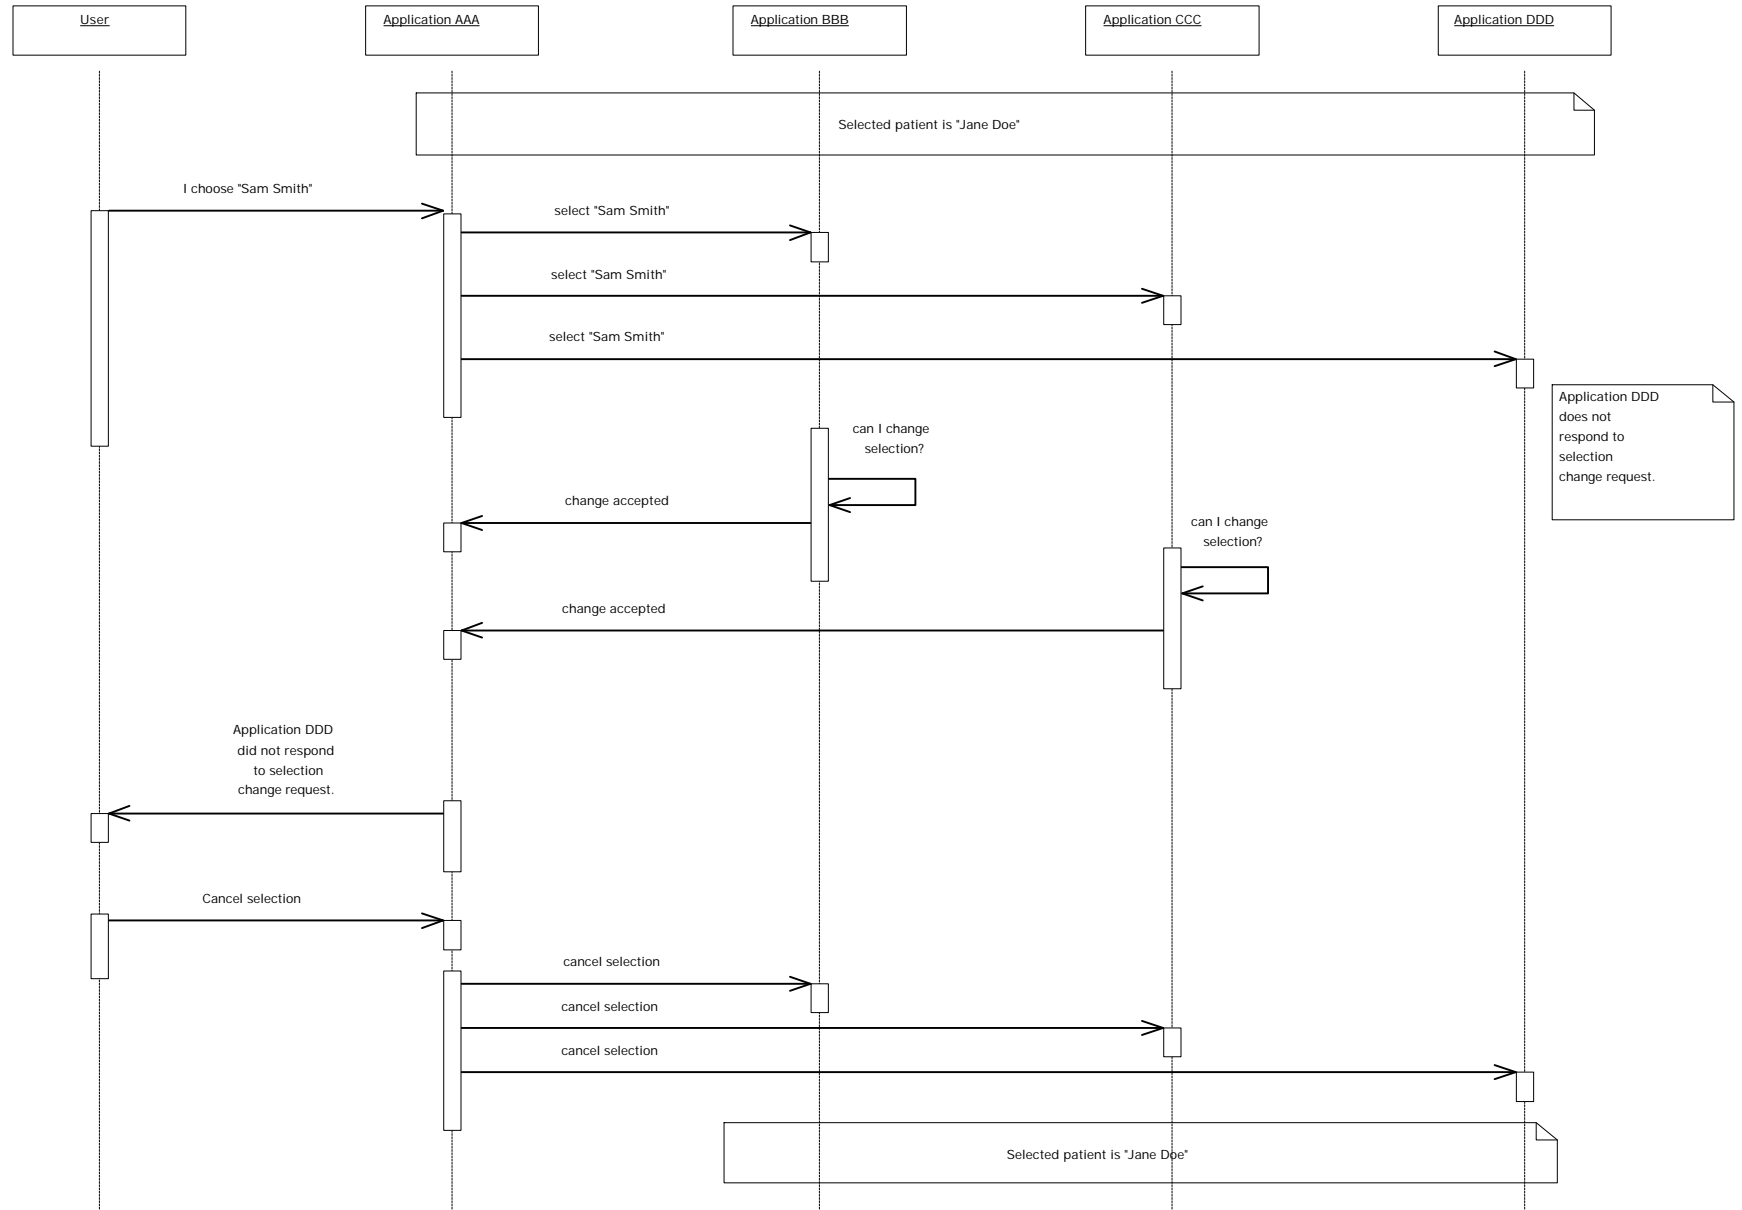
Context Management Specification, Technology and Subject-Independent Component Architecture



1

2 **Figure 8: User forces Application AAA to Become Out of Synchrony with other Context Participants**

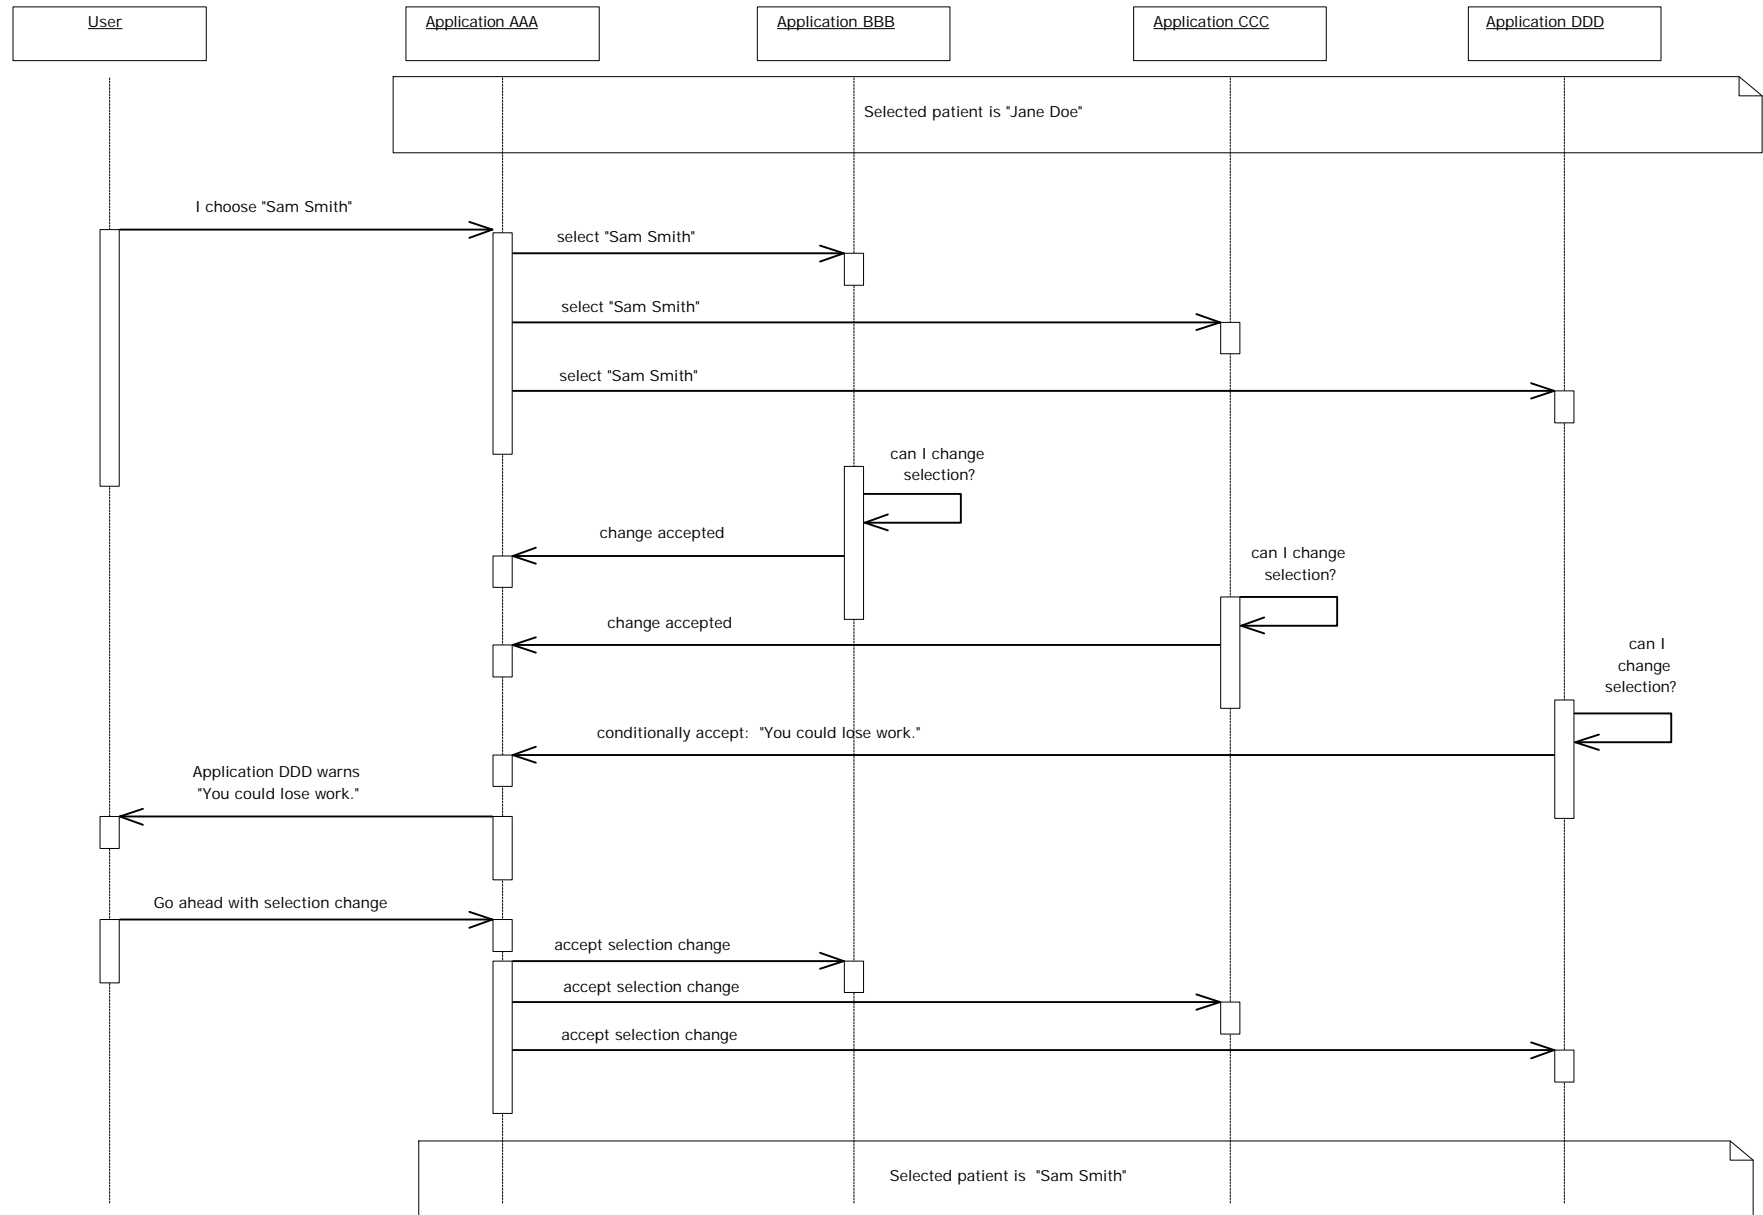
Context Management Specification, Technology and Subject-Independent Component Architecture



1

2

Figure 9: Context Participant Not Responding to Selection Change Request



1

2

Figure 10: User Accepts Consequences of going ahead with Patient Selection Change with all Applications

5.2 Context Management Responsibility

There are two fundamental schemes for architecting the responsibility for context management:

- **Distributed:** The responsibility for managing the common context is uniformly distributed among the applications. There is no central point of common context management.
- **Centralized:** The responsibility for managing the common context is centralized in a common facility that is responsible for coordinating the sharing of the context among the applications.

In the distributed model, applications must either all know about each other, or at least form a completely connected graph within which each application knows at least one other application. This is necessary in order for the applications to communicate context and control data among themselves.

Further, each application has the responsibility to act as a server for the common context in addition to acting as a client of the context. This is to offset the fact that there is no central point of ownership for the context, so each application must be capable of being an owner. This may be elegant, but it does introduce implementation complexities and burdens on all applications.

In the centralized model, applications only need to know about a common service or resource. This service off-loads from the applications much of the burden of maintaining and managing the common context. While a centralized service represents a single point of failure and a potential performance bottleneck, it is nevertheless the approach that is pursued in this document. The primary reasons include:

- It is simpler from the perspective of the application developer.
- The consequence of the service being a single point of failure is offset by the fact that the service and the applications it serves are typically co-resident on the same personal computer. Failures, if any, will be localized to a single user.
- The consequence of the service being a performance bottleneck is offset by the fact that the applications are far more likely to become the performance bottlenecks.

Given this basic system structure, the approaches for the other major architectural issues are summarized next.

5.3 Context Change Detection

There are at least two distinct categories of architectural approaches for realizing a common clinical context system:

- **Pull-model:** A shared component is used to maintain the shared context data. Applications update this resource to change the data. Other applications periodically poll the component to determine if the data has changed.
- **Push-model:** A shared component is used to maintain the shared context data. This component notifies applications whenever the data is changed. In order to receive a notification, an application must have first explicitly indicated its interest in being notified.

Both models have advantages and disadvantages. For example, the pull model is simpler to implement (e.g., does not require applications to handle asynchronous notifications), but can lead to performance problems due to polling even when the context data has not changed. Conversely, the push model can be the basis for better performance, but introduces additional implementation complexity.

Both models introduce the additional challenges of synchronizing concurrent access to the context data (e.g., to prevent two applications from attempting to change the data at the same time). In addition, both models must deal with failures modes that can occur when independent applications (i.e., applications that may be implemented as separate executables) are involved. For example, an application that crashes in the middle of changing the context data may leave the context data in an inconsistent state.

Given this analysis, the approach that is taken for the CMA is perhaps best described as a robust push-model. This is a push model that deals with synchronization and partial failure issues.

5.4 Context Data Representation

There are at least three distinct categories of architectural approaches for representing the common context data:

- **Fully-populated objects:** Objects are defined with properties and methods that model the real-world entities that they represent (e.g., a patient, a provider, etc.). These objects may be complex and involve a rich structure (e.g., are comprised of a logical network of objects).
- **Fully-populated messages:** Messages (as in “HL7 messages”) are used to convey detailed information about the context data.

- **Name-value pairs:** A set of name-value pairs represent only key summary information about the common context (e.g., just the patient's name and medical record number). The symbolic name for an item describes its meaning. The data types for the items come from a set of simple primitive data types.

The fully-populated object approach is perhaps the purest approach, but is subject to performance concerns. Copies of the objects could be produced and then communicated to each application every time the state of the primary copy changes. However, this involves the performance cost of marshaling the objects. The problem is further compounded by the fact that marshaling capability would need to be explicitly implemented in either CORBA or COM. (Java RMI implicitly supports the capability to communicate objects by value.)

The fully-populated message approach is actually a stylized way of marshaling objects. While it is appealing to think of leveraging existing healthcare standards such as HL7, it is non-trivial to implement the parsers and translators to create and interpret these messages. Even if such an implementation was commercially available, it is not clear that it would be desirable to require that all of the applications in a shared context system be able to support HL7 messages.

The name-value pair approach represents the compromise that is pursued in this document. Using simple primitive data types enables the values of the items to be easily communicated between processes. Performance concerns are mitigated because an application will be able to examine the values of only those items of interest in a single out-of-process access. (The application simply indicates the names of the items whose values it is interested in.) The approach is also readily extensible, as new items (i.e., new name-value pairs) can easily be added to the set of items.

All of the context data representation approaches described above are subject to establishing semantic agreement about the meaning of the data. This is true whether the context data is represented as objects, messages, or name-value pairs. The process for establishing this agreement is beyond the scope of the CMA, and is instead specified in a series of HL7 context management subject-specific data definition documents. These data definitions are key to implementing a plug-and-play common clinical context system.

5.5 Context Data Access

Any common context architecture must provide a way for an application that has just started to obtain its initial view of the common context. The pull-model implicitly solves this problem. With the push-model, there are two basic approaches:

- When the application joins the common context system, the necessary data is pushed to it.

- The data can be accessed from a well-known location, such as a file, or from the component that is responsible for pushing changes to the context system participants. This is, in effect, a specialized use of the pull-scheme.

The approach to this problem is linked to the approach by which applications access the context data for updating it, and the approach by which applications obtain the values for the context data when it has changed.

The options are straightforward:

- Each application maintains a copy of the context data. As changes occur, each application updates its local copy accordingly.
- A central “authentic” copy of the context data is maintained. Context data updates are directed by applications to this copy. Applications access this copy in order to inspect changes.

The approach in which each application maintains its own copy of the context data has an elegance to it. However, in the absence of an authentic copy, an application that has gotten out of synchrony with its peers may have a difficult time restoring its notion of the common context. Further, the communication costs of keeping all applications in synchrony can become significant, particularly as the complexity and size of the common context increases over time as additional common context items are defined.

The approach that is taken for the CMA is to maintain a single authentic copy of the common context for each common context system. Applications can choose to cache context data or they can simply access the authentic copy whenever they need to. Applications can also selectively read or write specific context data name-value pairs. Further, when the context changes, an application is only informed about the change and is not provided with the data that has changed. The application can selectively access this data when it needs to.

This approach was chosen as a balance between performance and complexity. Performance issues are addressed by enabling applications to have selective access to context data. Complexity issues are addressed by not forcing applications to maintain their own copy of the common context data.

5.6 Context Data Interpretation

In order for applications to apply common context data in a clinically consistent manner, they must interpret the meaning of the data in a uniform manner. With context items represented as name-value pairs, applications must be able to uniformly interpret both the meaning of the name and the value of a context item, or determine that it cannot correctly interpret the item.

Context data items logically represent two categories of information: data that *identifies* a real-world entity or concept (such as a specific patient or a specific encounter), and data that can be used to *corroborate* the identity data. Identity information is required in order to establish a common context between applications that involves a real-world entity or concept.

Corroborating data can be used by applications and/or users as a basis for checking further that the identified entity or concept is what was expected.

For example, a patient's name can rarely be used to uniquely identify a patient. Typically, a medical record number or similar identifier that is generally unique over some population of patients for one or more clinical systems is used. However, these identifiers are rarely meaningful to the user. Corroborating data might be comprised of the patient's name, sex, and data of birth. This data provides applications and/or the user with an additional means to check that the identified patient is the intended patient.

The clinical context is considered to have changed in a meaningful manner when identifier data is changed. Applications are notified of changes to the context when identifier data, and possibly corroboration data, are changed. Changes to corroboration data that are not accompanied by associated changes to identifier data are not meaningful and are rejected.

5.6.1 Establishing the Meaning of Context Data Item Names

Given this approach of organizing context data items into identity and corroborating data, there are two basic techniques for establishing the meaning of context item names:

- Apply a Context Management-specific information modeling process to identify and define candidate clinical context item names and meanings.
- Leverage names and their meaning as established by existing healthcare standards, such as the HL7 messaging standard.

The approach that is taken for the CMA is that existing HL7 messaging terms and their meaning will be used as the default source for clinical context item names. New item names and associated meanings will be created only when the HL7 messaging standard is not applicable. The standard set of clinical context data context item names are specified in separate HL7 context management data definition specification documents. Only the specified set of context data items shall be implemented by conformant systems.

The reason for this approach is that the value-added for HL7 context management is not in defining clinical content, but rather in enabling new forms of clinically-rooted desktop-based interoperability between independently-developed healthcare applications. There is little incentive to create new information models and develop new clinical concepts when there are existing concepts, such as those already specified for HL7 messaging, which can be leveraged.

5.6.2 Establishing the Meaning of Context Data Item Values

The abstract data types used to represent context data item values will also be leveraged from the HL7 messaging standard. These types may be represented as strings encoded using a simple subset of the HL7 character encoding rules. These types may also be mapped into convenient technology-specific data types. The actual clinical context data context item data types are specified in the HL7 context management data definition specification documents.

There are two basic approaches for establishing the meaning of context item values:

- Assume that each item has a value that can be globally interpreted by all of the applications that share a common clinical context.
- Provide multiple values for each item name such that each value represents that same real-world entity or concept. Each application can apply the value it understands.

In some cases, it is safe to assume that a context item's value can be globally interpreted by all applications. For example, if a patient's date of birth is defined to be a corroborating context data item, the value of this item has a single global interpretation.

5.6.3 Representing Context Subjects That Cannot Be Uniquely Identified

Unfortunately, it is not possible to assume that all context subjects, such as patients, can be identified using globally unique identifier values. For example, a patient cannot necessarily be globally identified using a single identifier, such as a medical record number.

However, in these cases, there may be multiple synonymous identifier values, each of which is pertinent to a subset of the applications that share a common context. For example, a hospital and its affiliated clinics may assign their own medical record numbers to the same patient population. Applications, such as master patient index systems, enable tracking and mapping between these values. The result is multiple distinct values that identify the same patient.

It is not the purview of the CMA to resolve global identification issues. It is within the scope of the CMA to at least recognize that multiple identifier values may be necessary. Therefore, the approach taken in this document is to support multiple identifier values for context items when necessary.

An item that can have multiple values is actually represented as multiple items that have a common name prefix but use a distinct name suffix. The prefix for an item, and the constraints on values for the suffix, is defined in the HL7 context management subject-specific data definition specification document within which the item is defined. The suffixes are configured into an application using an application-specific process when the application is installed at a site.

The values for such items are provided either by an application when it changes the clinical context, or by an external mapping agent. (See Chapter 8, Mapping Agent.)

Immediately following the item subject label is a short string that indicates the role of the item in terms of whether it represents identifier data or corroborating data. The string “id” shall indicate the role of identifier data. The string “co” shall indicate the role of corroborating data.

5.6.4 Context Subjects and Item Name Format

All context items are organized by subject. Each subject represents a real-world entity or concept that is identified as part of the overall common clinical context.

Standard subject labels are defined in the HL7 context management subject-specific data definition specification documents. The labels comprise the first part of each context data item name. Examples of possible subject labels are “Patient” and “User”. Item name elements are separated by a period (.). Words in multi-word item name elements are separated by an underscore (_).

The general format of a context data item name is:

Item_subject_label.role.item_name_prefix.optional_item_name_suffix

Examples of the name format for possible context data items is shown below. The name for the items that represent a patient’s medical record numbers (MRN) for both a hospital and its affiliated clinic (assuming that they use different medical record numbers):

“Patient.Id.MRN.St_Elsewhere_Hospital”

“Patient.Id.MRN.St_Elsewhere_Clinic”

The name for an item that represents a patient’s date of birth might be:

“Patient.co.date_of_birth”

The actual subject labels, item names, and rules for generating an item name suffix are specified in each the HL7 context management subject-specific data definition specification documents.

5.6.5 Standard Context Data Items

Each of the standard HL7 CMA subjects and associated context data items are defined in a corresponding HL7 context management subject-specific data definition specification document. This includes the two core subjects, patient, and user, and their respective context data items.

5.6.6 Non-Standard Context Data Items

Organizations, such as healthcare provider institutions and vendors, may define their own context data items. These items may be in addition to the standard items defined for the standard subjects. Non-standard subjects shall not be defined.

The names for such items shall only use the item role denoted by the string “zz”². This makes it simple to distinguish standard and non-standard context data items. However, it is not possible to indicate in an item’s name as to whether it is an identifier or corroborating data.

The item name prefix for a non-standard item can be the same as a standard item name prefix, although this approach is discouraged because it can be confusing. An organization should choose item name prefixes that are different from the standard item name prefixes.

Each such item shall always include an organizationally-defined suffix. This suffix shall denote the organization that defined the non-standard item. It can be the case that non-standard items defined by multiple different organizations will be part of the same system’s context data set. To prevent conflicts among data item names, an organization is encouraged to choose a suffix that is unlikely to be the same as a suffix defined by another organization.

The assignment, format, and content of this suffix is not currently managed or specified by HL7. In the future, HL7 will assign identifiers, per ISO/IEC 8824:1990(E) clause 28, that enable an organization to have a unique suffix or set of suffixes. Organizations that choose to use such an identifier will be guaranteed that their identifier is unique.

An example of a non-standard item representing the next of kin for a patient is:

“Patient.zz.next_of_kin.Galaxy_Medical_Systems”

Non-standard subjects shall not be defined using “zz” items. The use of subjects other than the standard subjects defined by HL7 is non-conformant.

Organizations that define and/or use non-standard items should do so with the understanding that applications that use these items may not easily interoperate with applications that do not use the items. However, the definition of non-standard context data items can be an expedient for implementing context management systems with specific, extended, capabilities. Nevertheless, organizations are encouraged to work with HL7 to define new standard context data items and subjects, and limit the use of non-standard items to interim solutions.

² The use of “zz” is motivated by the HL7 2.3 Data Interchange specification, in which Z segments represent non-standard message segments.

5.6.7 Representing “Null” Item Values

The value of a context identifier item or corroborating data item can be set to the distinguished value of *null* to indicate that the item does not have a valid value. This capability provides a means for an application to explicitly indicate it has not set a valid value for a particular context item. For example, setting the value of the identifier whose name is:

“Patient.Id.MRN.St_Elsewhere_Hospital”

to *null* indicates that the application has not set a valid value for this identifier.

The actual representation of *null* is technology-dependent and is specified in each of the HL7 context management technology mapping specification documents.

5.6.8 Representing an Empty Context Subject

A context subject is *empty* when a real-world entity or concept is not currently identified. For example, for the patient subject, this means that a patient is not currently identified.

An empty context subject is represented in either of two ways:

- There are no context identifier items.
- There are context identifier items, but the values for all of these items are null.

The initial state for all subjects in the context is that they do not contain any identifier items. See Section 7.6, Context Change Transactions. An application can explicitly establish an empty context. See Section 7.10.3, Application Behavior with Regard to an Empty Context.

5.6.9 Case Sensitivity with Regard to Item Names and Item Values

Context item names are case insensitive. This means that case is not to be used for the purposes of comparing names. Further, the case used to represent the same item name can be different for different applications, and the case used to represent a particular item’s name at one time need not necessarily be the same at a later time. For example, the item names:

“Patient.Id.MRN.St_Elsewhere_Hospital”

“patient.id.mrn.st_elsewhere_hospital”

“PATIENT.ID.MRN.ST_ELSEWHERE_HOSPITAL”

are all equivalent.

1 A context item whose value is represented as a character string is also case insensitive, unless
2 otherwise noted in the HL7 context management subject-specific data definition specification
3 document that defines the item.

4 However, for consistency with the situations in which item values are case sensitive, the case
5 used to represent the value for a particular item is preserved once the value has been set. The
6 casing for the item's value is maintained until a different value is subsequently established for
7 the item.

8 For example, the following flow of events is allowed:

- 9 1. An application sets the value of "Patient.Id.MRN.St_Elsewhere_Hospital" to
10 "RS779238XZW".
- 11 2. An application gets the value of "Patient.Id.MRN.St_Elsewhere_Hospital" as
12 "RS779238XZW".
- 13 3. An application sets the value of "Patient.Id.MRN.St_Elsewhere_Hospital" to
14 "AS119292RUH".
- 15 4. An application gets the value of "Patient.Id.MRN.St_Elsewhere_Hospital" as
16 "AS119292RUH".
- 17 5. An application sets the value of "Patient.Id.MRN.St_Elsewhere_Hospital" to
18 "rs779238xzw".
- 19 6. An application gets the value of "Patient.Id.MRN.St_Elsewhere_Hospital" as
20 "rs779238xzw".

21 The following flow of events is not allowed:

- 22 7. An application sets the value of "Patient.Id.MRN.St_Elsewhere_Hospital" to
23 "RS779238XZW".
- 24 8. An application gets the value of "Patient.Id.MRN.St_Elsewhere_Hospital" as
25 "rs779238xzw".

6 Component Model

The architecture for a common clinical context system is described in terms of components and the interfaces they must implement in order to be participants in the system. Only the components and interfaces that are germane to the establishment and maintenance of a common clinical context for a clinical desktop are described.

A role is described for each component, and the policies that govern the intended use of the interfaces are detailed. These policies can be thought of as the patterns of allowed interactions between components. Both normal and exceptional interactions are described.

The key components in a common clinical context system are: a clinical context manager, one or more context participant applications, and an optional mapping agent for each context subject.

The context manager coordinates the applications each time there is a context change. It is also the “owner” of the authentic context for the system. The context participant applications set and/or get the context from the context manager. They must follow the policies established later in this document in order to behave as proper context management “citizens.”

A mapping agent is a service component that from the perspective of an application is a transparent participant in a context change. A mapping agent’s primary role is to add additional subject-specific context identifier items to the context data. This is useful when a subject is known to the various context participant applications via multiple distinct identifiers, but only one or a few of these identifiers are known to the application that sets the context.

Additional context management components are also defined, but serve in supporting roles. All of the necessary components are detailed later in this document.

The context manager does not need to know about the functionality or specific features implemented by any of the applications. Conversely, all applications perceive the context manager through a uniform set of interfaces and capabilities. Further, the applications do not need to know about each other in order to participate in the same context system. Finally, a mapping agent is transparent to applications, as it interacts only with the context manager.

Applications and the context management components can all be independently implemented and still interoperate as long as they comply with the CMA specification. The CMA specification is in turn predicated upon an underlying component model, described next.

6.1 *Component and Interface Concepts*

The clinical context manager and the applications that participate in a common context system are modeled in the architecture as components. The component model that is used is a high-level hybrid of the component models defined by Microsoft for its Component Object Model (COM) and by the Object Management Group for its Object Management Architecture (OMA).

6.1.1 Interfaces and References

In the hybrid model, components have one or more formally-defined object-oriented interfaces. Each interface defines a semantically related set of operations (methods) that the component is capable of performing. The interfaces implemented by a component represent the only way that other components can interact with it. Each interface is denoted by a reference that can be resolved at run-time to access the component instance that implements the interface.

Each method has a name and a set of inputs, outputs, and exceptions. The inputs enable a component's clients to parameterize the behavior of the method each time they request that it be performed. The outputs enable the component to convey to a client the results that pertain to having properly performed the method. The exceptions enable the component to convey to a client the fact that something unexpected was encountered during the course of performing the method (such as an error condition). A method completes by returning outputs or by raising exceptions. Methods need not have inputs, outputs, or exceptions.

The methods defined for an interface are invoked using a binary calling sequence. This means that the component that issued the call does not need to be aware of how the component that services the call is implemented. The components might be implemented using different tools and libraries, and even different programming languages. Further, components can interact with each other in a location independent manner. A component only needs a reference to another component's interface to perform calls against the component. Knowledge of the physical location of a component that services a call is not needed.

6.1.2 Interface Interrogation

The interfaces that a component implements can be determined by other components at run-time through direct interrogation. The interrogator uses the symbolic name of the interface, or an identifier that denotes the interface, to indicate the desired interface. If the interface exists, the component being interrogated returns a reference to the interface. Otherwise an error indication is returned.

It is assumed that all of the interfaces defined in this document include a common method that enables interface interrogation. The name and signature for this method is the same for all components implemented using a particular technology. The details of this method vary for different implementation technologies and are not specified in this architecture document.

6.1.3 Principal Interface

Every component implements at least one well-known interface, referred to as the component's *principal interface*. The principal interface includes the same interface interrogation method as a component's other interfaces. The name of the principal interface is the same for all components implemented using a particular technology. The principal interface enables components to perform initial interface interrogations because the name of the principal interface is known a priori, and because all components implement it.

The details of the principal interface and the methods that it supports vary for different implementation technologies and are not specified in this architecture document.

6.1.4 Interface Reference Registry

An interface reference registry is a service that contains references to component interfaces. Components can use the registry to obtain interface references to each other. A reference can be used to access a component via the referenced interface. Each reference is denoted in the registry by a symbolic name and/or description. This enables components to locate references of interest based upon a symbolic and/or logical description of the reference of interest.

It is assumed that an interface reference registry is provided by the underlying implementation technology. The means by which interface references are denoted and placed into the registry, and the means by which components access the registry to retrieve the references, are technology-dependent.

The registry is assumed to be a well-known service that logically resides on each clinical desktop. This means that each component on a desktop has an a priori technology-specific means for knowing how to locate the desktop's registry. This provides all components on a desktop with a common means to obtain references to each other.

6.1.5 Interface Reference Management

To ensure orderly system behavior, components must have a means of knowing whether or not other components possess references to any of its interfaces. This enables a component to determine when it needs to be in a running state (because there is at least one other component that possess a reference), and when it can terminate (because no components possess a reference). The means by which this is accomplished is technology-specific.

It is assumed that each component that holds an interface reference performs an implicit or explicit action, which is technology specific, that indicates it wants to use a particular interface reference that it has obtained (e.g., from the interface reference registry). It is also assumed that a component performs an implicit or explicit action, which is technology-specific, when it no longer intends to use a particular reference. The latter action is referred to as *disposing* an interface reference.

7 Patient Link Theory of Operation

Patient Link enables the user to select a patient once, from any Patient Link-enabled application, as the means for automatically “tuning” all of the Patient Link-enabled applications in the common context system to the same patient.

Patient Link also establishes the foundation for all other context management “links”. For this reason, many of the fundamental CMA principles and rules are explained in this chapter, but are framed in terms of Patient Link so as not to become too abstract, and therefore hard to understand.

7.1 Patient Link Component Architecture

The following context management interfaces for Patient Link are modeled and illustrated in Figure 11: Patient Link Component Architecture:

- **ContextManager** (CM) - implemented by the context manager; used by applications to join/leave a common context system and to indicate the start/end of a set of changes to the common context data.
- **ContextData** (CD) - implemented by the context manager; used by applications to set/get the data items that comprise the common context.
- **ContextParticipant** (CP) - implemented by an application that wants to participate in a common context system; used by the context manager to inform an application that the context has changed.
- **ImplementationInformation** (II) – implemented by the context manager and mapping agent; used by applications, context management components, and tools, to obtain details about a component’s implementation, including its revision, when it was installed, etc.

Formal definitions of these interfaces, as well as example interactions between the components via these interfaces, are presented later in this document.




Figure 11: Patient Link Component Architecture

7.2 Patient Subject

The context subject of *Patient* is defined for Patient Link. The context data identifier item for this subject is an alphanumeric patient identifier, such as a medical record number. The patient's name is not used as an identifier.

This identifier is unlikely to be universally unique. However, it is assumed that a population of patients across which the identifier is unique can be established. Each such population is referred to as a *site*, as it is typical that each population of patients corresponds to a physical site within an overall healthcare institution.

Consequently, a single patient may be identified using multiple patient subject identifier items. Each item is differentiated by a different site-specific suffix. An application shall be configurable such that it can be instructed on-site as to which suffix (of suffices) it is to use when it interacts with the context manager to set or get patient context data.

The format of a patient subject identifier item name includes a site-specific suffix. Use of this suffix, and the values that may be assigned to this suffix, is at the discretion of each healthcare institution at which a context management system is deployed.

In addition to identifier items, the patient subject also supports corroborating data items. The actual names, meaning, and data types used to represent the values for both patient subject identifier items and corroborating data items are defined in the document *Health Level-Seven Standard Context Management Specification, Data Definition: Patient Subject*.

An example of a patient subject identifier item appears below:

Patient Subject Identifier Item		
Example Item Name Format:	Example Item Name:	Example Item Value:
Patient.Id.MRN.site_name	Patient.Id.MRN.St_Elsewhere_Hospital	RAS1958-12939213-122

7.3 Patient Mapping Agent

An optional patient mapping agent is also part of the common context system. The patient mapping agent maps the identifiers for patients. Whenever an application sets the patient context, the context manager instructs the patient mapping agent (if present) to provide any additional identifiers it knows for the patient. The site-suffix for each of the mapped identifier items denotes the site for which the patient identifier is valid, for example:

Patient Subject Identifier Item	
Examples Item Names:	Example Item Values:
Patient.Id.MRN.St_Elsewhere_Hospital	123-456-789Q36
Patient.Id.MRN.General_Hospital	6668-3923-987122

Mapping agents are described in more detail in Chapter 8.

7.4 Context Change Transactions

All changes to the common context are governed by a context change transaction that is initiated by an application but is coordinated by the context manager:

- An instigating application initiates a context change transaction and sets the patient context within the context manager. This context contains the identity of the patient.
- The context manager consults the patient mapping agent (if present) and it adds data to the context manager's patient context. This data includes additional identifiers by which the patient is known.

- The context manager surveys the other applications, and if the transaction completes, they obtain pertinent patient context data from the context manager.

The high-level events that transpire when a user selects a patient are summarized in Figure 12. This description assumes that a patient mapping agent is present. The patient mapping agent is presumed to know the identifiers for all patients for all applications within the common context system. (See Chapter 8, Mapping Agents.)

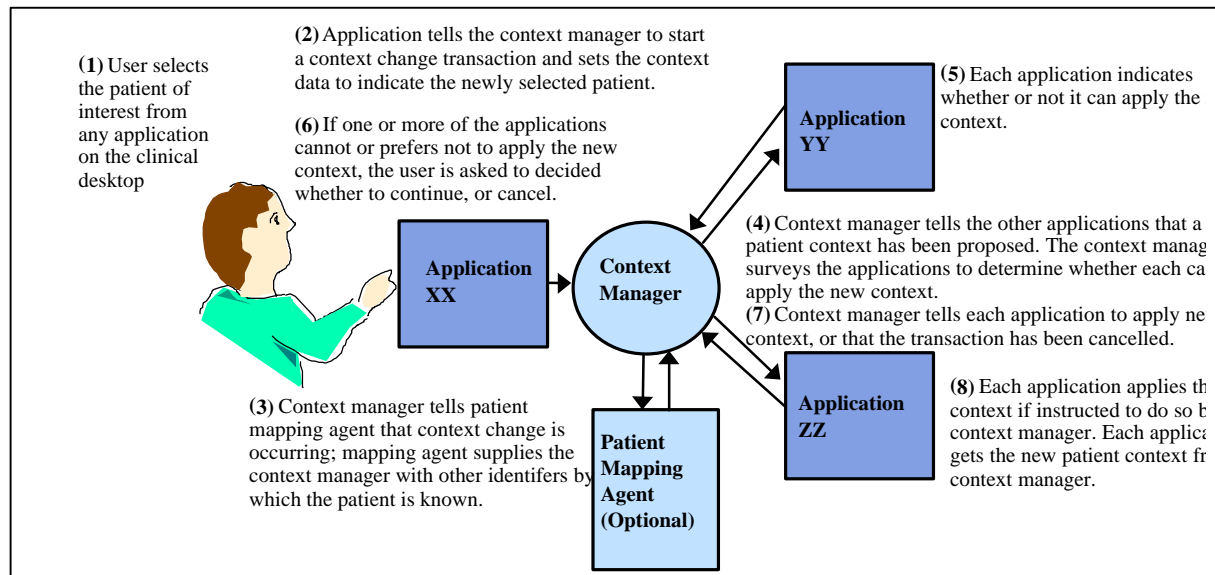


Figure 12: Patient Link Context Change Process

The details for how this process works and the responsibilities of the applications and CMA components are described next.

7.5 Joining the Common Context System

Applications join a common context system via the context manager for the system. The context manager's ContextManager interface is used for this purpose. The application obtains a reference to this interface by interrogating the context manager's principal interface. A reference to the context manager's principal interface is obtained from the desktop's interface reference registry.

An application typically retrieves the current common context data from the context manager's ContextData interface in order to establish its initial context. A reference to the context manager's ContextData interface is obtained by interrogating the context manager's principal interface or by interrogating the context manager's ContextManager interface. The context data is represented as a set of name-value pair items.

7.6 Context Change Transactions

Once it is a participant within a common context system, the context manager will inform the application of context data changes through the application's ContextParticipant interface. This data can be changed by any of the participants in the common context system. A participant executes a context change transaction to effect a context change. The transaction is coordinated by the context manager and involves the instigator of the transaction as well as the other participants.

The ContextManager interface is for beginning and ending a context change transaction. The ContextData interface is used for setting the new context data.

When a context change transaction is started, the context manager creates a transaction-specific version of the context data. This version of the context data is initially empty and does not contain any name-value pair items. This is to prevent data from the current context from becoming mixed with the data for the new context. Items are added to the transaction-specific context data during the course of the transaction.

This version of the context data is updated during the course of the transaction and is intended to be visible only to the application that instigated the transaction. All other applications continue to view the context data as it was when most recently published. The published context data is replaced with the context data set during the course of the transaction when the transaction completes successfully.

Prior to the first context change transaction, the published set of context data items is empty. Items are added during the course of subsequent transactions.

While the context manager serves as a holder for the current context data, its semantic understanding of the meaning of this data is intended to be minimal. Further, the specific items that constitute the context data are not assumed to be hardwired into the context manager implementation. This enables new context items to be defined over time without requiring changes to context manager implementations. This includes context items that represent identifier data as well as corroboration data.

Only one context change transaction is allowed at a time. Once it has started a change transaction, the instigator of the transaction is free to update the context data via the context manager's ContextData interface.

7.7 Transactional Consistency

In order to ensure that changes to this set of items are self-consistent, a participant must explicitly begin and end a context data change transaction. All of the context change operations that are performed within the scope of the transaction are treated as a single logical unit of work. When the transaction completes, either all of the changes are published, or none of them are. Other participants that access the ContextData interface to read the context data values will see the values as they were

prior to the transaction. Only the instigator of the transaction will see the values as they are during the course of the transaction. This prevents other participants from accidentally seeing inconsistent values.

This capability relies upon the proper use of context coupons, which are monotonically increasing identifiers that are assigned each time a change transaction begins. The context manager provides the instigator of a transaction with the context coupon when it is started. All other participants can only obtain from the context manager the coupon for the most recently committed transaction. A coupon is also provided as a parameter to most of the methods defined for the ContextData interface, thereby enabling the manager to determine whether it should respond in terms of the transaction-in-progress or the most recently committed transaction.

When the instigator of the context changes is done, it informs the context manager that the changes have been completed. A context manager may unilaterally decide to terminate a transaction and undo the changes if an application fails to indicate that it is done with its changes in a timely manner. (The context manager decides how long “timely” is. How this value is determined is an implementation decision.)

7.8 Context Change Notification Process

When the instigator completes the context changes, the context manager initiates a two-step change notification process wherein it determines whether to publish the shared context data changes. This process is inspired by the two-phase commit protocol used in many database systems to ensure transaction consistency. For the purposes of managing a common clinical context, the protocol has been simplified.

In the first step of the process, the context manager surveys the applications. Each application is informed that there are a candidate set of context data changes and is asked to indicate whether it can accept these changes. At this point, applications are provided with the context coupon value for this change transaction. This enables the applications to access the context data changes in order to consider specific data values as part of their decision about whether to accept the changes. This is accomplished via the context manager’s ContextData interface. It is possible for a participant to obtain just the values that have changed.

The context manager gathers the results of the survey and provides them to the application that instigated the context change. Depending upon the survey responses the application may be free to go ahead and publish the changes, or it may need to solicit guidance from the user about how to proceed. This guidance is required when there is at least one surveyed application that:

- is unable to apply the context change because it is blocked (e.g., it is a single threaded application that has a modal dialog open); these applications are referred to as “busy”

- might lose work performed by the user if it applies the context changes (e.g., the user was in the process of entering data that would not be applicable in the new context); these applications are referred to as having “conditionally accepted” the context changes.

For each application in one of these states, the user is provided with a description that identifies the application and explains its situation.

When user guidance is required, the following choices are offered:

- **Cancel** - the context change is canceled; the context changes are not published.
- **Break Link** - the context changes are applied just to the application with which the user initiated the context changes. This application essentially breaks away from the common context system until the user explicitly instructs the application to rejoin the system. The application that has broken away displays a distinct visual cue indicating that its context may be different from the other applications (e.g., it might display a warning message in a prominent location)³.
- **Apply** - the context data changes are applied to all of the applications, including those that indicated that they might lose work performed by the user; *this choice is allowed only when there are no busy applications.*

It is the responsibility of any application that enables the user to instigate a context change to present, when necessary, a dialog that obtains the user’s guidance as described above. The appearance of the dialog and the commands that the user can choose from are specified in each of the HL7 context management technology-specific user interface specification documents. This will ensure a consistent and familiar set of interactions for users across CMA-conformant applications.

The ability for any one application to require the user’s direct involvement in mediating context changes provides an important efficiency and safety feature.

The efficiency feature addresses the fact that changing the context may cause an application to lose work performed by the user. This work may be in the form of data entered but not yet saved by the user, or may be in the form of an expensive computation (such as a lengthy database retrieval) that would need to be re-performed in light of a context change. Allowing the user to decide how to proceed in these circumstances minimizes the likelihood that the user will unintentionally lose work.

The safety feature addresses the fact that it may not always be possible to force an application to accept changes to the context data. Specifically, this is the case for blocked, or busy, applications.

³ A specific visual cue will be recommended within each of the HL7 context management technology-specific user interface specification documents.

If context changes were automatically applied piecemeal to just the applications that could respond, applications could become out of synchrony with regard to their clinical context, without the user being aware of the situation. For example, the user might assume that after a context change, all of the applications are displaying data for the same patient when in fact they are displaying data for different patients. The approach described above avoids this problem. This is because the only time that an application can become out of synchrony with regard to the clinical context used by the other applications is when the user has explicitly instructed it to break away.

In the second step of the two-step change notification process, the applications in the common context system are informed about whether or not the context changes are to be applied. If all of the surveyed applications indicate that they accept the changes, then the changes are applied and are reflected as the new context state. If the user indicates that the changes should be canceled, then the changes are discarded.

Once a participant has been informed that the context data has changed, it is free to inspect the data to obtain the new values if it has not already done so (again, using the context manager's ContextData interface). The participants can also assume that all of the other participants are applying the same context data.

In either case, the context change transaction completes when all of the applications have been informed of the outcome of the survey. If the context manager is unable to inform an application of the survey outcome, it will keep trying periodically, unless the manager determines that the application has terminated. The periodic attempt to notify a non-responsive application does not prevent the transaction from completing, nor will it prevent a new transaction from being started.

7.9 Leaving a Common Context System

When an application terminates, it explicitly leaves the common context system by informing the context manager via its ContextManager interface. At this time, the context manager shall dispose of any application interface references that it possesses, and the application shall dispose of any context manager interface references that it possesses.

A diagram of the overall common context system model is presented in Figure 13, followed by component interaction diagrams that represent typical common context data update transactions.

7.10 Behavioral Details

7.10.1 Application Behavior When it Cannot Cancel Context Changes

It is possible that an application that instigated a context change transaction cannot easily implement the capability to cancel the transaction. In this case, it is acceptable for the application to not offer canceling the transaction as an option to the user. The details of how this appears to the user are

specified in each of the HL7 context management technology-specific user interface specification documents.

7.10.2 Application Behavior When it Does Not Understand Context Identifiers

It is possible that an application is unable to interpret any of the context identifier items that were set when the current context was established by another application. For example, the selected patient might not be a patient known to the application.

An application that is unable to interpret any of the identifiers shall still participate in the context change transaction. This situation is not a basis for the application to prevent the transaction from proceeding. Specifically, the application shall not use the surveying process to reject the context change.

However, at the completion of the transaction, the application shall clearly indicate to the user that it is unable to apply the current context. The application shall not show any patient data. The details of how this indication appears to the user are specified in each of the HL7 context management technology-specific user interface specification documents.

7.10.3 Application Behavior with Regard to an Empty Context

The context is empty when a context system is first initialized. (See Section 5.6.8, Representing an Empty Context Subject). When this is the case, all of the applications in the context system shall clearly indicate to the user that there is no current context. The details of how this indication appears to the user are specified in each of the HL7 context management technology-specific user interface specification documents.

7.10.4 Surveying Details

During the context change survey, the context manager informs each of the applications in the common context system (except for the application that instigated the changes) that there are pending context data changes. When an application is surveyed, it shall create a visual cue that indicates it is about to change its clinical context *before* responding to the survey⁴. It shall not change its context yet. The context-changes-pending indication shall only be removed once the context manager has informed the surveyed application about how to proceed.

Under normal circumstances, the application will eventually be notified by the context manager about whether or not the context changes should be applied. However, if the context manager is unable to inform the application about how to proceed (e.g., because the application blocked after responding to the survey but before being notified that the context changes have been accepted), the user will at least

⁴ A specific visual cue recommended within each of the HL7 context management technology-specific user interface specification documents.

be able to determine that the application may not be in synchrony with the other applications. This is because the application is presumably still displaying a visual cue that indicates it might change its clinical context. The fact that this cue is still being displayed *after* the context has changed clues the user that there is a problem with the application.

An application can explicitly respond to a context change notification survey by indicating one of the following:

- **Accept:** It is willing to accept the context data changes and to change its internal state accordingly if the changes are published.
- **Accept-Conditional:** It is in the midst of a task that might cause work to be lost if the user does not complete the task; if the changes are published it is willing to terminate the task, accept the context data changes and change its internal state accordingly.

If the changes are subsequently published, an application can defer changing its internal state until some time in the future (for example, when it regains the focus for user-inputs). However, it must offer a visual cue that indicates it not in synchrony with the new context. For example, it might blank out its data display or minimize itself.⁵

An application that cannot interpret the context data (e.g., does not know who the patient is) should accept the changes. However, the application should clearly indicate to the user (e.g., by displaying a message) that it cannot apply the current context data.

The context manager infers an implicit response from an application under the following conditions:

- **Terminated:** the context manager has determined that the application has terminated without first informing the context manager
- **Busy:** the context manager has determined that the application is still running but is unable to answer the survey (e.g., the application is single-threaded and has a modal dialog open)

It is not possible for a surveyed application to explicitly reject, and therefore prevent, a context change.

The context manager gathers the survey responses and returns them to the application that was used to instigate the context change transaction. Applications that have responded with *accept-conditional* are expected to also provide a succinct but informative description of the consequences to the user of applying the context changes. The context manager then prepends the name of the application (provided by the application when it joined the common context system) to the description. This description is shown to the user by the instigating application.

⁵ A specific visual cue is recommended within each of the HL7 context management technology-specific user interface specification documents.

The context manager also provides the instigating application with a succinct but informative description about any applications that are busy. This description includes the name of the application. This information is provided by the context manager on behalf of these applications, as they are unable to do so for themselves. This description is also shown to the user by the instigating application.

Applications that have terminated do not affect the survey process. The context manager considers such applications to no longer be part of the common context system. Any information that the manager is maintaining about terminated applications is discarded.

Applications that have suspended their participation in the context are not involved in the survey process.

Applications that have joined the system but indicated that they do not want to participate in surveys are not involved in the survey. However, they are informed along with the other participants whenever the decision to accept the changes is published. (They are not informed about decisions to cancel changes, as this information would be irrelevant.)

7.11 Common Clinical Context Use Model

The Common Clinical Context Use Model (Figure 13) illustrates a system with four actors (Authorized User, Healthcare Application, Context Manager, and a System's Administrator) applying forces on three use cases. The use cases are Lifecycle of Common Context, Context Selection Change, and Abnormal Termination of Common Context.

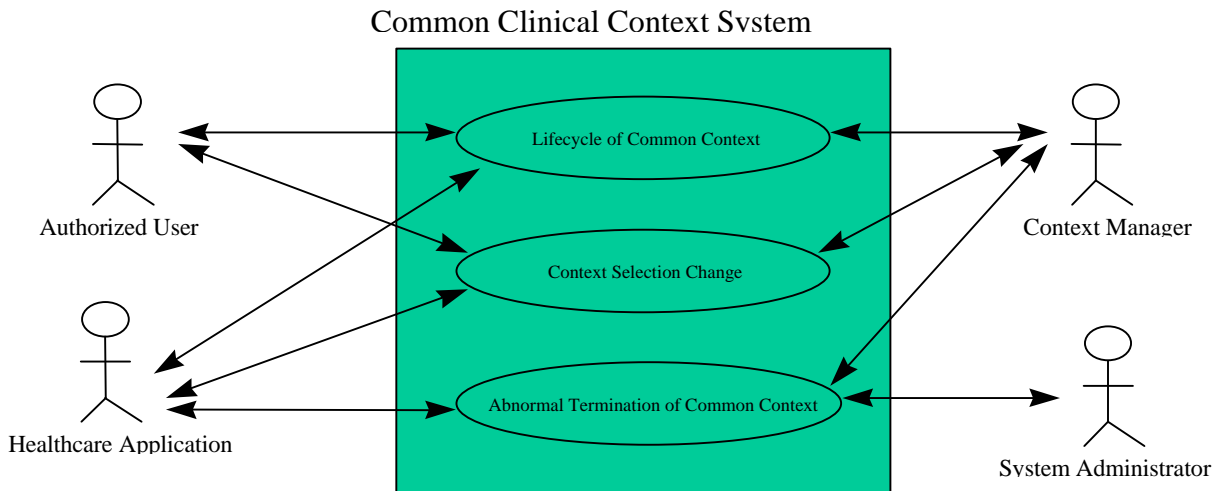


Figure 13: Common Clinical Context Use Model

The common clinical context system is presented by providing a diagram of each use case followed by interaction diagrams illustrating different behavioral flows of the associated use case. Each use case has an associated description, which is provided below. Further, for brevity the specific interface names (ContextManager, ContextParticipant, and ContextData) are not used; their abbreviations are used instead (CM, CP, and CD). Also, the word “interface” is abbreviated to “iface”. The diagram notes (illustrated as a sheet of paper with corner folded over) are from a software developer’s perspective, not the user of the application.

7.11.1 Lifecycle of Common Context

A common context does not initially exist. An application must establish the common context. The common context ceases to exist when there are no longer any applications participating in the common context. Figure 14, Interaction Diagram 1, and Interaction Diagram 2 illustrate this use case.

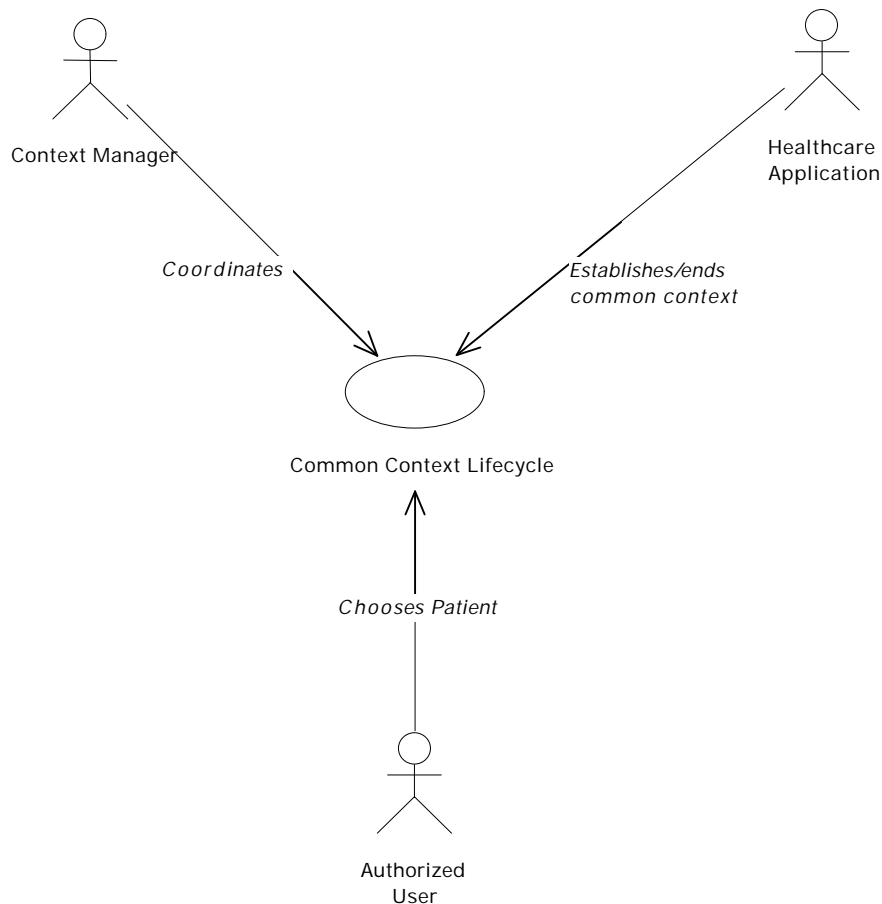
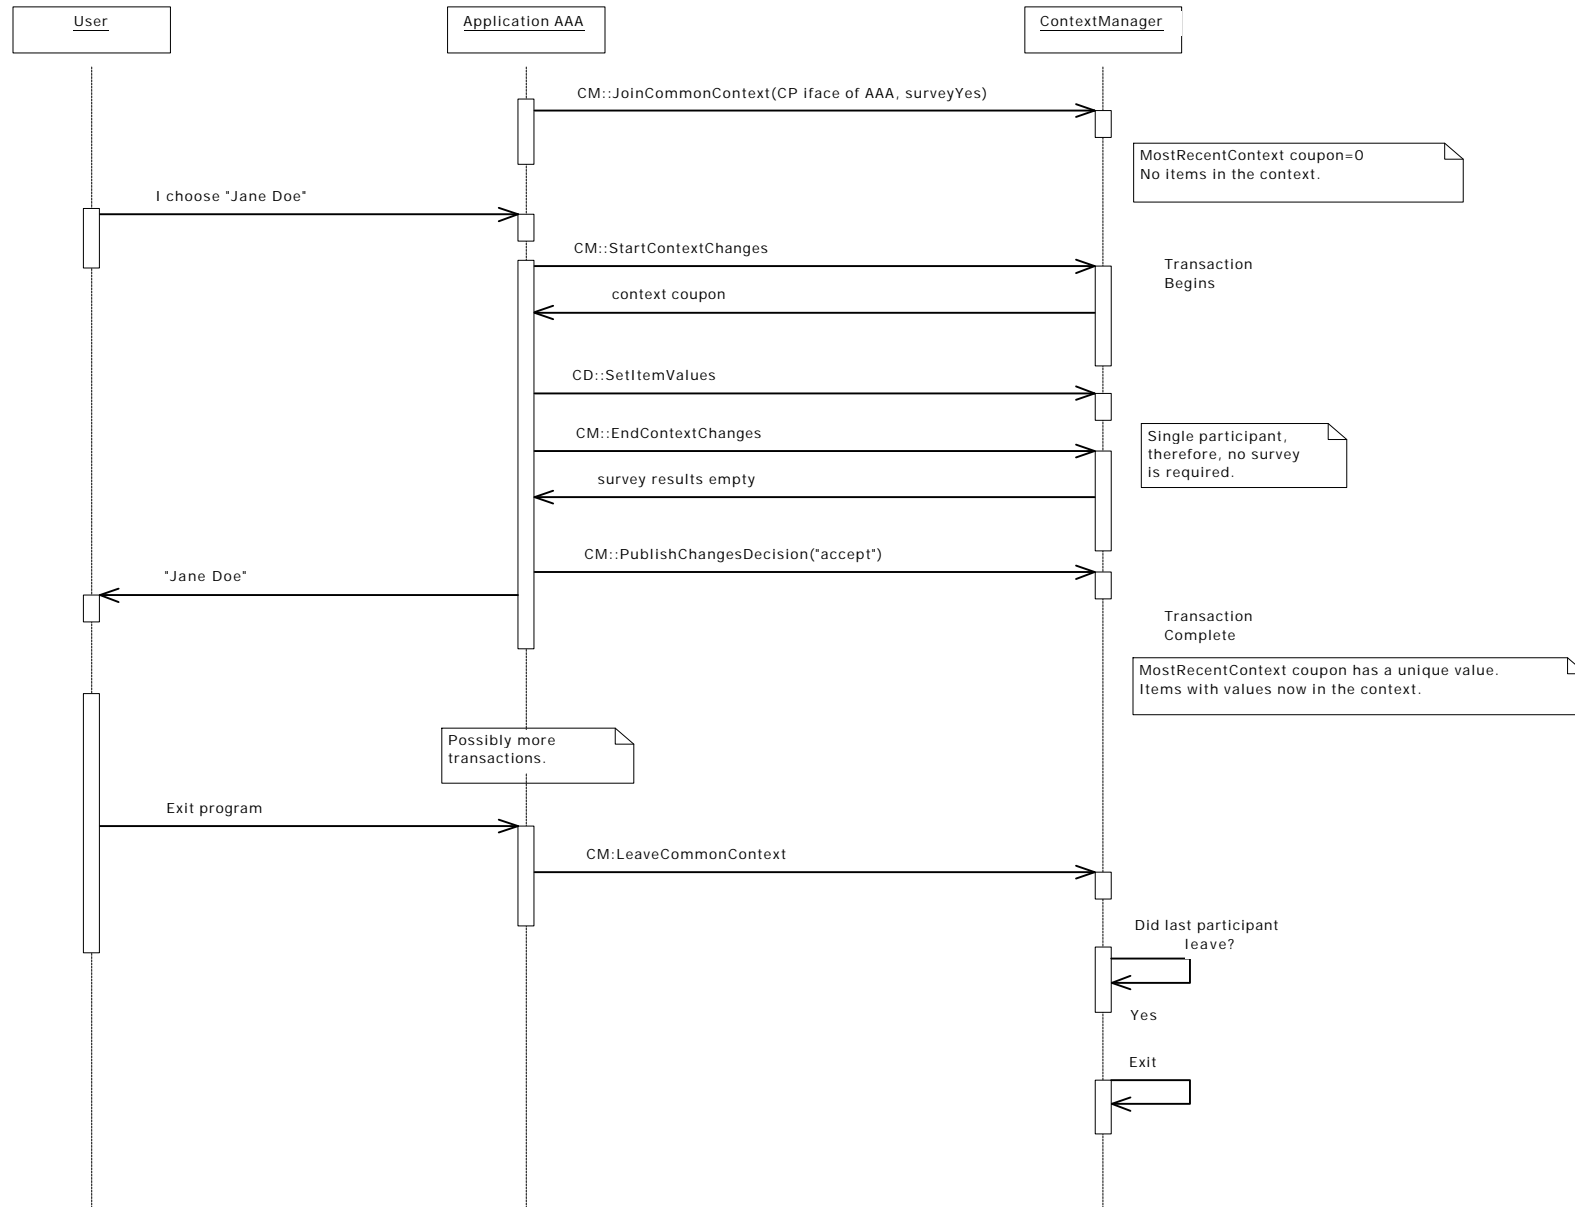


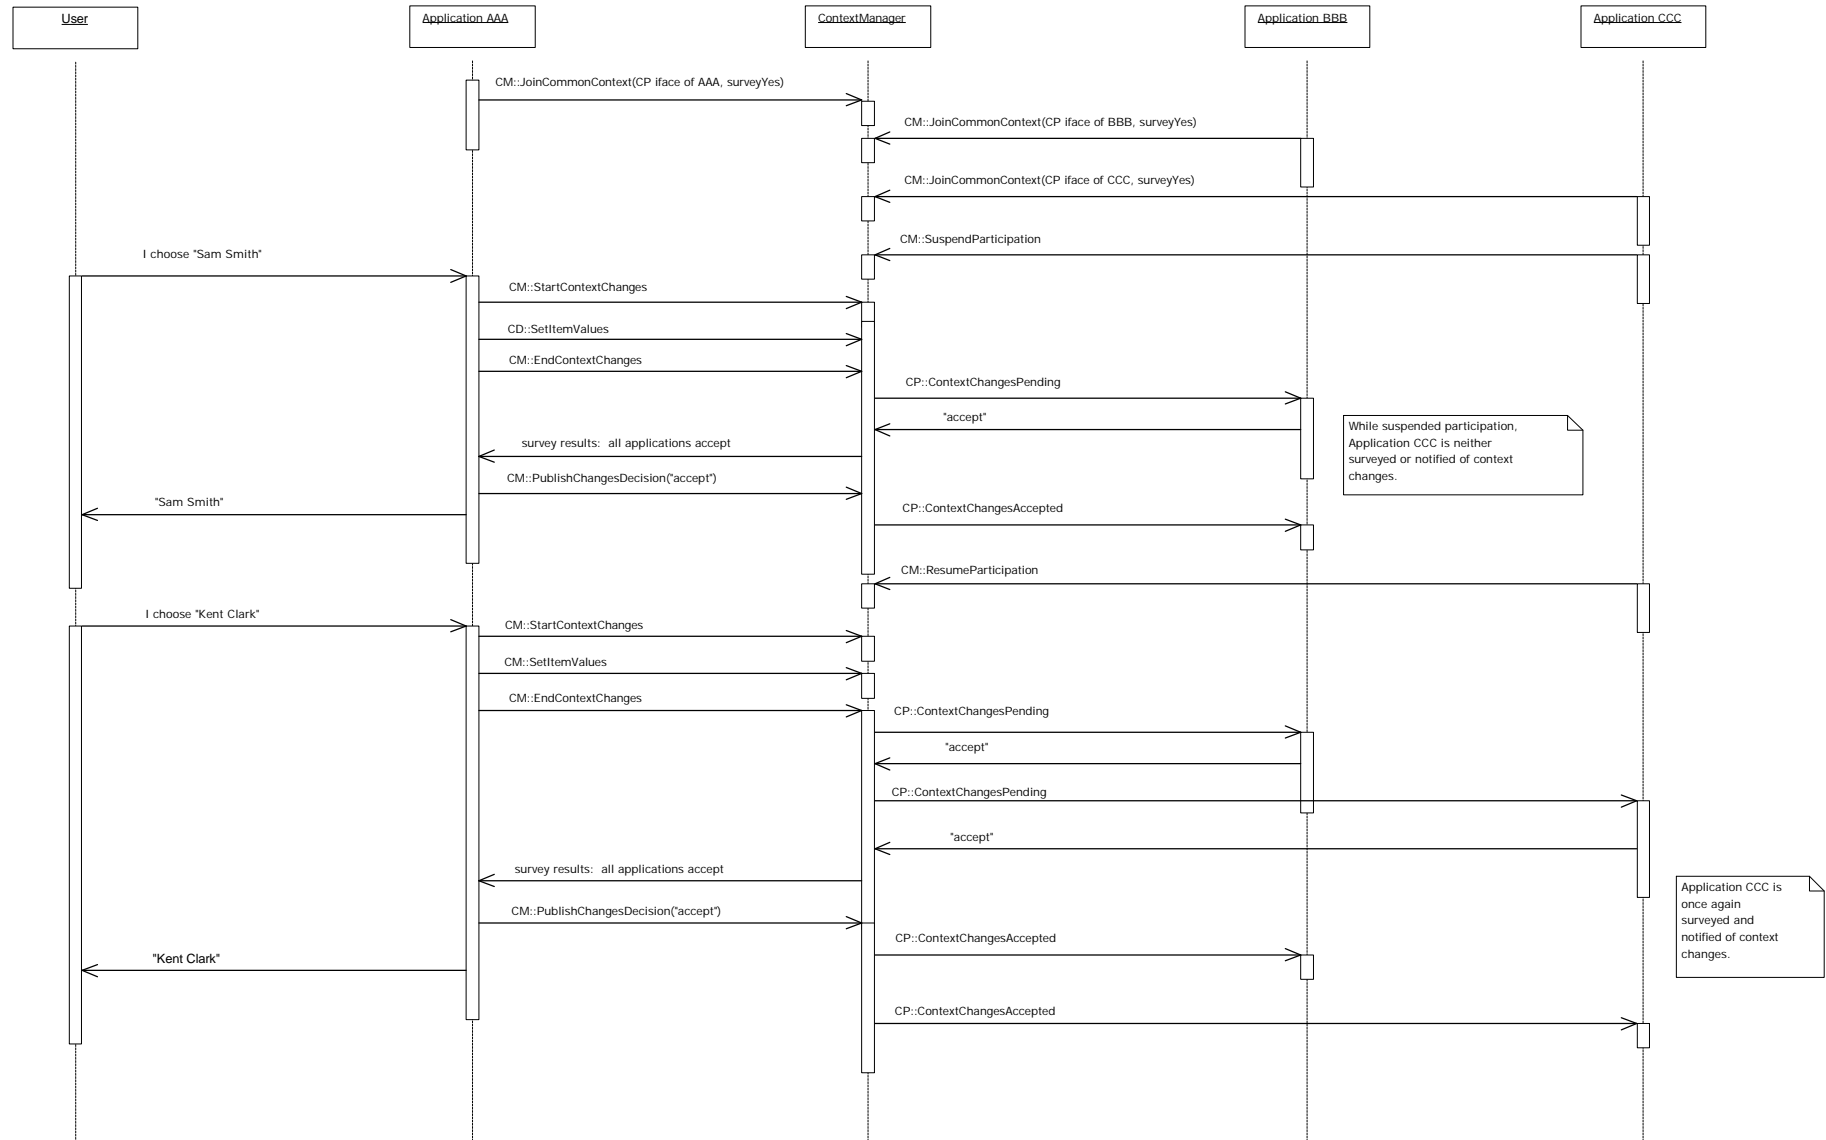
Figure 14: Common Context Lifecycle Use Case

Context Management Specification, Technology and Subject-Independent Component Architecture



Interaction Diagram 1: Common Context Lifecycle

1



2

3

Interaction Diagram 2: Suspending/Resuming Context Participation

7.11.2 Context Selection Change Use Case

The Context Selection Change use case assumes a patient context has been established. The user is currently focused on one application, while several other healthcare applications may be executing on the same host machine. The user chooses to change the selected patient from “Jane Doe” to “Sam Smith”.

Figure 15 illustrates this use case. There are several possible instances of this use case which are provided in Interaction Diagram 3 through Interaction Diagram 10.


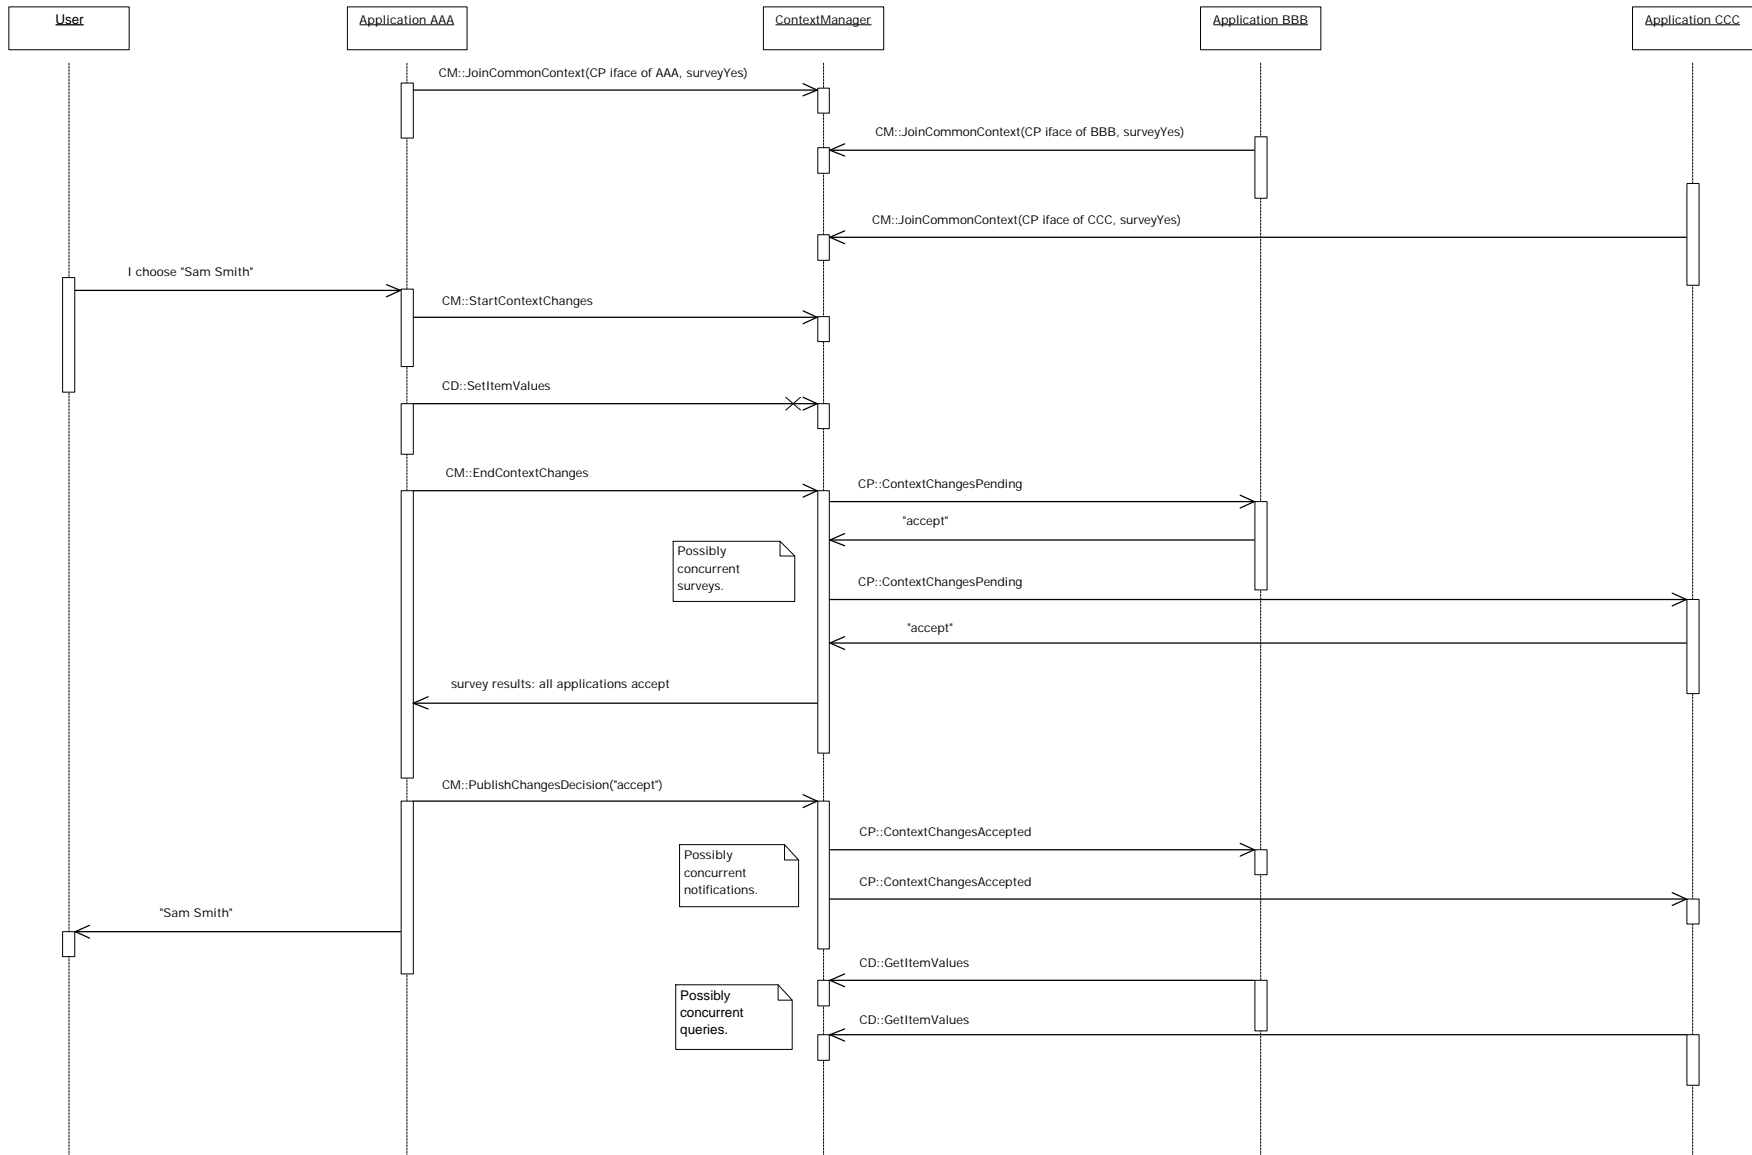


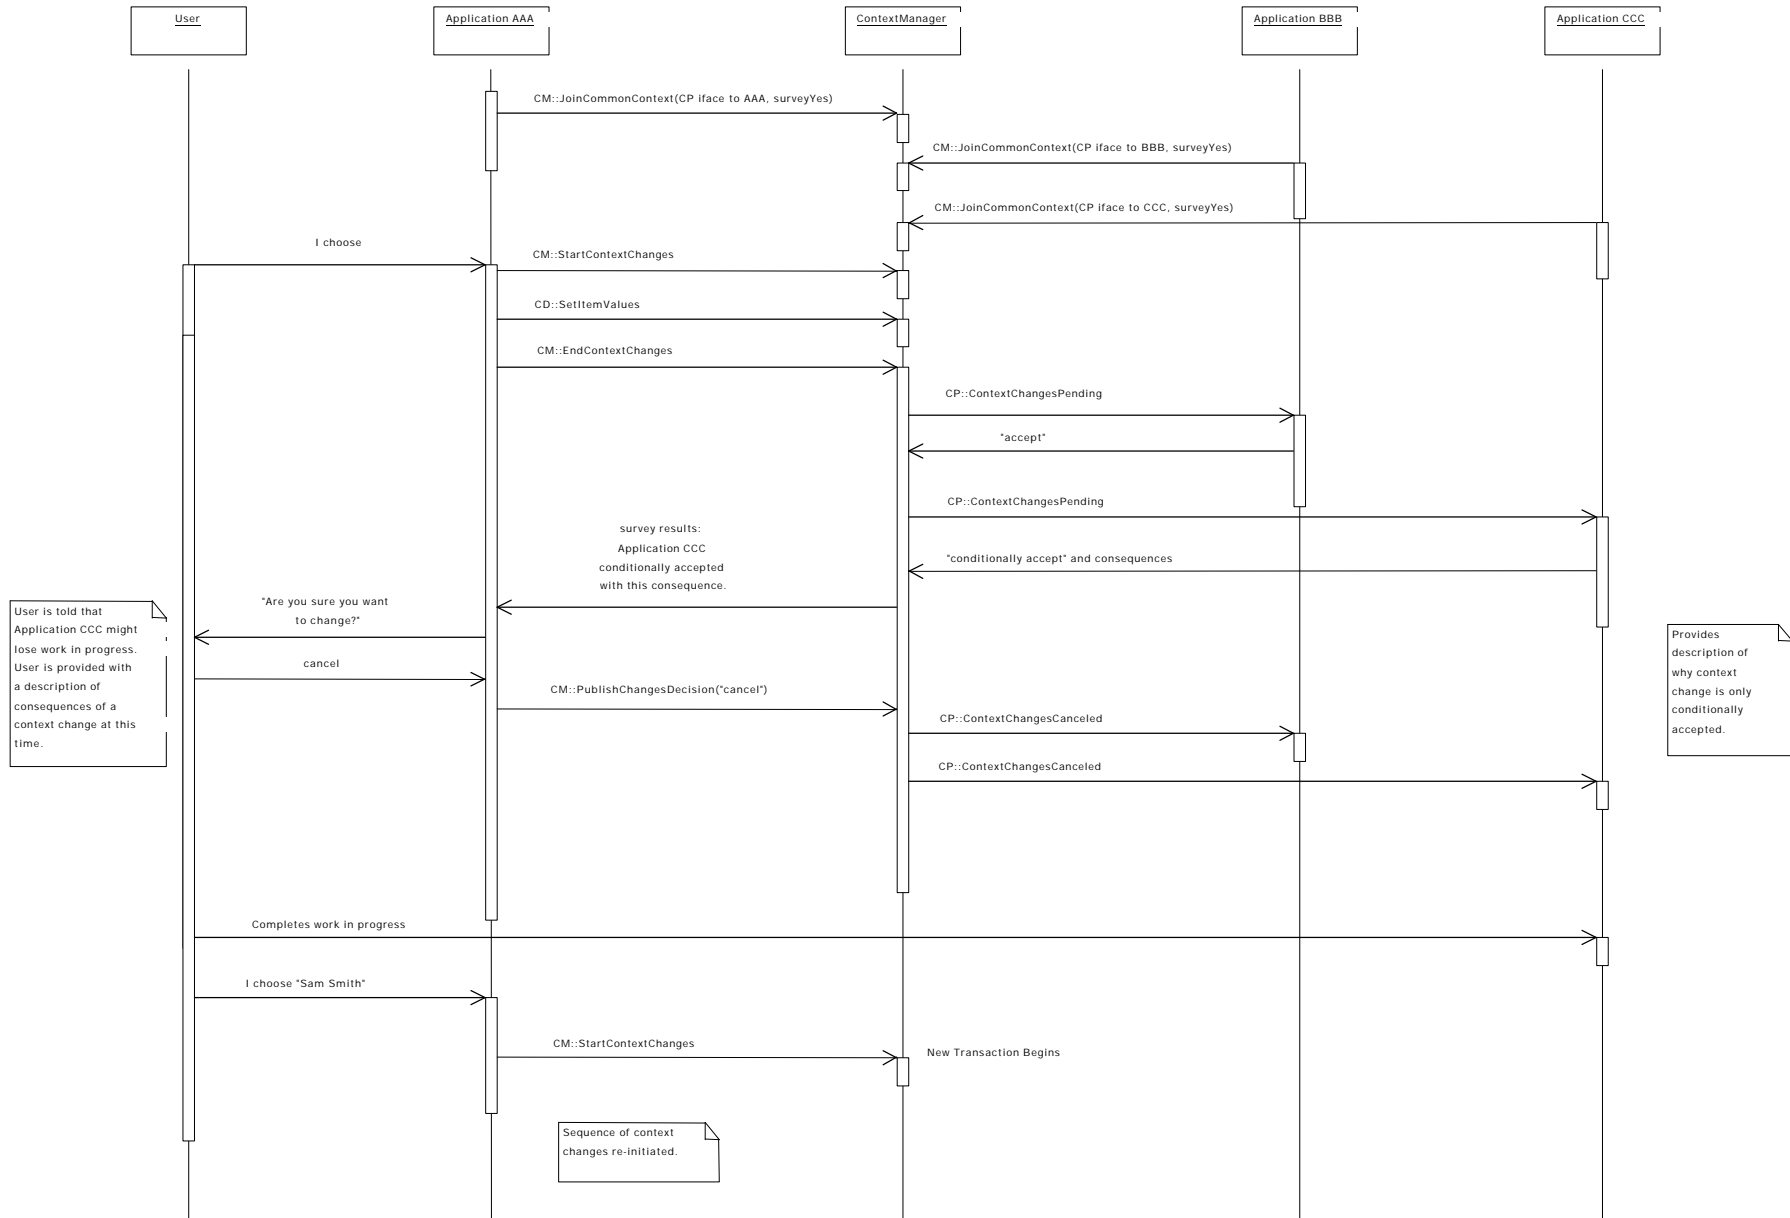
Figure 15: Context Selection Change Use Case

Context Management Specification, Technology and Subject-Independent Component Architecture



Interaction Diagram 3: All applications accept the changes

Context Management Specification, Technology and Subject-Independent Component Architecture

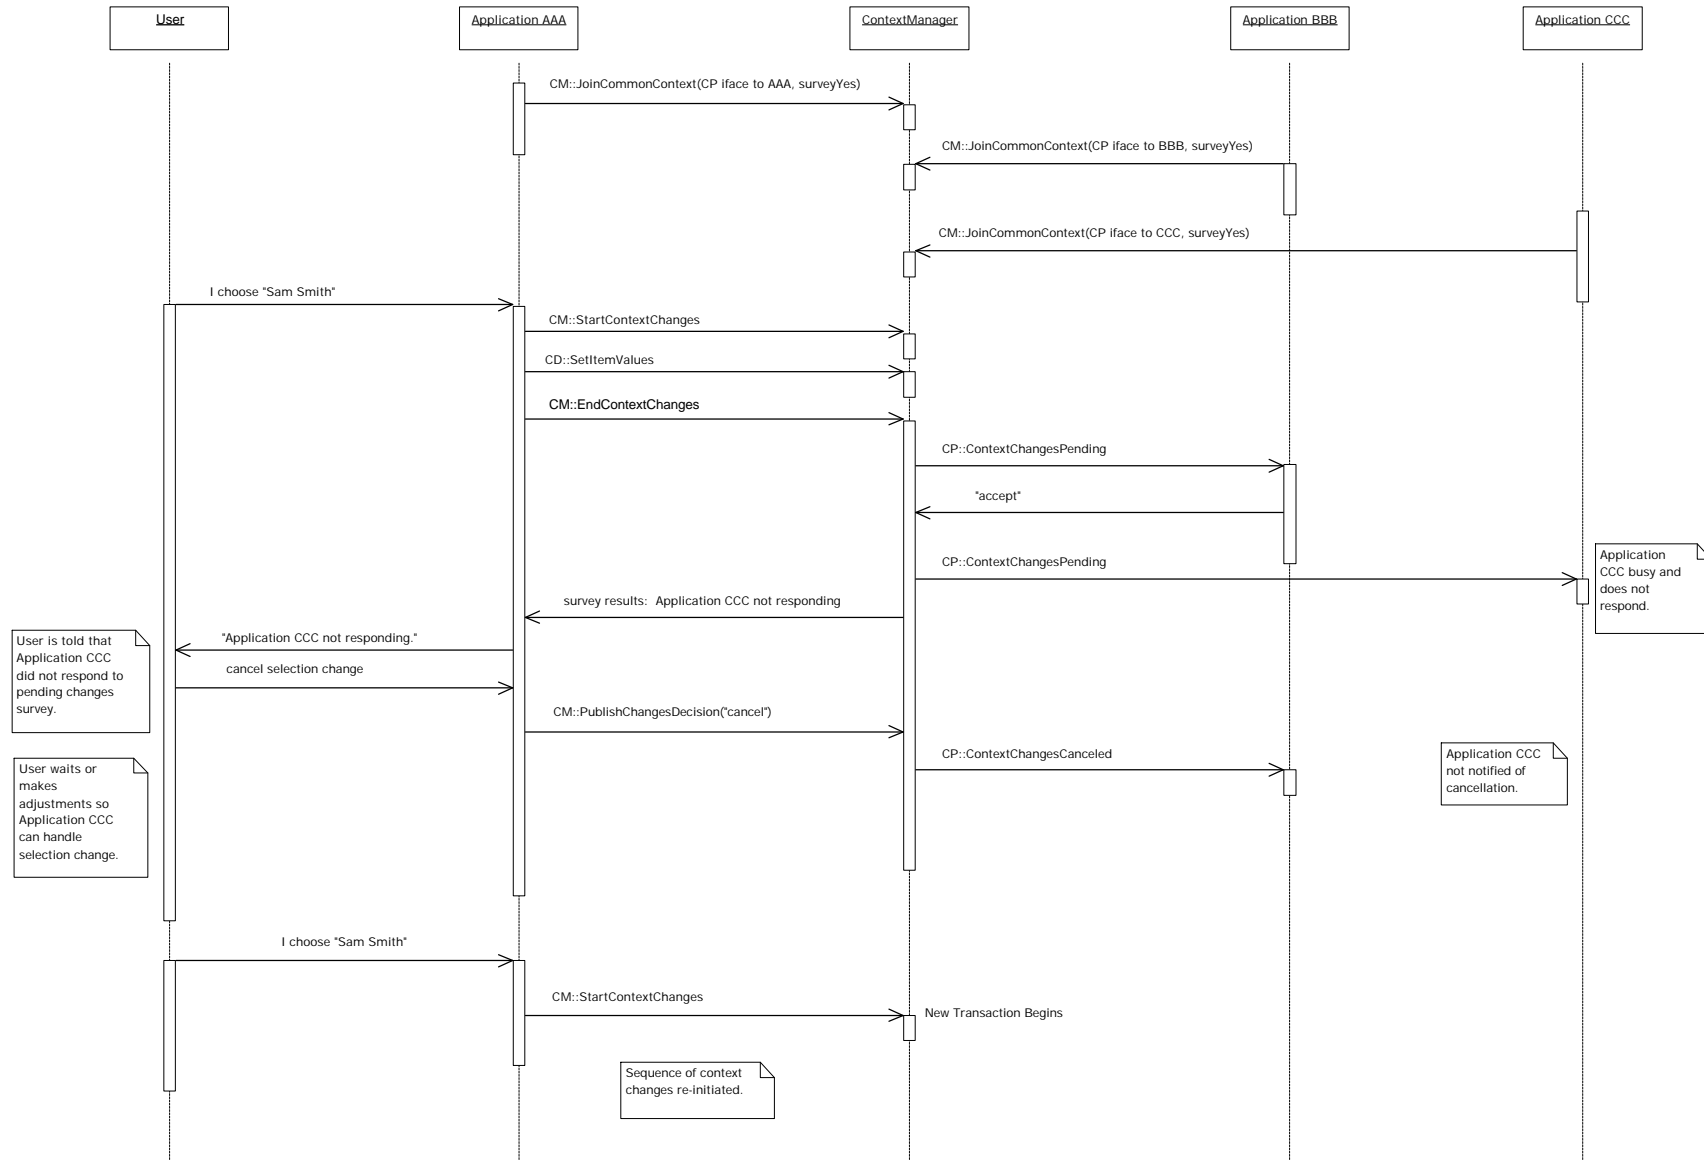


1

2

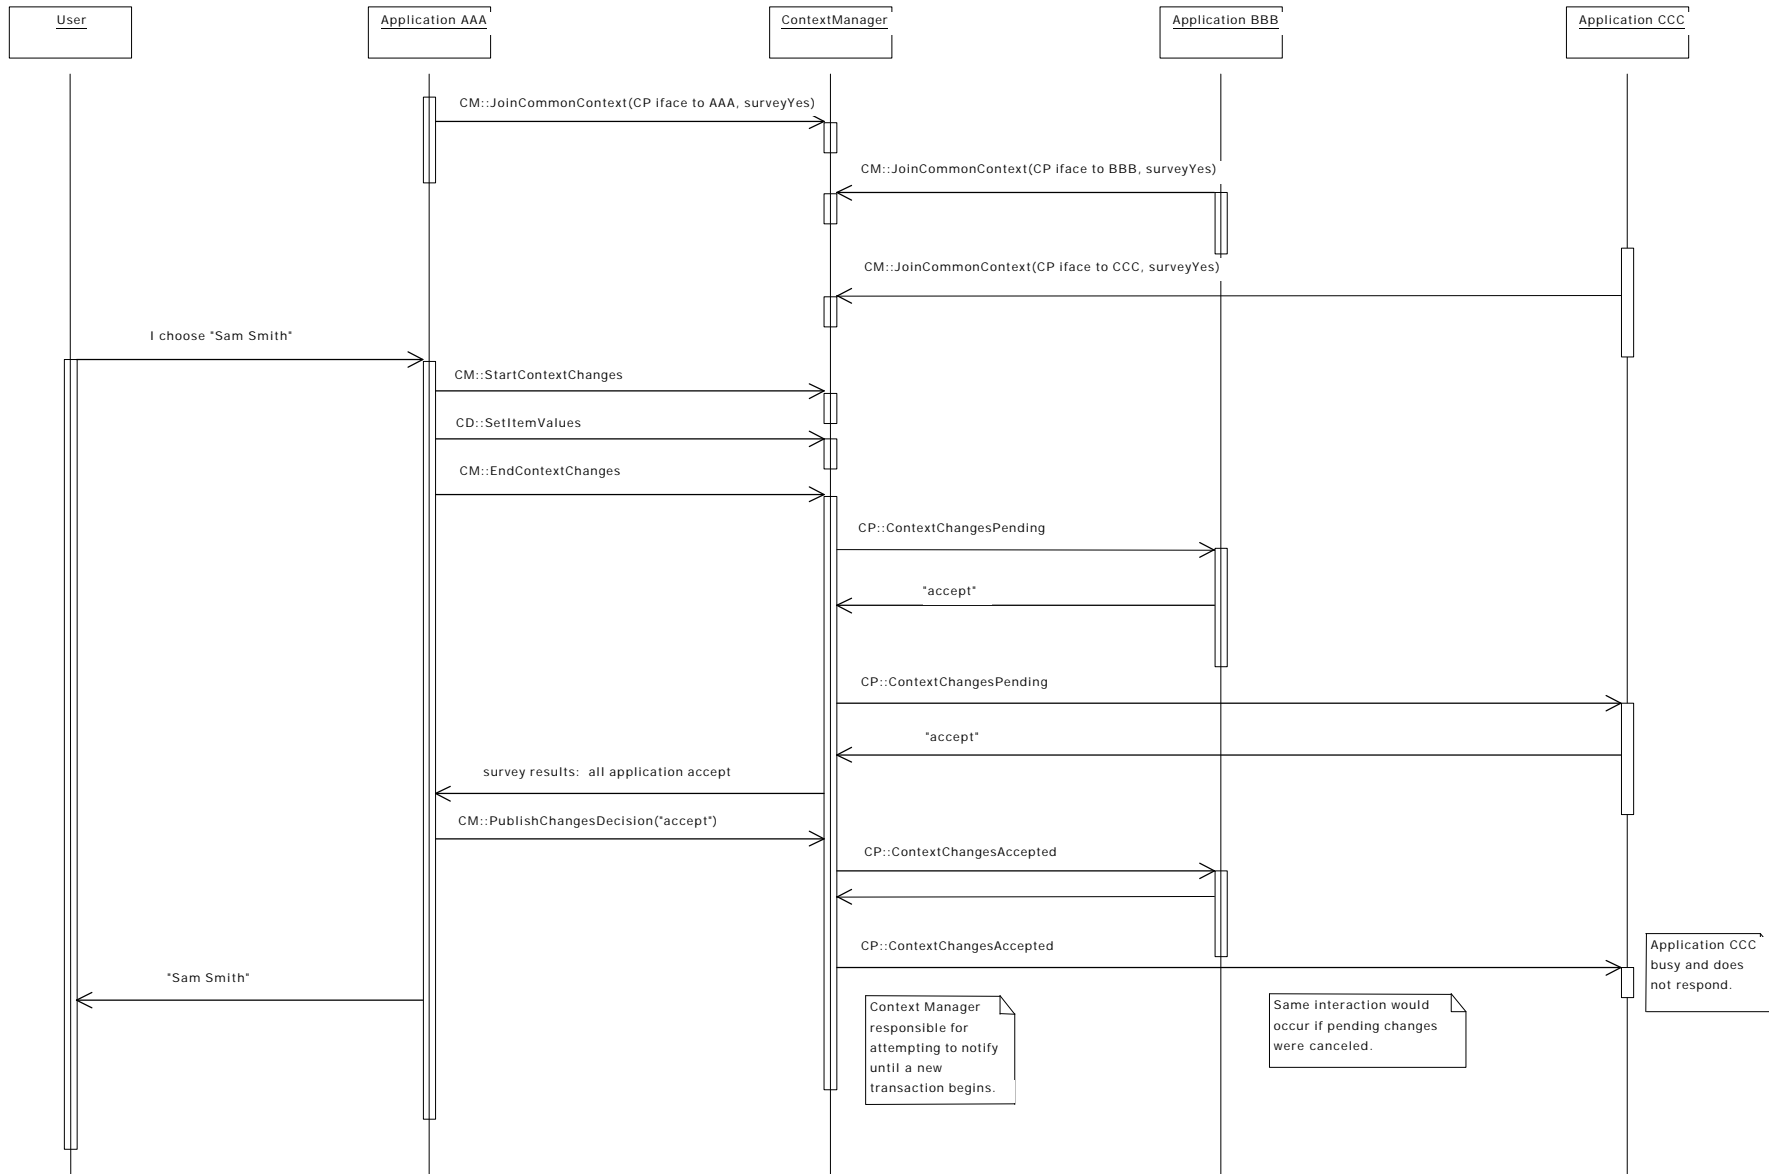
Interaction Diagram 4: An application conditionally accepts the changes; user decides to cancel changes

Context Management Specification, Technology and Subject-Independent Component Architecture



Interaction Diagram 5: An application does not respond to survey

Context Management Specification, Technology and Subject-Independent Component Architecture



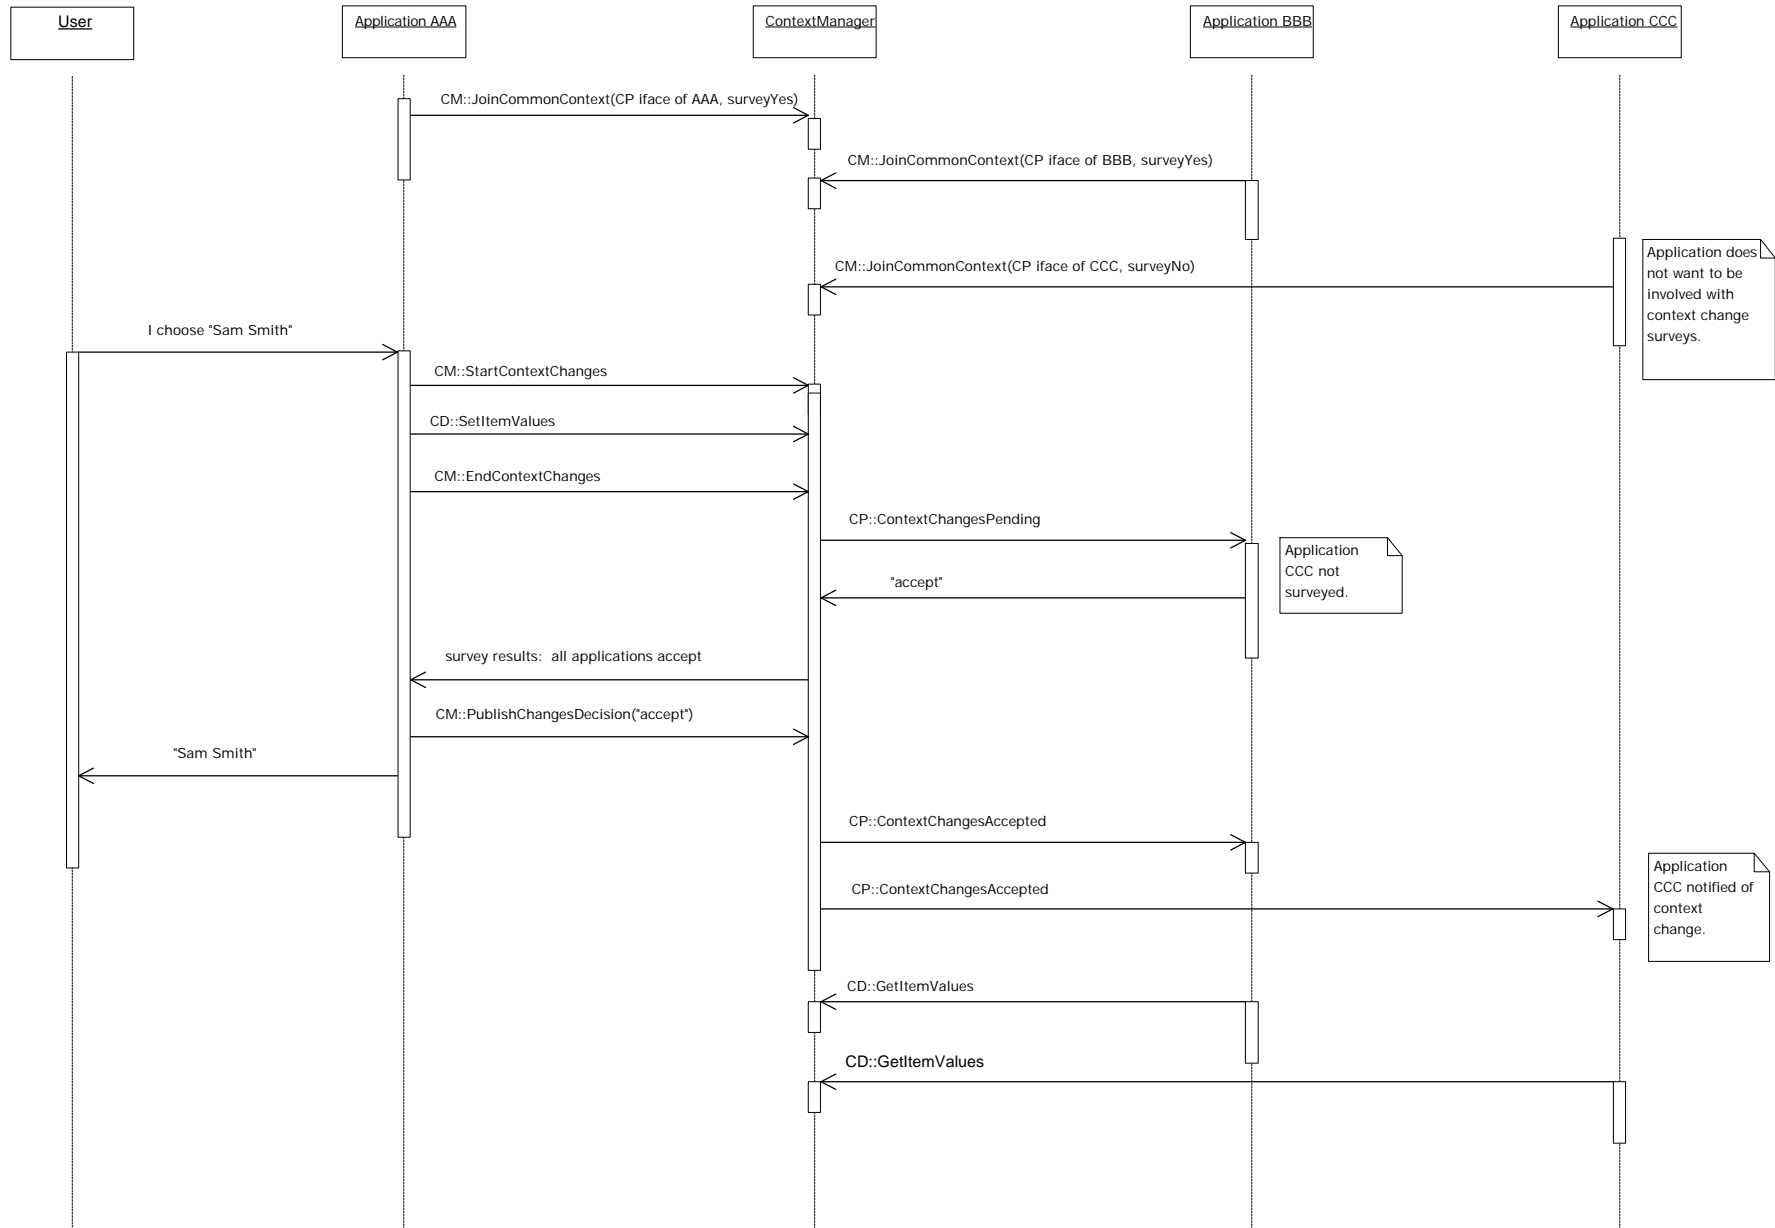
Interaction Diagram 6: An application does not respond to change notification

Context Management Specification, Technology and Subject-Independent Component Architecture



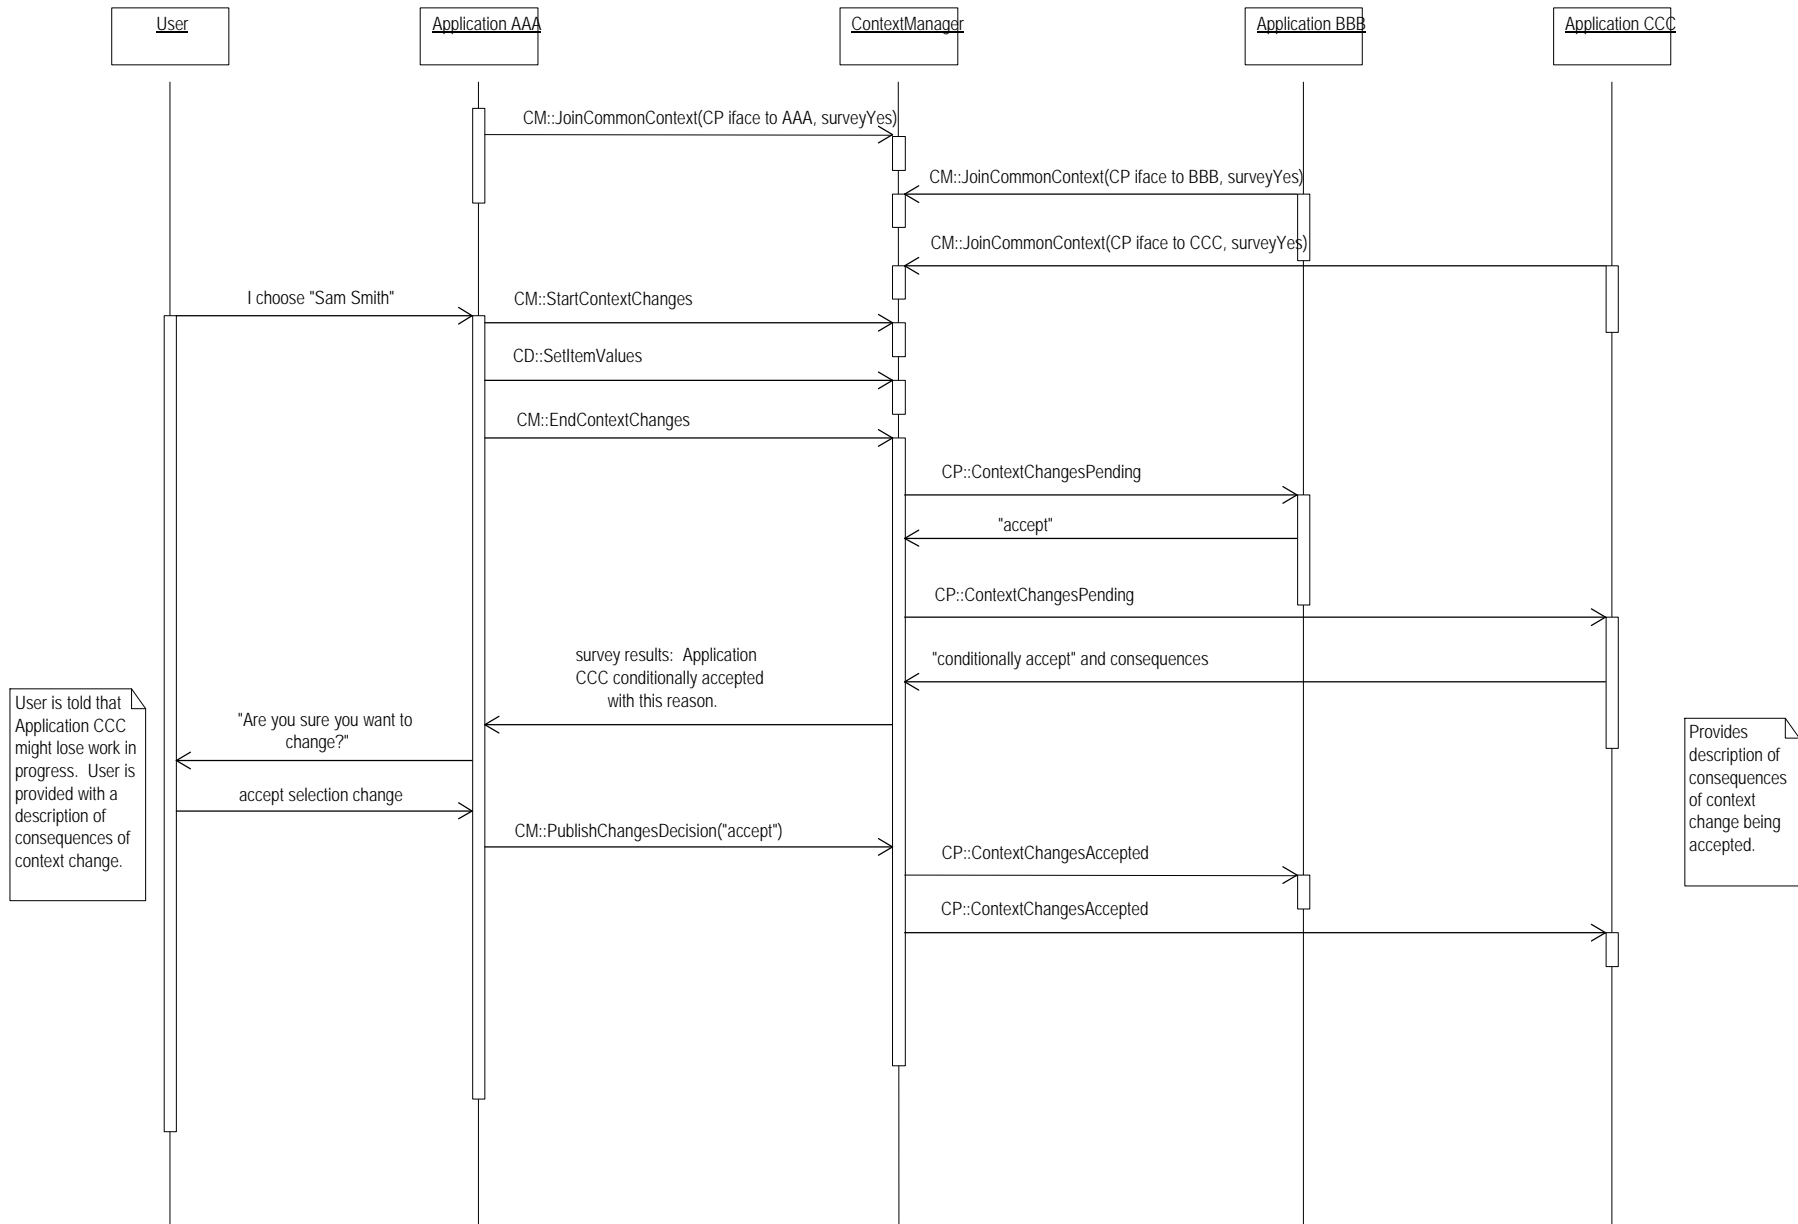
Interaction Diagram 7: An application responds after context change transaction has completed

Context Management Specification, Technology and Subject-Independent Component Architecture



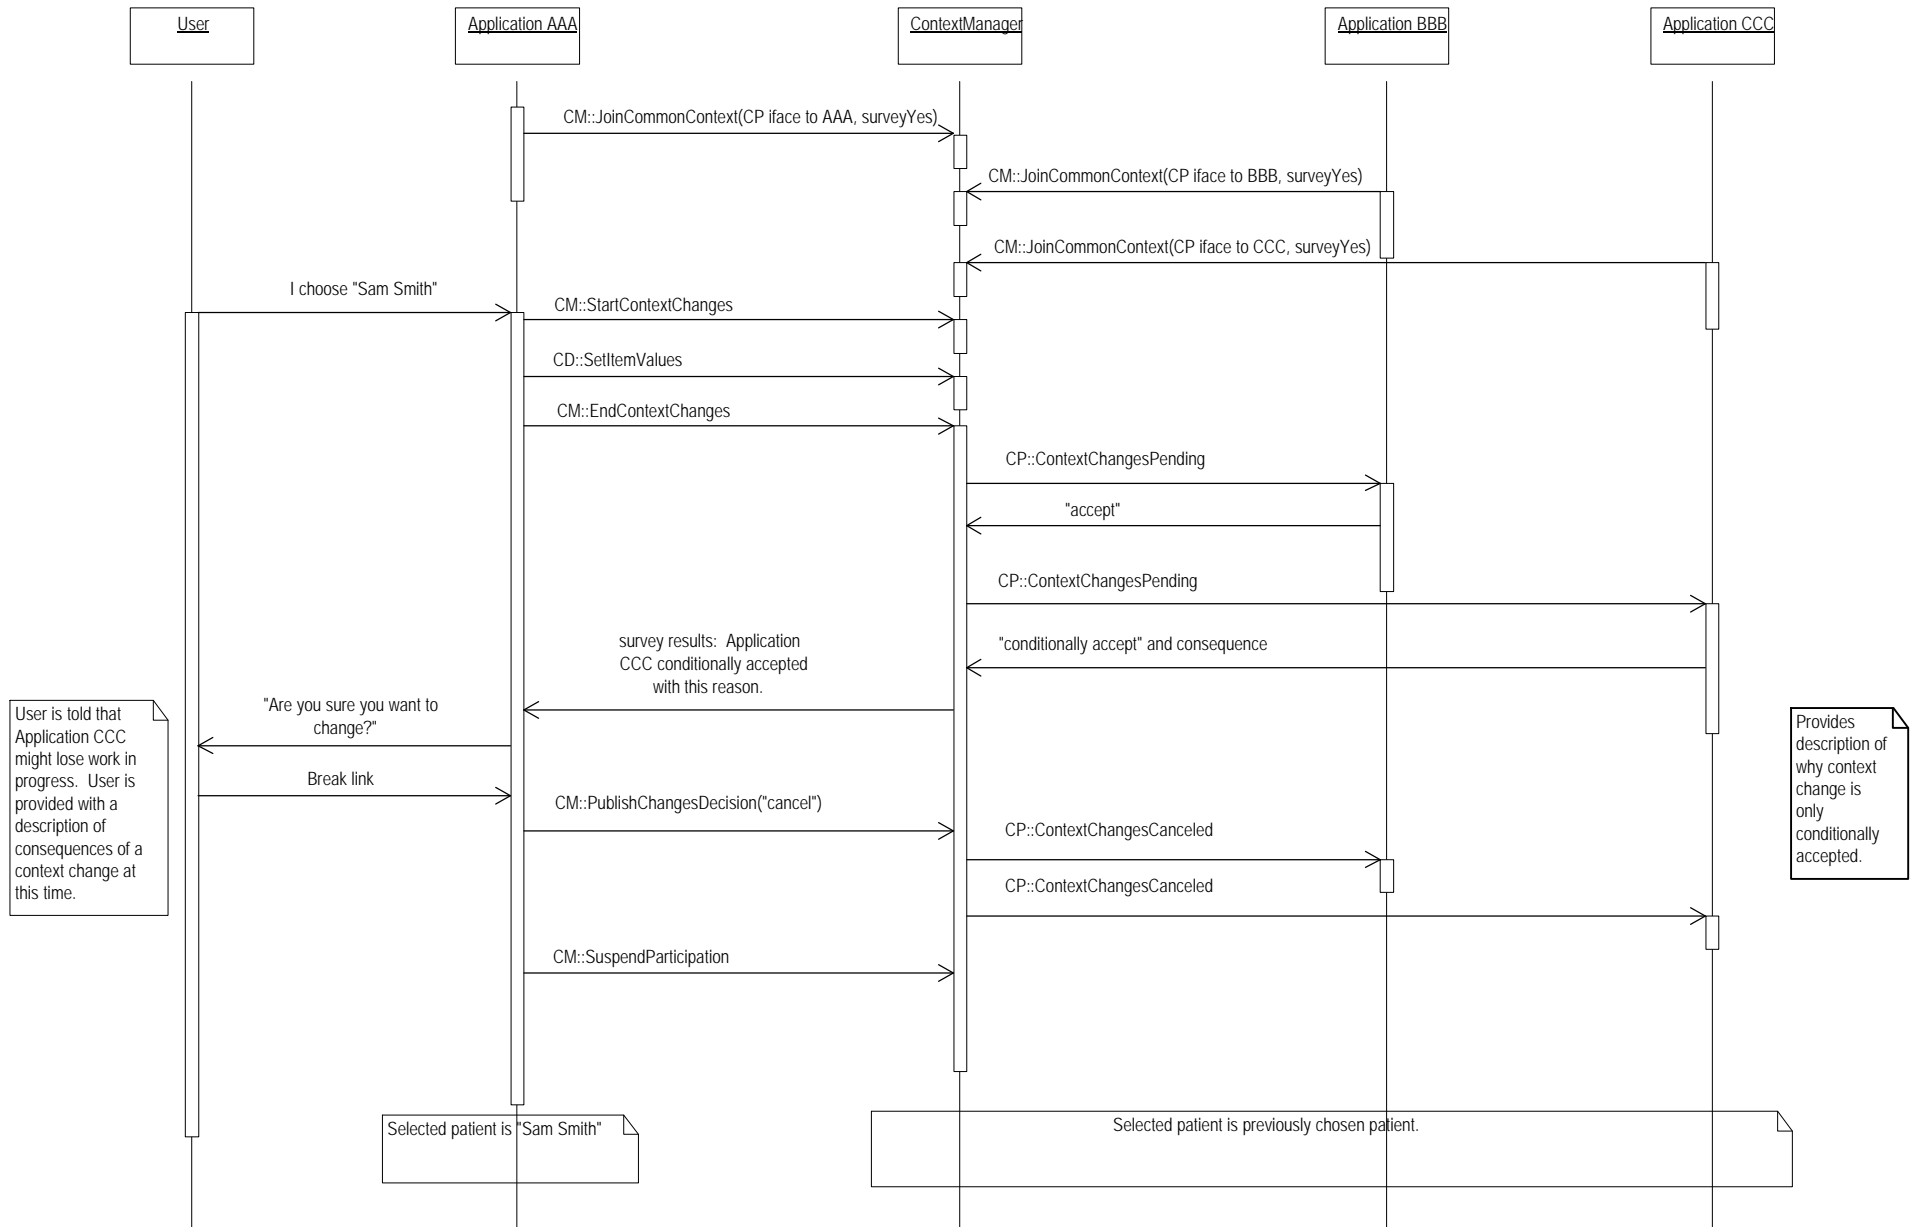
Interaction Diagram 8: A non-surveyed application participates in context change

Context Management Specification, Technology and Subject-Independent Component Architecture



Interaction Diagram 9: An application conditionally accepts the changes; user decides to accept consequences of change

Context Management Specification, Technology and Subject-Independent Component Architecture



Interaction Diagram 10: An application conditionally accepts the changes; user breaks link with common context

7.11.3 Abnormal Termination of Common Context Use Case

The Abnormal Termination of Common Context Use Case involves a system administrator forcing the termination of the context manager through some action. The common context participants are notified of the termination of the common context.

Figure 16 illustrates the abnormal termination use case while Interaction Diagram 11 captures an instance of this case.

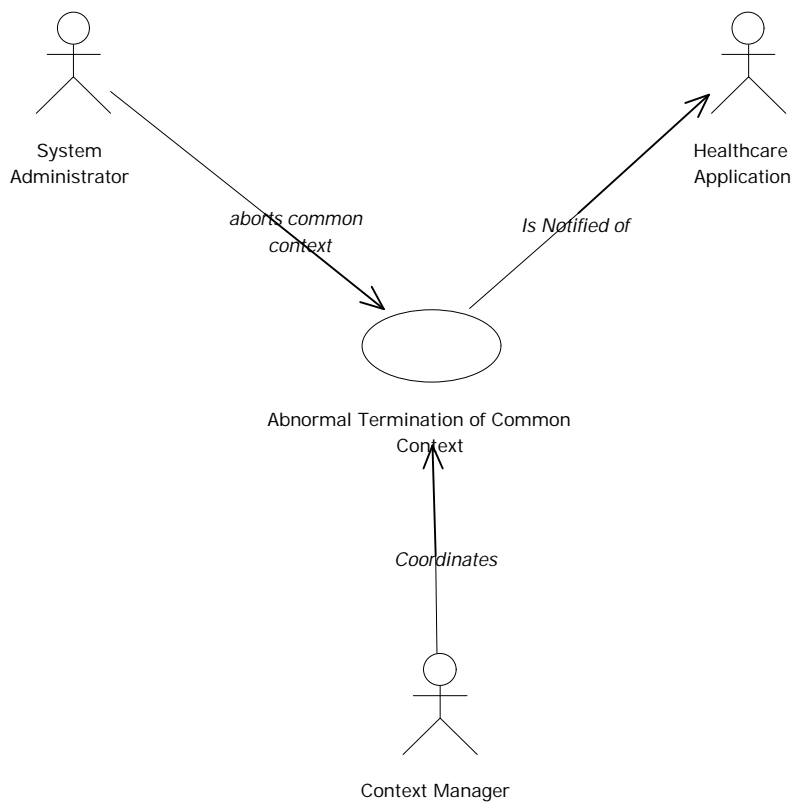
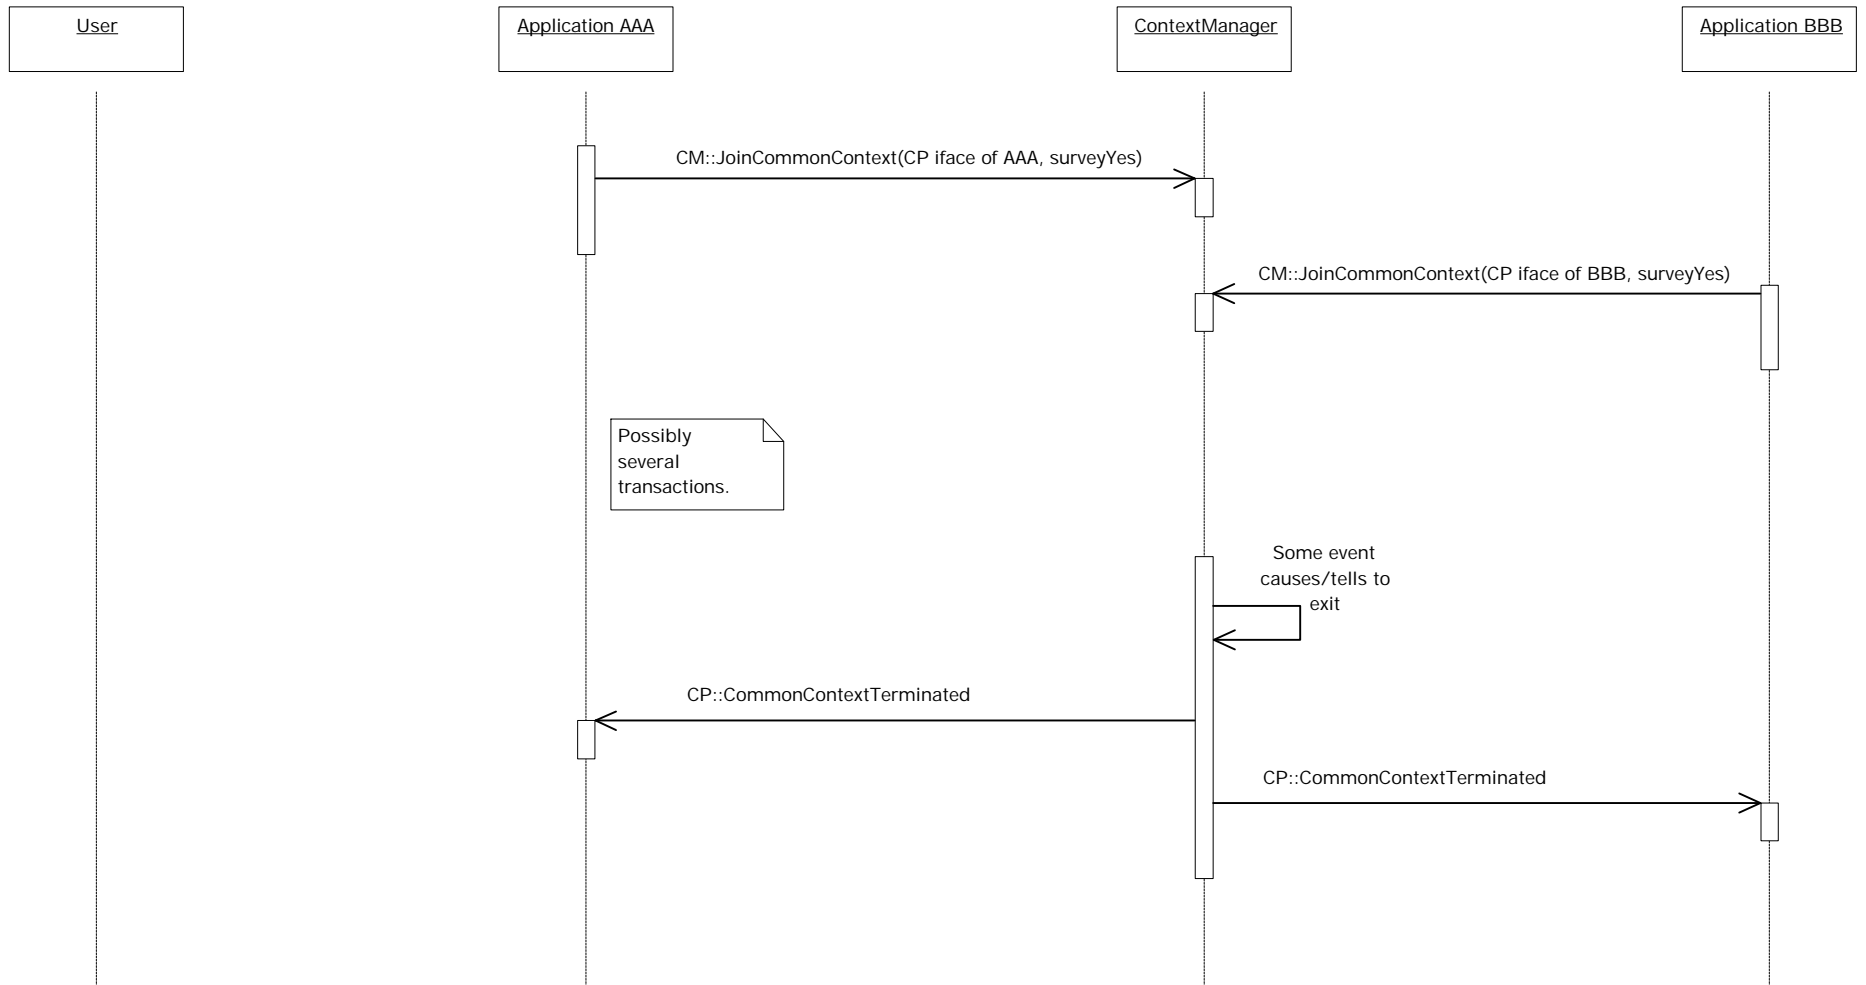


Figure 16: Abnormal Termination of Common Context Use Case



Interaction Diagram 11: Abnormal Termination of Common Context

7.12 Stat Admissions

A stat admission occurs when an application needs to enable the user to record information about a patient even if an identifier for the patient is not known. In this case, the application should indicate to the user that it is breaking its participation in the patient context, and then break its participation upon user confirmation. This is because it is not possible for the application to identify the patient, which is needed in order to change the common context. The only reasonable recourse is for the application to break its participation in the common context.

7.13 Optimizations

There are several optimizations that have been designed into the specification. These optimizations are reflected in the interface specifications described in Chapter 11:

- An application can indicate that it never wants to participate in the survey conducted by the context manager when the context data changes. The context manager will assume that such applications always accept the changes. Read-only data displays represent a class of applications for which this capability is useful.
- An application can selectively suspend its participation in the surveying process without actually leaving the common context. This enables an application to perform computational tasks without being interrupted by context changes. This also enables an application to minimize its use of computational resources if it is in a state (e.g., minimized) in which responding to context changes provides no benefit to the user. The application can subsequently resume its participation in the common context.
- An application can obtain just the context data values that were altered by the most recent change transaction. This capability will become increasingly useful as additional common context data items are defined.
- Multiple common context items can be accessed by an application in a single invocation of a context manager method. This optimizes performance by reducing the number of calls an application needs to make to access context items.
- When an application is notified about a context change, it is also provided with the context coupon value that it needs in order to access the context data. This simplifies the design of applications because they do not necessarily need to keep track of context coupon values.
- Context managers can be implemented to conduct the change survey and the subsequent change notifications in a concurrent manner, thereby decreasing the amount of time it takes to complete these computations.

Additional optimizations, such as enabling applications to indicate their interest in only being notified when specific context data items change are candidates for future enhancements.

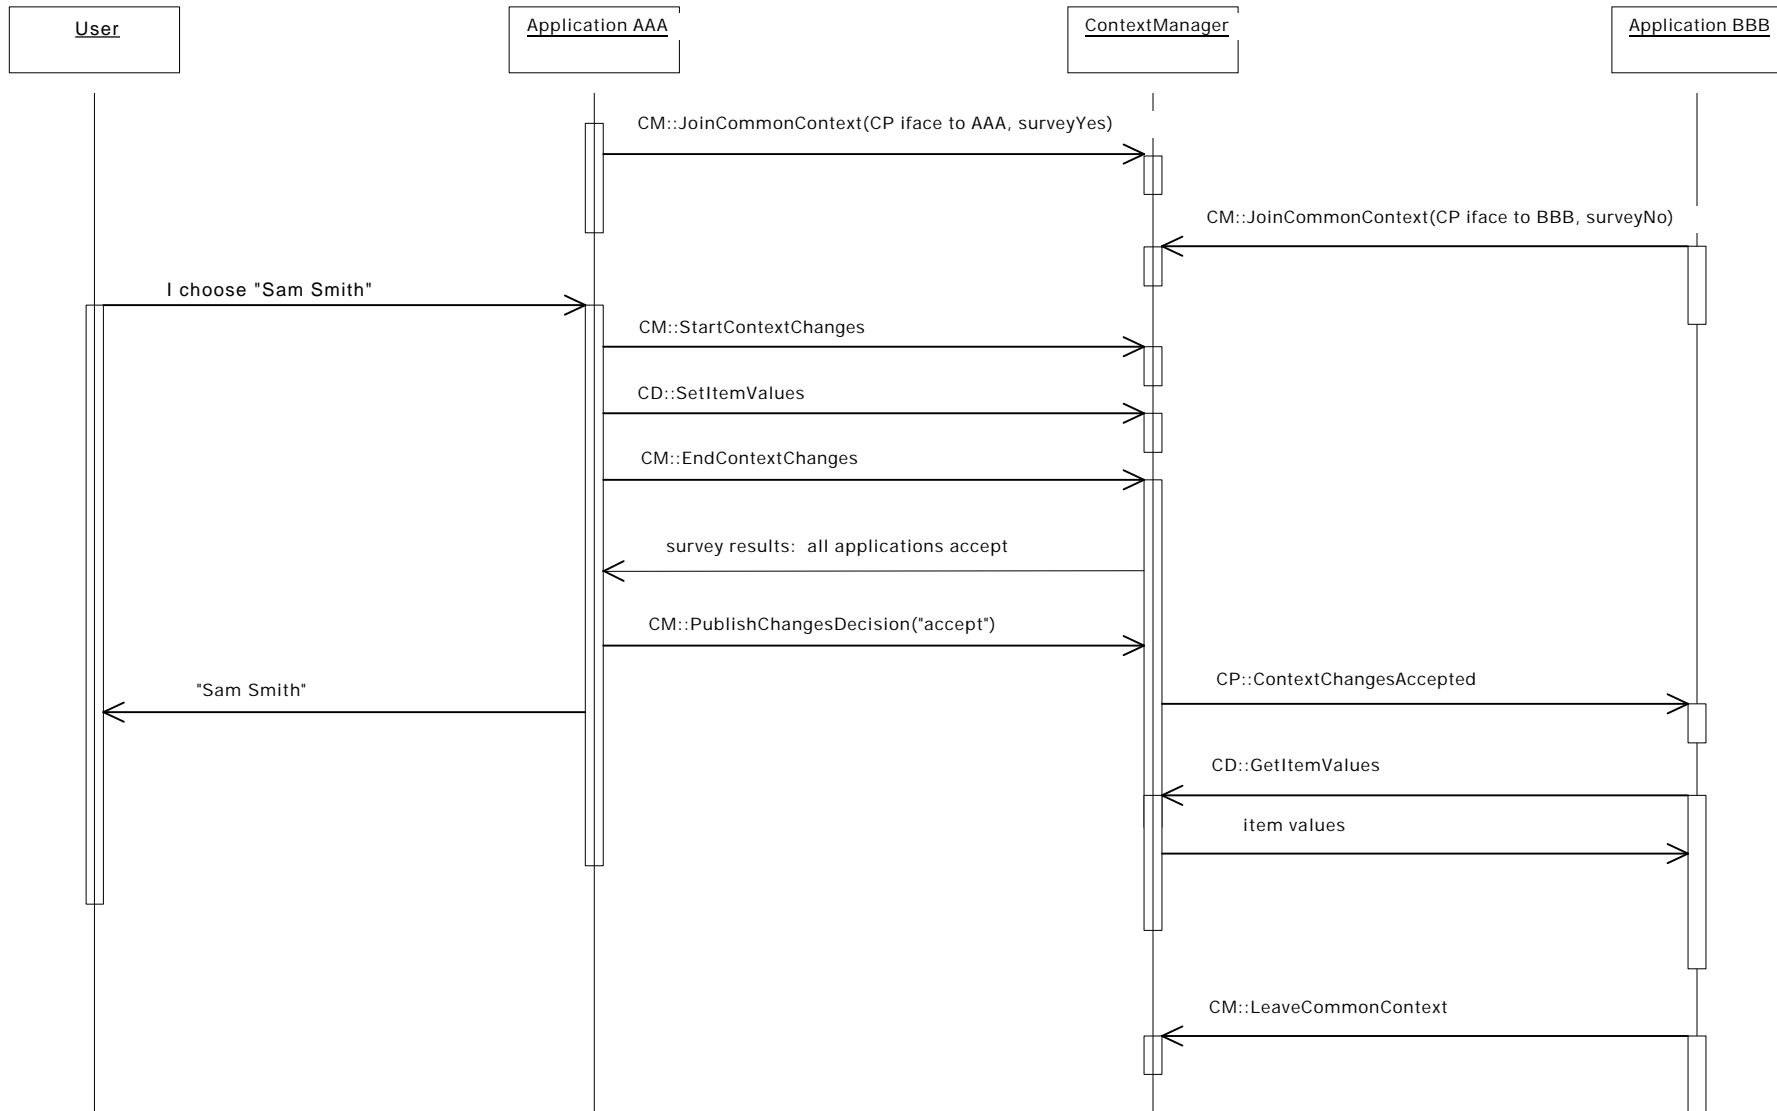
7.14 The Simplest Application

The responsibilities that an application must implement in order to behave properly as a participant in a common context system depends upon the application's functionality. Applications that need to participate in the context change survey must implement straightforward but non-trivial behaviors. However, for many applications it will suffice to implement a very small set of behaviors. Specifically, the simplest participants are those that do not participate in the survey, do not set the context data, and only want to be informed when context changes have been accepted. These applications only need to do the following:

1. Join the common context system via the context manager's ContextManager interface.
2. Implement the ContextParticipant method that enables the application to be informed about accepted context changes.
3. Access the context data via the context manager's ContextData interface.
4. Leave the common context system upon termination, via the context manager's ContextManager interface.

As Interaction Diagram 12 illustrates below, this amounts to implementing one method for ContextParticipant. (The others can be stubbed with trivial default behaviors.) It also requires using two ContextManager methods: one to join and one to leave a common context system. Finally, it requires using one ContextData method to access the context data. The application does not necessarily need to keep track of the value of the context change coupon, as the context manager each time a change occurs provides the correct coupon value to the notified application. The result is that simple applications are not penalized for being co-participants with applications that have more sophisticated needs.

Context Management Specification, Technology and Subject-Independent Component Architecture



Interaction Diagram 12: Simplest Application

8 Mapping Agents

A mapping agent in a common context system provides a means to automatically supply multiple synonymous identifiers for the same real-world entity or concept even when only one identifier is known to the application used to instigate a context change. This mapping is performed in a manner that is transparent to the user and to the applications in the context system.

For example, multiple medical record numbers within a healthcare enterprise might identify a patient. However, each application might only be able to denote a particular patient via just one of these identifiers. When the user selects a patient using such an application, the application sets the new patient context using the patient identifier it knows. The context manager automatically delegates the task of mapping the provided identifier to additional identifiers to a mapping agent. A master patient index system might serve as the basis for implementing a mapping agent capable of mapping patient identifiers.

Mapping agents are not necessarily needed in order to realize a useful and correctly functioning common context system. Specifically, mapping agents are not needed when each real-world entity or concept has a single identifier that is already known to all of the applications in the common context system. For example, there are healthcare enterprises that have a uniform way to identify their patients.

The specification contained in this chapter is for a Patient Link mapping agent. However, other kinds of mapping agents are envisioned for other types of common clinical context data.

Therefore, an attempt has been made to specify the mapping agent in a way that will enable forward compatibility with future CMA capabilities, such as additional context subjects.

8.1 Assumptions and Assertions

It is not an objective of the CMA to define how mapping agents should work or to prescribe or assume a particular mapping agent implementation. Instead, a mapping agent is treated as an abstraction. Interfaces are defined that enable mapping agents to be connected to context managers for the purpose of aiding in the mapping of context identifiers between multiple identifier spaces.

Additional assumptions and assertions include:

- When present, the mapping agent is the authority within a common context system on the mapping between context identifiers.

- 1 • A mapping agent does not allow an identifier to map to more than one real-world
2 entity or concept (e.g., a patient mapping agent does not allow a patient identifier to
3 map to more than one patient).
- 4 • There is at most one mapping agent per context subject per clinical desktop. (Behind
5 the “scenes” mapping agents may work together, or may be implemented using a single
6 common service. However, this is not visible to the context manager or the context
7 participants.)
- 8 • A context manager does not know about the mapping agent implementation; a context
9 manager only “sees” a mapping agent through its CMA-defined interface.
- 10 • Context participant applications do not “know” about the mapping agent (or even if
11 there is one); the mapping agent does not “know” about context participant
12 applications.
- 13 • The mapping agent may reside on a computer that is remote from the computer (s)
14 upon which the context manager(s) they serve reside; however, these computers must
15 be connected by a LAN or WAN whose performance is LAN-equivalent.
- 16 • Mapping agents are an optional component of a CMA context management system.

17 **8.2 Interfaces**

18 The following interfaces are defined for and implemented by mapping agents:

- 19 • MappingAgent (MA) - used by a context manager to inform a mapping agent that the
20 clinical context has changes pending and that the mapping agent should perform its
21 context data mapping responsibilities
- 22 • ImplementationInformation (II) - used by a context manager to obtain details about
23 who implemented the mapping agent, when it was installed, etc., for the purpose of
24 creating detailed error reports

25 In addition, mapping agents to set/get context data items uses the context manager
26 ContextData interface.

27 The mapping agent interfaces are modeled and illustrated in Figure 11: Patient Link
28 Component Architecture.

8.3 Theory of Operation

Assume, first, that one or more context participants have already joined the same common context and that they are connected to the context manager. Further, assume that the context manager already has an interface reference to a mapping agent's MappingAgent interface. How these references are obtained is described in Section 8.3.1, Initializing a Context System When a Mapping Agent is Present.

Given these conditions, a context participant instigates a context change transaction via the context manager's ContextManager interface, sets the new context data via context manager's ContextData interface, and then indicates it is done setting the data via the context manager's ContextManager interface.

At this point, before the other context participants are surveyed, the manager informs the mapping agent that the context data has changes pending, via the mapping agent's MappingAgent interface (which is similar to an application's ContextParticipant interface). The mapping agent blocks the context manager's method return until the mapping agent has completed its mapping tasks. The proposed context data items that are available to the mapping agent are exactly as the instigating participant set them.

The mapping agent reads the proposed context data via the context manager's ContextData interface, and may set one or more *additional* context data identifier or corroborating items via this same interface. The objective is for the mapping agent to *enhance* the proposed context by providing *additional* identifier or corroborating data in a manner that is transparent to the application that instigated the transaction.

Applications (including the instigating application) are not allowed to set context item values after the instigating application has completed its changes. However, the context manager allows the mapping agents to make changes because it knows it is a mapping agent that is setting the item values. How the context manager knows that it is a mapping agent will be described later.

Once the mapping agent has completed its mapping tasks, the context manager surveys the context participants and processing of the context change transaction is performed as usual. With this approach, all of the synonymous values for an identifier will be set before the other applications are informed via a context manager-initiated survey that the context has been changed.

However, if the instigating application has set multiple values for a context identifier, and the mapping agent detects an inconsistency among these values, then it informs the context manager that the context change transaction has been invalidated. This is because the mapping agent is the authority in a context system when it comes to mappings between identifiers. Allowing the transaction to proceed could create confusion about the context among the other context participants.

1 The details about the conditions under which a mapping agent can invalidate a context change
2 transaction are described in 8.3.5 Conditions for Mapping Agent Invalidation of Context
3 Changes.

4 When the mapping agent invalidates a context change transaction, the context manager does
5 not survey the participating applications. Instead, the context manager informs the instigating
6 application that the transaction has been invalidated. The instigating application then asks the
7 user to intervene to decide how to proceed.

8 The user can decide (via a dialog presented by the application that was used to instigate the
9 context change) whether to cancel the context change or to break the instigating application
10 away from the common context system. In either case, the context change transaction is
11 terminated and the context changes are discarded. Additional identifiers are not mapped and
12 the other applications are not surveyed.

13 This approach gives the user the option of applying the context changes to just the application
14 used to instigate the context change while also preventing the other applications from becoming
15 confused about the context.

16 The details of this situation are described in 8.3.6 Treatment of Mapping Agent Invalidation of
17 Context Changes.

18 **8.3.1 Initializing a Context System When a Mapping Agent is Present**

19 A mapping agent and the context manager it serves must be connected to each other. There are
20 two ways in which this can be accomplished. Either the context manager connects to the
21 mapping agent, or the mapping agent connects to the context manager. The order in which this
22 connection occurs has significant impact on complexity and computing resource utilization.

23 The mapping agent could conceivably locate and connect to a context manager the same way a
24 context participant does. This requires that the connection be made *before* the first time a
25 context participant application sets the context. This is so that the mapping agent can be
26 instructed by the context manager to perform its mapping tasks.

27 A consequence of this approach is that a context manager will execute even if it is not actively
28 servicing any context participants. Further, the requirement that the connection be made *before*
29 the first time a context participant application sets the context introduces initialization-
30 sequencing complexities.

31 In general there is no way to know when the first context participant will connect to a context
32 manager, so the only prudent recourse would be to launch the context manager and the
33 mapping agent as part of the boot-up process for the desktop they serve. This would
34 complicate the installation process for context managers and mapping agents.

1 The alternative is for the context manager to connect to the mapping agent. This approach
2 enables the connection to be deferred until the mapping agent is needed to service a context
3 participant. However, a means by which context managers can locate the necessary mapping
4 agent must be established.

5 Fortunately, the fact that there is only one mapping agent per context subject per clinical
6 desktop enables the location process to be easily implemented using the desktop's technology-
7 specific desktop interface reference registry. Specifically, a reference to a mapping agent's
8 principal interface is entered into the desktop's interface reference registry. The symbolic name
9 and/or description of the interface within the registry indicates the context subject that the
10 mapping agent maps. The context manager obtains this reference and uses it to interrogate the
11 mapping agent to obtain references to its other interfaces, such as MappingAgent.

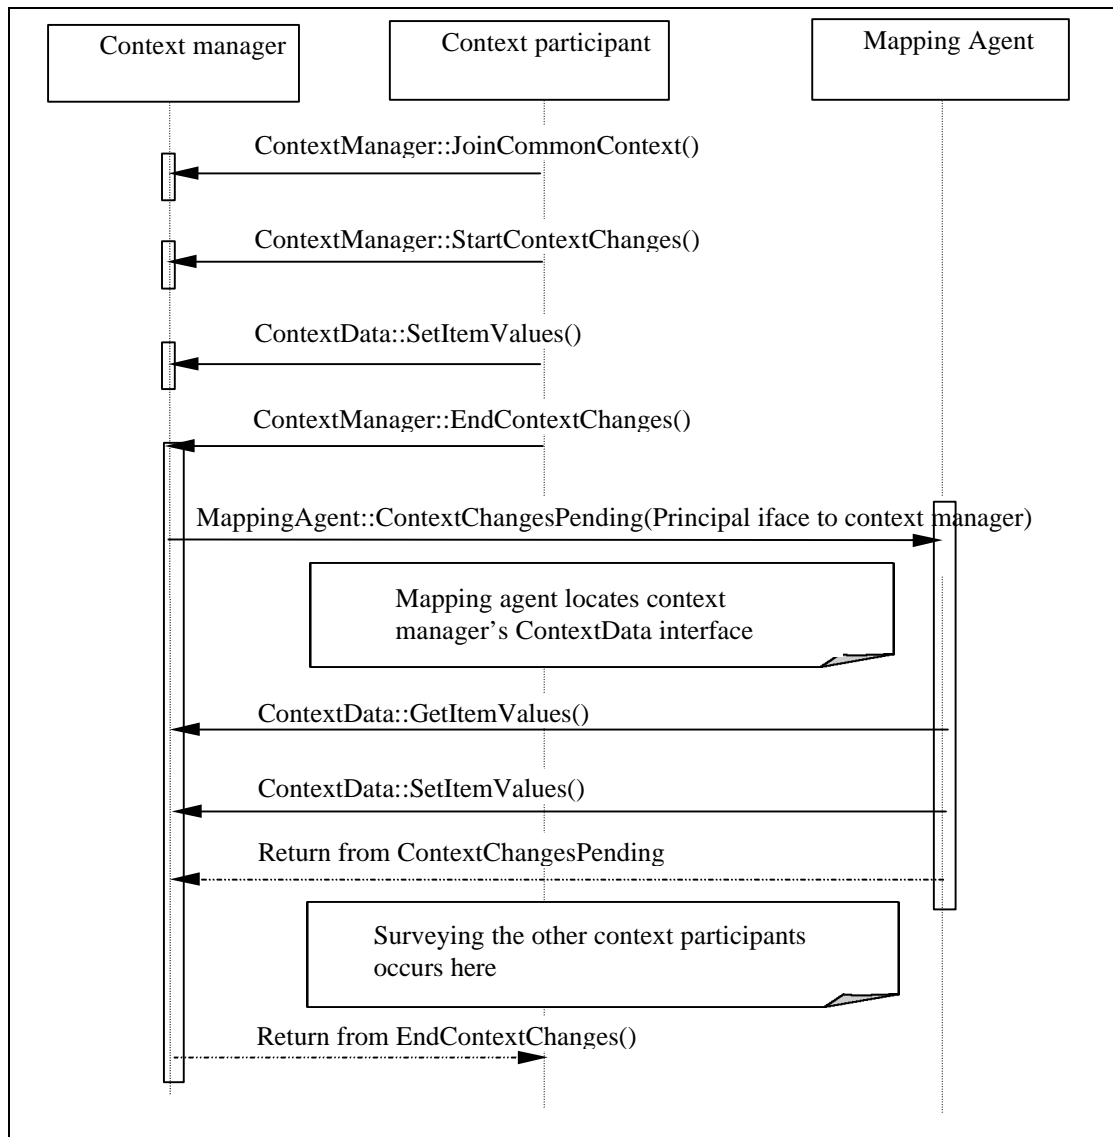
12 An additional benefit of the manager-connects-with-agent approach is that it is not even
13 necessary for distinct connect/disconnect methods to be defined. Instead, the context manager
14 simply informs the mapping agent whenever the context manager has changes pending. The
15 context manager explicitly provides a reference to its principal interface to the mapping agent.
16 The mapping agent then interrogates the context manager via its principal interface to obtain a
17 reference to other context manager agent interfaces, such as the interface ContextData.

18 The sequence of events is shown in Interaction Diagram 13: Context Change Transaction with
19 Mapping Agent.

20 **8.3.2 Terminating a Context System When a Mapping Agent is Present**

21 To enable the orderly termination of the context system, the context manager shall implicitly or
22 explicitly dispose of any mapping agent interface references that it possesses prior to
23 terminating. The mapping agent shall dispose of any context manager interface references that
24 it possesses when it has completed its mapping actions for a context change transaction. The
25 means by which these disposals are effected is technology-specific.

26 The consequence of these disposals is that at the end of a context change transaction, only
27 context participant applications will possess context manager interface references. If there are
28 no participants, then the context manager can properly terminate. (Participants dispose of any
29 context manager interface references that they possess prior to terminating. See Section 6.1.5,
30 Interface Reference Management.) This also means that once the context manager terminates,
31 the mapping agent can also properly terminate.



Interaction Diagram 13: Context Change Transaction with Mapping Agent

8.3.3 Distinguishing Between Mapping Agents and Context Participants

When a mapping agent is informed that a context change is pending, the context manager provides it with two coupons. One coupon denotes the context change transaction; the other denotes the mapping agent. The mapping agent coupon is not the same as any of the coupons assigned by the context manager to the context participants.

The mapping agent shall use the coupon that denotes it whenever it sets context data via the ContextData interface. The context manager uses this coupon to determine that a mapping agent, and not a context participant, is setting the context data. Only a mapping agent is

1 allowed to set context data after the instigator of the context change has indicated that it has
2 completed the context changes.

3 **8.3.4 Mapping Agent Updates to Context Data**

4 A mapping agent only adds data to the context. A mapping agent can add additional context
5 identifier items. It can also add additional corroborating data items. These updates are
6 primarily for the benefit of the context participants other than the application that instigated the
7 context change.

8 This is because it cannot be assumed that the instigating application will re-read the context
9 data once it has completed its context changes. In contrast, the other applications do not read
10 the new context until they are surveyed, which occurs after the mapping agent has added data
11 to the context.

12 If a mapping agent was allowed to change the values for context items that have been set by
13 the instigating application, it could be confusing to the user. This is because the user might see
14 differences between the context data as displayed by the instigating application and as
15 displayed by the other context participant applications.

16 Given this concern, a mapping agent shall not alter the values of any of the context data items
17 that have already been set by the instigating participant as part of the proposed context. Any
18 attempt to alter existing context data items by the mapping agent shall result in the context
19 manager raising an exception.

20 A mapping agent shall not delete any of the context data items. Any attempt to delete context
21 data items by the mapping agent shall result in the context manager raising an exception.

22 **8.3.5 Conditions for Mapping Agent Invalidation of Context Changes**

23 A context subject is comprised of multiple identifier and corroborating data items, each of
24 which is represented as name/value pairs (see Section 5.4, Context Data Representation, and
25 Section 5.6, Context Data Interpretation). It is the responsibility of every application that sets
26 these items to ensure that they are self-consistent. However, there are a variety of potential
27 item name and/or item value inconsistencies that a mapping agent must be able to detect.

28 Specifically, if an application has set multiple values for a context identifier item, and the
29 mapping agent determines that these values *do not* all identify the same real-world entity or
30 concept (e.g., patient), the mapping agent shall invalidate the context change transaction.

31 Specifically, a mapping agent shall invalidate a context change transaction when:

- The instigating application sets more than one value for the same context identifier item, but the mapping agent determines that at least two of these values identify different patients.
- The instigating application sets more than one value for the same context identifier item, but the mapping agent knows that at least one of these values conflicts with a value known to identify the patient.

There are situations in which the mapping agent must not invalidate a context change transaction even though there are apparent context item inconsistencies. A mapping agent must not flag what it believes to be inconsistencies when in fact the suspect items might represent reasonable application behaviors.

The following scenarios illustrate the desired mapping agent behaviors. Assume that there are two patients, each with identifiers for two sites, and the mapping agent is able to map the patient identifiers for both sites:

	Patients and Their Site-Specific Identifiers	
Institution	John Doe	Jim Smith
St. Elsewhere Hospital	123-456-789Q36	155-213-424Y82
St. Elsewhere Clinic	2888-91922-W928	18291-81293-D812

The first two scenarios represent inconsistencies that the mapping agent must respond by invalidating the context change transaction. The last three scenarios represent inconsistencies that the mapping agent must ignore:

	What the instigating application does ...	Example ...	What the mapping agent does ...
1	Sets two identifier values, both with the intent of denoting John Doe, but the values erroneously denote John Doe and Jim Smith.	<i>Item identifies John Doe:</i> [Patient.Id.MRN.St_Elsewhere_Hospital, 123-456-789Q36] <i>Item erroneously identifies Jim Smith:</i> [Patient.Id.MRN.St_Elsewhere_Clinic, 18291-81293-D812]	Invalidates the context change transaction because the first identifier value denotes John Doe, while the second denotes Jim Smith. Mapping is not performed.
2	Sets more than one identifier pair, both with the intent of denoting John Doe. The first value is John Doe's hospital identifier, but the second value is not John Doe's clinic identifier.	<i>Item identifies John Doe:</i> [Patient.Id.MRN.St_Elsewhere_Hospital, 123-456-789Q36] <i>Item does not identify John Doe:</i> [Patient.Id.MRN.St_Elsewhere_Clinic, 0000-00000-0000]	Invalidates the context change transaction because while the first identifier value is John Doe's hospital identifier, the second value is known not to be John Doe's clinic identifier. Mapping is not performed.
3	Sets only one context identifier item and the name of the item is not known to the mapping agent.	<i>Item name not known to mapping agent:</i> [Patient.Id.MRN.General_Hospital, 6668-3923-987122]	Ignores this situation and does not inform the context manager about inconsistencies. Mapping is not performed.
4	Sets more than one value for a context identifier item, and one or more of the item names are not known to the mapping agent.	<i>Item name known to mapping agent:</i> [Patient.Id.MRN.St_Elsewhere_Hospital, 123-456-789Q36] <i>Item name not known to mapping agent:</i> [Patient.Id.MRN.General_Hospital, 6668-3923-987122]	Ignores this situation and does not inform the context manager about inconsistencies. Mapping is performed.
5	Sets the corroborating data to values that are different (or incomplete) as compared to the corroborating data known to the mapping agent	Application sets corroborating data containing the identified patient's name to "Jack Doe" but mapping agent knows the identified patient as "John Doe".	Ignores this situation and does not inform the context manager about inconsistencies. Mapping is performed.

1

2 In summary, detectable inconsistencies between identifier values are the only reason that a
3 mapping agent should invalidate a transaction. Transactions must not be invalidated when
4 unknown identifier names are used by an application or because of corroborating data
5 inconsistencies.

6 8.3.6 Treatment of Mapping Agent Invalidation of Context Changes

7 Applications that instigate context change transactions and then explicitly set more than one
8 identifier during a context change transaction shall explicitly handle the situation in which a
9 mapping agent invalidates a context change transaction. (Applications that set only one
10 identifier do not need to handle this situation.)

11 An instigating application is not provided with a means to distinguish between the invalidation
12 of a context change transaction and the presence of a busy application. These are clearly

different situations, but are to be handled by an instigating application in the same way. The application shall present a dialog that clearly indicates that a problem has been encountered while attempting to change the common context.

The dialog shall include a description of the problem that was encountered. The dialog shall also enable the user to cancel the context change or to break the link between the instigating applications and the other applications.

When the mapping agent has invalidated a transaction it shall not be possible for the user to force a common context change. If the user decides to break the link between the instigating application and the other applications, instigating application shall only apply the context change to itself. This application shall break away from the common context and shall clearly indicate to the user that it is not participating in the common context.

If the user cancels the context change, then the instigating application shall indicate this fact to the context manager. Both the instigating application and the context manager shall discard the current transaction. The context manager shall not survey the other applications.

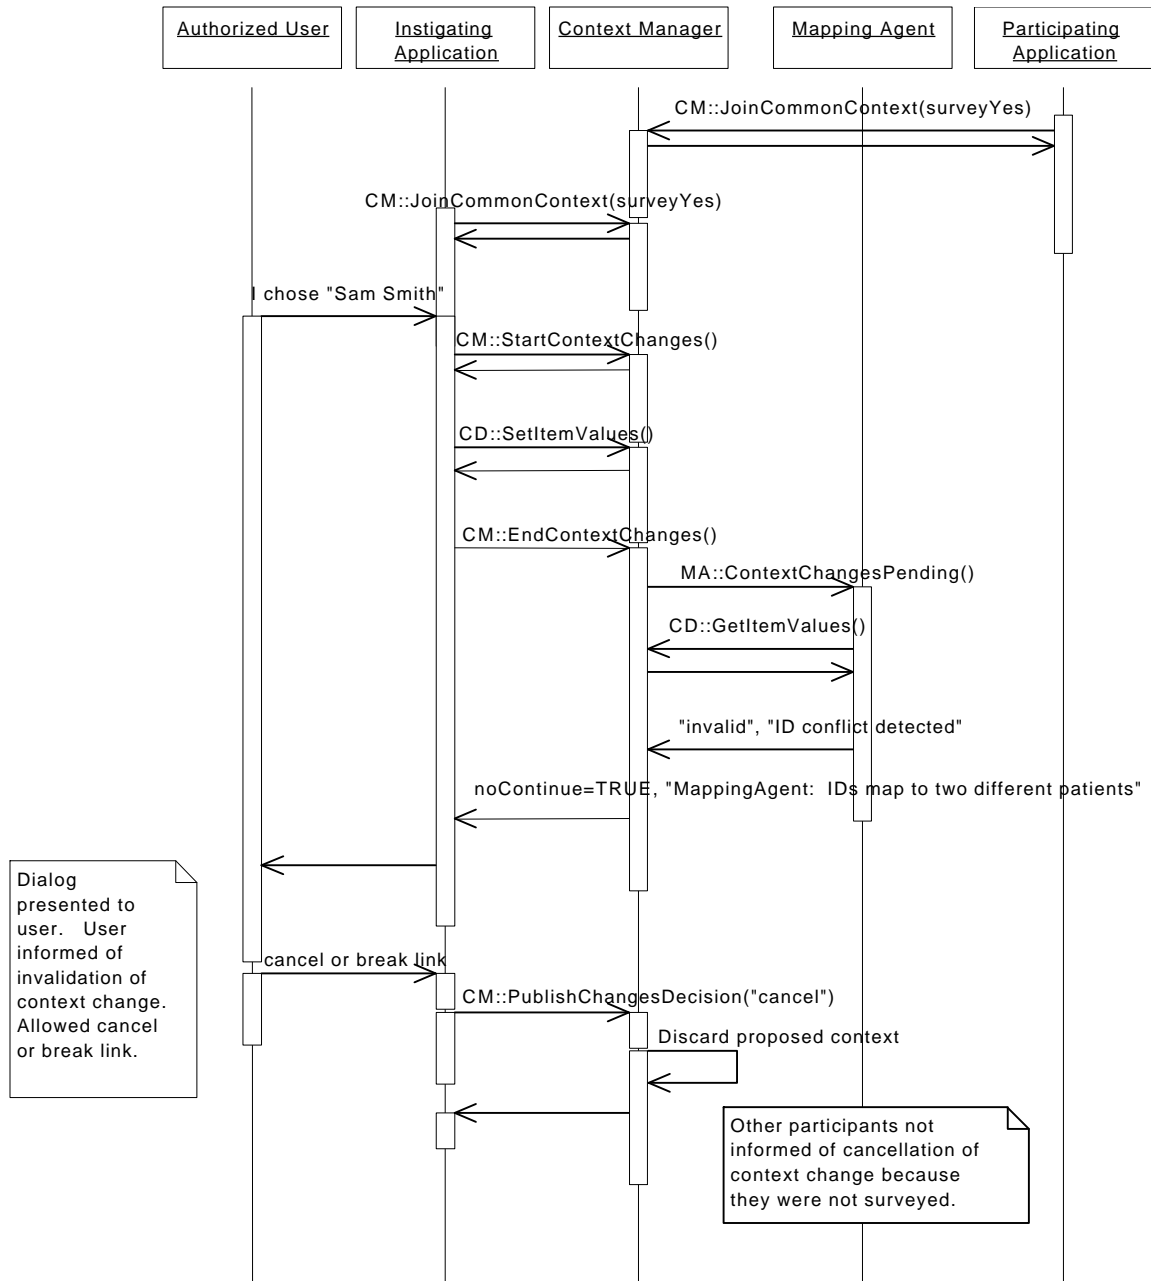
Independent of the reason for which the mapping agent invalidated the transaction, the context manager shall always provide to the instigating application the same user-friendly description of the problem that was encountered. This is in order to keep things simple for the user, who is unlikely to be concerned about the details of what went wrong. This description shall be included in the dialog by the instigating application.

The appearance of the dialog and the commands that the user can choose from are specified in each of the HL7 context management technology-specific user interface specification documents. The wording for the user-friendly description that is included in the dialog is also specified in these documents. This will ensure a consistent and familiar set of interactions for users across CMA-conformant applications.

The sequence of events that occur when a mapping agent invalidates a context change transaction is shown in Interaction Diagram 14: Mapping Agent *Invalidates* Context Change Transaction.

8.3.7 Mapping Null-Valued Identifiers

A mapping agent shall not perform any mapping when the context subject is empty (See Section 5.6.8, Representing an Empty Context). The net effect is that the context subject remains empty, and all of the applications see the context as such.



Interaction Diagram 14: Mapping Agent Invalidates Context Change Transaction

8.3.8 Initializing Mapping Agents

Different mapping agent implementations may require different initialization methods. For example, a mapping agent might need to authenticate the current user in order to enforce security policies. Other than being automatically launched by a context manager, the additional steps needed to initialize a mapping agent are implementation issues and are not addressed by

1 this specification. (Future versions of the CMA specification may provide standardized ways
2 of initializing mapping agents.)

3 It can be the case that different mapping agent implementations will require different explicit or
4 implicit actions on the part of the user to complete their initialization tasks. An example of an
5 explicit user action is signing on to the mapping agent via a mapping agent-supplied dialog. An
6 example of an implicit user action is signing on to a context participant application that relays
7 its authentication of the user to the mapping agent; this obviously implies a relationship with
8 the mapping agent that goes beyond this specification.

9 **8.3.9 Handling Mapping Agent Failures**

10 A context manager must be able to detect and handle the failure of a mapping agent.
11 Specifically, a context manager shall behave in a robust manner even if its calls to a mapping
12 agent's MappingAgent interface do not return in a timely manner.

13 The recourse, after a timeout has occurred, is for the context manager to continue with the
14 normal processing of the context change transaction. If the mapping agent has indeed failed,
15 then some of the context participants may not be able to interpret the next context. However,
16 this fail-soft approach still enables the user to perform useful work until the mapping agent
17 failure is corrected.

18 Finally, even if a mapping agent has failed, a context manager shall continue to try to access
19 the mapping agent during subsequent transactions on the prospect that the failure has been
20 corrected. In doing so, the context manager may need to obtain a new interface reference for
21 the mapping agent (because the old reference may no longer be valid).

22 Note that this policy of continually attempting to access a failed mapping agent also applies
23 even when a context manager is first launched. It may be the case that a mapping agent
24 becomes available after the context manager has begun executing. (See Section 8.3.8,
25 Initializing Mapping Agents, for one explanation of why this might happen.) A context
26 manager that does not locate and initiate a mapping agent when it is launched shall
27 nevertheless keep trying between and/or during context change transactions. It is an
28 implementation decision as to how the performance impact of this policy is minimized.

29 **8.4 Mapping Agent Effect on Application Security Policies**

30 Mapping agents may implement their own security policies in terms of what context data it will
31 map for a particular user. Mapping agent security policies can differ from the policies of the
32 participating applications. A mapping agent's policies might effect what patients a user can, or
33 cannot, access.

When the mapping agent's policy is more restrictive than one or more of the participating application's, a mapping agent might elect to *not* map an identifier because doing so would violate the security rules known to the mapping agent. When the mapping agent's policy is less restrictive than one or more of the participating applications, each application's own security policy will be the predominating policy for the current change transaction.

A mapping agent that elects to *not* map an identifier because of security concerns shall not indicate this fact to the user. The user will simply observe that access to the selected patient is not possible through one or more of the participating applications. These applications do not know that the identifier for the selected patient has not been mapped because of the mapping agent's security policy. Instead, it looks to the applications as though a patient has been selected but the identifier(s) by which the patient is known to the applications has not been provided. These applications behave as specified for in 6.5.1 Application Behavior When it Cannot Cancel Context Changes.

8.5 Identifying Mapping Agent Implementations

Context managers use a mapping agent's ImplementationInformation interface to provide system administrators with a description of the mapping agent implementation it is using. This information can help system administrators diagnose run-time problems that involve mapping agents.

The ImplementationInformation interface shall be supported by all mapping agent implementations. A context manager shall not interact with a mapping agent that does not support this interface.

8.6 Performance Costs and Optimizations

When present, a mapping agent will be involved in every context change transaction. This adds an overhead to the context change transaction in the form of the added communication between the context manager and the mapping agent, and for the time it takes for the mapping agent to validate the identifiers and provide any additional mappings for the identifiers. However, these costs are viewed as being worth the benefits of the semantic integrity that a mapping agent brings to a context system.

In some cases, a mapping agent will be implemented using an underlying application that provides its own user interface for patient selection. This type of mapping agent is, in effect, both a mapping agent and a context participant application. In the case in which this underlying application is used to instigate a context change, performing identifier validations and mappings is superfluous. It is possible to optimize the mapping agent implementation so that it does not perform identifier validations and mappings when it knows that it was essentially itself that instigated a context change.

1 However, the only information that is readily available to the mapping agent that could help it
2 determine this fact is the context change coupon. This coupon is provided by the context
3 manager to an application when the application starts a context change transaction. This
4 coupon is also provided by the context manager to the mapping agent via its MappingAgent
5 interface during each context change transaction.

6 It is an implementation decision as to how the portion of an application that implements a
7 mapping agent obtains the value of the context coupon from the portion of the application that
8 instigates a context change transaction.

9

10

9 User Link Theory of Operation

This chapter describes CMA support for User Link. With User Link, a user can securely sign on to any User Link-enabled application on a desktop using just one logon name and one means of authentication (such as a password) in order to securely sign on to all User Link-enabled applications on the desktop.

User Link extends CMA support for Patient Link in several ways:

- It introduces another context subject. Managing multiple subjects requires additional context management policies beyond those defined for Patient Link.
- It introduces the user subject as the second foundational CMA context subject.
- It introduces security capabilities that not only enable the creation of secure User Link context management systems, but that also serve as a foundation for future subjects that require security.

In order to accomplish this, the Patient Link architectural approach is leveraged (i.e., context manager, context participants, and mapping agent) to create a single context per desktop. The context is extended to include the user subject in addition to the patient subject.

The Patient Link interfaces ContextManager, ContextParticipant, MappingAgent, and ImplementationInformation interfaces are used. Two additional security-related interfaces are defined: SecureContextData, which is modeled upon the Patient Link ContextData interface, and SecureBinding, which enables a trusted relationship to be established between User Link-enabled applications and components.

Additional User Link capabilities include:

- The provider institution decides which applications are to be trusted to authenticate users.
- There can be multiple ways to authenticate users, including passwords, biometrics, etc.
- In keeping with the CMA philosophy, the User Link approach is conceived for low re-engineering costs.

The architecture that supports these capabilities is described next.

9.1 *User Link Terms*

The following terms are used to describe the User Link theory of operation:

- **User Link-enabled application** - an application that implements the CMA User Link capability.
- **Sign on** – the act of identifying oneself to an application, prior to initiating a user session, in a manner that can be authenticated by the application, typically involving a secret password or a biometric reading (such as a thumb-print scan).
- **Log-off** – the termination of a user’s session with an application.
- **Empty context** – a context is not defined for a particular subject, either because no context identifier items are present in the context data (as is the case when a context manager is first initialized) or because the values of all of the identifier items for the subject that are present in the context data are *null* (as is the case when an application explicitly indicates that the context is empty).

9.2 *Desktop Assumptions*

The following assumptions are made about the clinical desktop upon which User Link-enabled applications are deployed:

- Logging-off from an application does not require user authentication.
- The desktops upon which User Link-enabled applications are deployed may reside in physically unsecured locations.
- While recommended, it may not be the case that appropriate security precautions have been taken to restrict the types of operating system-level actions, such as installing new programs, that users can perform on desktops that reside in physically unsecured locations.

In summary, the CMA is intended to be no less secure than the User Linked applications would be were they not User Linked. In general, User Linked applications will be substantially more secure.

9.3 *User Subject*

The context subject of *User* is defined for User Link. The context data identifier item for this subject is the user’s logon name. A logon name denotes a user to an application. A user’s logon name is generally different from their given name.

This identifier is unlikely to be universally unique. However, it is assumed that a population of users across which each logon name is unique can be established. Each such population is referred to as an *application*, as it is typical that within an overall healthcare institution each population of users corresponds to a particular application.

Consequently, a single user may be identified using multiple user subject identifier items. Each item is differentiated by a different application-specific suffix. An application shall be configurable such that it can be instructed on-site as to which suffix (or suffices) it is to use when it interacts with the context manager to set or get user context data.

The format of a user subject identifier item name includes an application-specific suffix. Use of this suffix, and the values that may be assigned to this suffix, is at the discretion of each healthcare institution at which a context management system is deployed.

In addition to identifier items, the user subject also supports corroborating data items. The actual names, meaning, and data types used to represent the values for both user subject identifier items and corroborating data items are defined in the document *Health Level-Seven Standard Context Management Specification, Data Definition: User Subject*.

An example of a user subject identifier item appears below:

User Subject Identifier		
Example Item Name Format:	Example Item Name:	Example Item Value:
User.Id.Logon.application_name	User.Id.Logon.3M_Clinical_Workstation	robs

9.4 User Authentication Data Is Not Part of the User Context

The data used to authenticate a user is *not* included as part of the user context data. This data is typically a password, but it can be any data that is used to authenticate a user, such as a biometric sample. Instead, each application is expected to be able to sign on a user given just the application-specific logon name for the user.

This approach substantially reduces security risks because the data used by an application to authenticate the user remains private to the application. If this data were part of the user context, it would be vulnerable to undesired access. However, in order for applications to tune to the user context, they must trust that the context data is authentic. The means by which this is accomplished is referred to as the “chain of trust” and is described below.

9.5 *User Link Common Context System Description*

Consistent with the CMA, on each desktop there are applications that are user context participants, and there is a context manager. The applications perform context change transactions to indicate who the user is.

However, in contrast to the way in which patient context is communicated in a Patient Link system, the user context is communicated throughout the common context system in a secure manner. This is to prevent people from accidentally or maliciously gaining access to applications that are User Linked.

The necessary security is achieved by adding capabilities to the CMA that enable the realization of a “chain of trust” among the User Link-enabled applications and User Link components. With the chain of trust, User Link-enabled applications and User Link components work together to ensure that only authorized users are allowed access to a common context system.

The chain of trust not only simplifies the overall solution, but results in a system that is more secure than would be the case if authentication data were part of the common context, and were therefore vulnerable to security attacks directed against the context manager or mapping agent.

The chain of trust is specified in Chapter 10.

9.5.1 *User Mapping Agent*

An optional user mapping agent is also part of the common context system. The user mapping agent maps the logon names for users. The user mapping agent is similar to, but distinct from, the patient mapping agent (although a single mapping agent implementation could fulfill both roles).

Whenever an application sets the user context, the context manager instructs the user mapping agent (if present) to provide any additional logon names it knows for the user. The application suffix for each of the mapped identifier items denotes the application for which the mapped logon name is valid, for example:

Examples Item Names:

Example Item Values:

User.Id.Logon.3M_Clinical_Workstation	robs
User.Id.Logon.Medicalogic_Logician	rob_seliger
User.Id.Logon.HP_CareVue	r_seliger

9.5.2 Context Management Interfaces

The context management interfaces defined for User Link are similar to the ones defined for Patient Link. A context participant still implements ContextParticipant (CP). The context manager still implements ContextManager (CM), but it also implements the following new interfaces:

- SecureContextData (SD) - Similar to the ContextData interface defined for Patient Link, this interface is used by applications to securely set/get the values for the items (logically represented as name-value pairs) that comprise the clinical context.
- SecureBinding (SB) - Used by applications to establish a secure communications binding with the context manager before using the SecureContextData interface.
- ImplementationInformation (II) – Originally defined for the patient mapping agent, this interface is added to the context manager so that applications, other components, and tools, can obtain details about the context manager implementation, including its revision, when it was installed, etc.

The interfaces implemented by the user mapping agent are MappingAgent (MA) and ImplementationInformation (II). These are the same interfaces as defined for the patient mapping agent.

9.5.3 Authentication Repository

In order to make it practical to re-engineer existing applications to support the chain of trust, the CMA authentication repository component is defined. This repository enables applications to securely store and retrieve application-specific user authentication data. The repository is used by applications that do not have a built-in means to easily sign on a user given only a logon name.

The authentication repository implements the following interfaces:

- AuthenticationRepository (AR) - Used by applications to securely interact with the repository to store and retrieve user authentication data.
- SecureBinding (SB) – Used by applications to establish a secure communications binding with the repository before using the AuthenticationRepository interface. This is the same interface that the context manager implements.
- ImplementationInformation (II) – Originally defined for the patient mapping agent, this interface is added to the authentication repository so that applications, other components, and tools, can obtain details about the authentication repository, including its revision, when it was installed, etc.

9.5.4 Overall User Link Component Architecture

The overall User Link architecture (including the Patient Link Architecture) is illustrated in Figure 17: User Link Component Architecture. (A description for how to interpret the notation used in this diagram appears in the Appendix: Diagramming Conventions.)

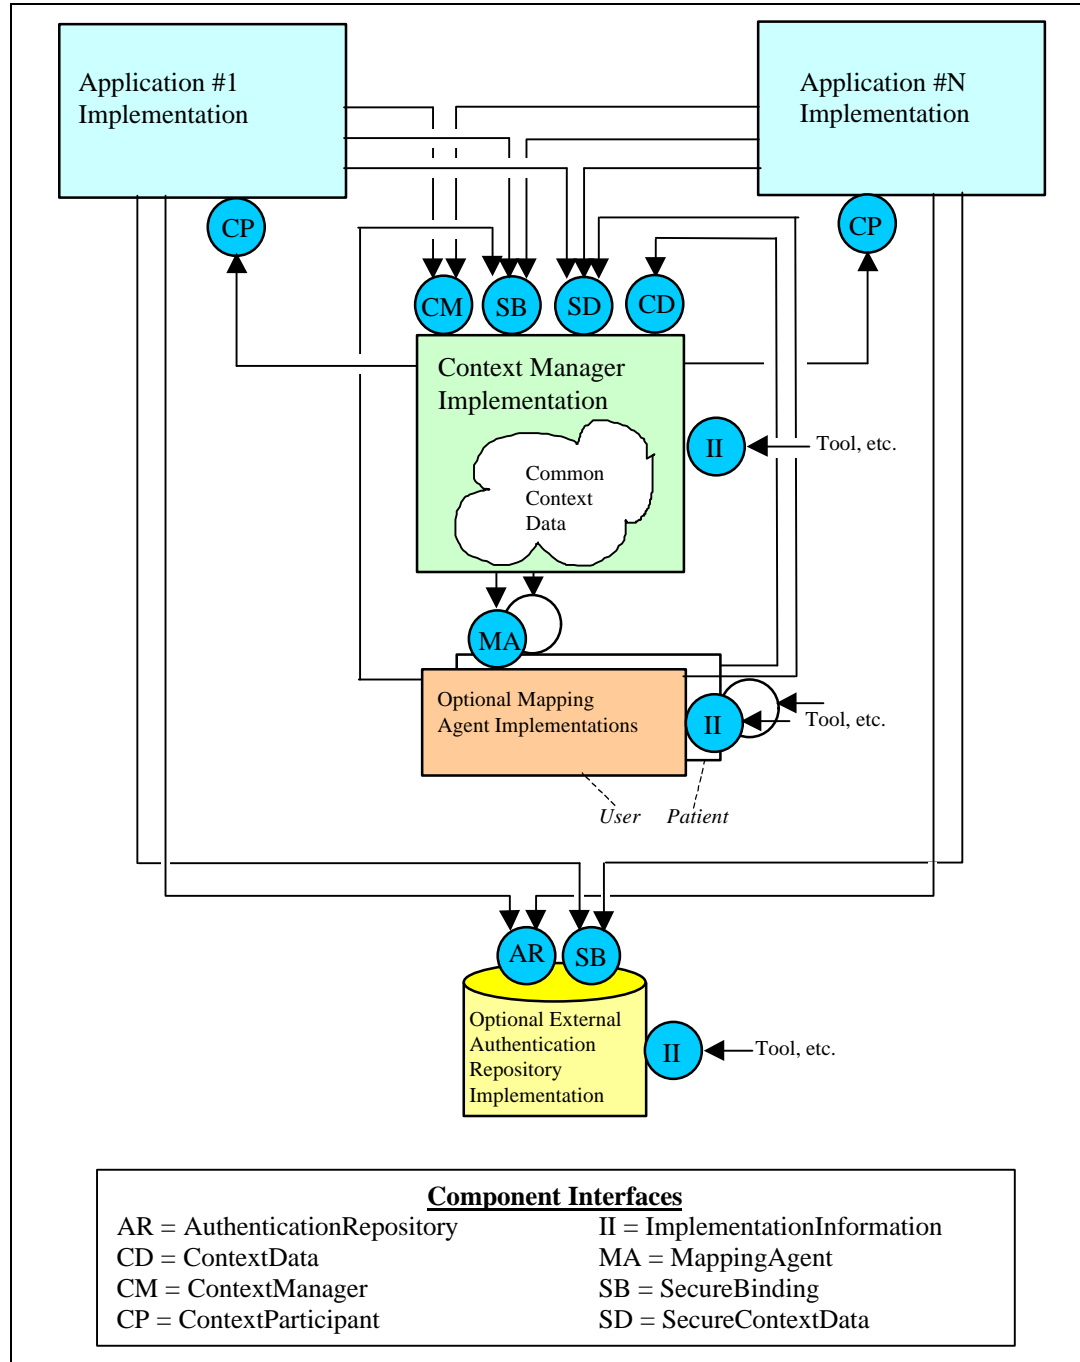


Figure 17: User Link Component Architecture

9.6 User Link Sign-On Process

The process for performing a context change transaction to set the user context is essentially the same as defined for Patient Link for setting the patient context:

- An instigating application initiates a context change transaction and sets the user context within the context manager. This context contains just the identity of the user. It does not include the data used to authenticate the user.
- The context manager consults the user mapping agent (if present) and it adds data to the context manager's user context. This data includes additional logon names by which the user is known.
- The context manager surveys the other applications, and if the transaction completes, they obtain pertinent user context data from the context manager.

The high-level events that transpire when a user signs-on are summarized in Figure 18: User Link Sign-On Process. This description assumes that a user mapping agent is present. The user mapping agent is presumed to know the logon names for all users for all applications. (See Section 9.19, Populating the User Mapping Agent.) The description omits most of the details pertaining to the surveying of the participant applications by the context manager. This process is identical to the process defined for Patient Link. (See Chapter 7.)

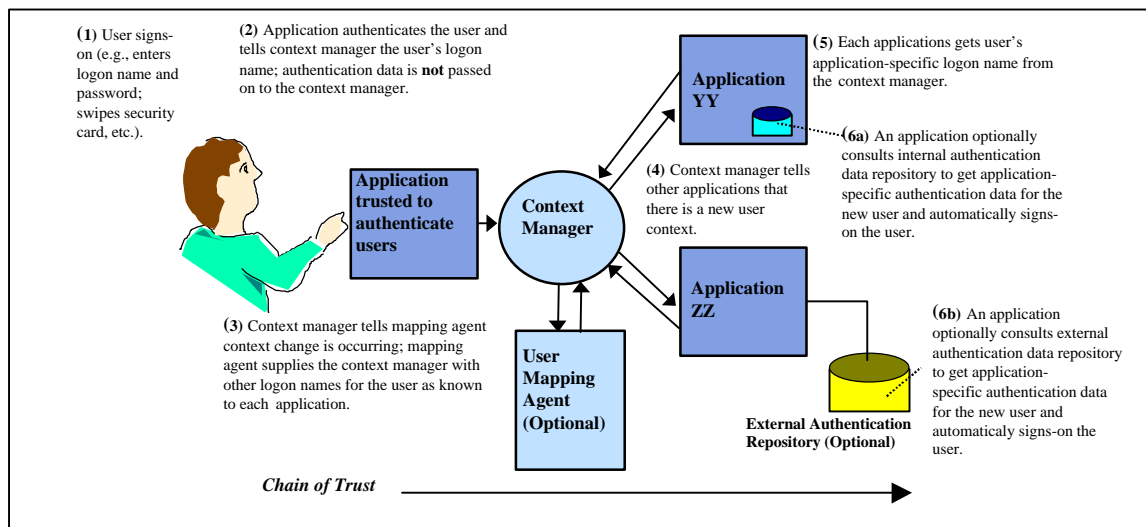


Figure 18: User Link Sign-On Process

9.7 Designating Applications for User Authentication

Any User Link-enabled application can serve as the means by which a user signs-on to all of the User Link-enabled applications on a desktop. To serve in this capacity, the User Link-enabled application shall provide a mechanism for establishing and authenticating the user's identity.

1 The CMA does not specify an application's user authentication mechanism, visual appearance,
2 or implementation. The authentication mechanisms can vary among applications. Applications
3 can be created whose sole purpose is to enable user authentication for desktops comprised of
4 User Linked applications.

5 However, even though any User Link-enabled application has the potential to be used for
6 signing on to a desktop of User Linked applications, the provider institution designates the
7 specific application or applications it trusts for this task. Only the designated applications shall
8 be allowed by a context manager to complete a context change transaction that involves a
9 change to the user subject.

10 The one exception to this rule is that any application can set the user subject to empty. This is
11 so that any application can be used to log-off from a desktop of User Linked applications. (See
12 Section 9.14, Logging-Off and Application Termination.)

13 A context manager implementation-specific configuration process is used for indicating the
14 designated applications for a particular desktop. One, several, or all of the User Link-enabled
15 applications on a desktop can be designated for this purpose. The designated applications for a
16 desktop can differ among desktops. It is recommend that a healthcare institution analyze the
17 use cases for their clinical applications to determine how to best deploy User Link.

18 The decision criteria for a provider institution's choice of whether to designate an application
19 for authenticating users is based upon whether they trust the application's security capabilities
20 as it pertains to user authentication. For example, it might *not* be a good choice to designate an
21 application that maintains user passwords in plain text (which can easily be read by
22 unauthorized users).

23 **9.8 Signing on to Applications Not Designated for Authenticating** 24 **Users**

25 A User Link-enabled application that has *not* been designated for authenticating users on a
26 particular desktop shall not allow the user to sign on to the application or the desktop. The user
27 must sign on to a designated application in order to sign on to a linked but non-designated
28 application. The user must break a non-designated application's link with the common context
29 in order to sign on to just the application.

30 If the application has not been designated for authenticating users and it is the first to be
31 launched on the desktop, the user must either launch an application that has been designated
32 for authenticating users, or the user must break the link of the non-designated application. The
33 user can then sign on to just the non-designated application.

34 The CMA does specify a means by which an application can determine whether it has been
35 designated for authenticating users. See Section 11.3.7.1, InitiateBinding. This enables an

1 application to determine whether it has been designated before a user attempts to sign on to the
2 application. An application can use this information to present or hide its user interface user
3 sign on controls accordingly.

4 **9.9 Application Behavior When Launched**

5 When a User Link-enabled application is launched on a desktop, it should join the common
6 context system established for the desktop. The application should set its user context to match
7 the current user context. If the application is Patient Link-enabled, it should also set its patient
8 context to match the current patient context.

9 **9.10 Multiple Context Subjects**

10 User Link introduces user as an additional common context subject. This creates the need to
11 define what happens to one context, such as the user context, when another context, such as
12 patient context, changes. The simplest approach is to assume that there are no dependencies
13 between subjects.

14 With this assumption, it should be possible for an application to independently set the context
15 data items for just one subject or for both subjects during the course of a single context change
16 transaction. For example, at the end of the transaction the application has changed the user
17 context, the patient context, or both contexts. A context that is not altered by the application
18 shall remain as it was prior to the transaction. The details of managing multiple context
19 subjects are described in the following sections.

20 **9.10.1 The Effect of Multiple Subjects on the Meaning of “Link”**

21 Even though there are multiple subjects in a common context system (e.g., patient and user),
22 there is only one link that coordinates the CMA-compliant applications on a desktop. This
23 means that when an application is linked, it must “tune” to all of the subjects it is capable of
24 dealing with. For example:

- 25 • An application that is only Patient Link-enabled tunes to just the patient context.
- 26 • An application that is only User Link-enabled tunes to just the user context.
- 27 • An application that is both Patient Link-enabled and User Link-enabled tunes to both
28 the patient context and the user context.

29 Conversely, when the user breaks an application’s link, then the application shall no longer be
30 tuned to any context subject.

Independent of the number of context subjects it supports, a single visual cue is provided by an application to indicate whether or not it is linked. The appearance of this cue is defined in the each of the HL7 context management technology-specific user interface specification documents.

9.10.2 Context Manager Support for Multiple Context Subjects

Even though context subjects such are logically independent, there are nevertheless relationships between subjects. These relationships require that context manager implementations have an understanding of multiple subjects and potentially the inter-relationships between the subjects. Further, some applications may need to be aware that they are dealing with multiple context subjects. There are two basic ways to address these issues:

- Maintain a context manager per subject.
- Support multiple context subjects within a single context manager.

The first approach has the advantage that context manager implementations can be specialized to support a single subject. For example, this would enable a Patient Link context manager from one vendor to be used with a User Link context manager from another vendor. The disadvantages are that applications would need to deal with two context managers.

Further, the context managers would need some way to cooperate in order to coordinate transactions that affect multiple subjects (such as a user context change). This coordination would probably require the definition of additional context manager interfaces. This coordination would also increase the complexity of the failure scenarios because of the increased opportunity for partial failures (e.g., one context manager fails while the other context manager continues to function).

The second approach has the advantage that it enables the complexities of dealing with multiple subjects to be hidden within the implementation of the context manager. Additional context manager interfaces are not required, and partial failure scenarios are avoided.

This approach also has the advantage that applications only need to deal with a single context manager.

The second approach has the disadvantage that context manager vendors would need to support all subjects within their context managers. However, it the CMA philosophy to push complexity into the context manager whenever it simplifies the creation of new applications and the reengineering of existing applications. The second approach is the one that shall be pursued in this document because, from the perspective of an application, it is simpler than the first approach.

9.10.3 Effect of Multiple Subjects on Context Change Transaction

For application flexibility and backwards compatibility, it is highly desirable that:

- An application does not have to know about both the user and patient subjects in order to set the context pertaining to just one subject.
- Either or both the user and patient subjects can be updated within a single context change transaction.

However, these desires raise the question of how to treat context data for a subject that is not “touched” during a transaction by the instigating application? There are two approaches:

1. At the completion of the transaction, the untouched subject is *empty*, meaning that it does not contain any context items.
2. At the completion of the transaction, the untouched subject is *unaffected*, meaning that it contains the same items and item values as it did before the transaction.

The first approach is essentially consistent with the existing behavior defined for Patient Link. Specifically, the context manager ensures that each context change transaction begins with an empty context (i.e., no context items). With two subjects, only the subject that is touched during a transaction will contain items at the completion of the transaction.

However, a problem arises with this approach. An application that is only Patient Link-enabled might be co-resident with applications that are Patient Link and User Link-enabled. If the application that is only Patient Link-enabled changes the patient context, the user context shared by the other applications will be lost (i.e., it will be empty).

Applications could be required to know about both subjects and to explicitly copy the subject that is not to be changed from the current context to the new context. However, this creates a burden on the application developers. It is also a substantial impediment to backward compatibility.

The second approach avoids this problem, but requires changes to the behavior of applications or to the behavior of the context manager. To ensure backward compatibility, changing the behavior of applications is ruled out. This eliminates the option of requiring applications to indicate which context subject or subjects it intends to set. (Further this would require changes to the context manager’s interfaces.)

A simpler solution involves a change to the context manager’s behavior that is nevertheless backwards compatible with applications that are only Patient Link-enabled. This solution is described in Section 9.10.4, Context Manager Treatment of Multi-Subject Context Data.

9.10.4 Context Manager Treatment of Multi-Subject Context Data

As is currently the case with Patient Link, when a context change transaction is started, the context manager creates a transaction-specific version of the context data. This version of the context data is initially empty and does not contain any user subject or patient subject context items.

The application that instigated the transaction then establishes the new context by setting context data item values for the user and/or the patient subjects. The application then informs the context manager that it has completed its context changes. The context manager shall then copy the items from the previous context to the new context for any subject that the instigating application did not touch. This shall occur before the context manager surveys the context participants.

The net effect is that the instigating application sets context items for whichever subject(s) it knows about. If a subject was “untouched” by the application, then the items for the subject are automatically post-filled by the context manager to reflect the values as they were *before* the context change transaction.

For applications that are only Patient Link-enabled, this post-filling behavior emulates the existing behavior defined for Patient Link. For applications that are User Link as well as Patient Link-enabled, this behavior enables the user and patient subjects to be managed independently.

With these new rules, an application can just set subjects based upon the user’s explicit gestures, such as selecting a patient, signing on, or both. As with Patient Link, an application only needs to set the user (or patient) subject context items that it is capable of setting. For example, an application may not be able to set all of the corroborating data for a subject. Similarly, a participant application does not have to deal with all subjects, or show all of the context data items defined for a subject.

9.10.5 Effect of Multiple Subjects on Mapping Agents

For simplicity, each context subject (e.g., patient, user) shall have at most one corresponding mapping agent.

When a context change transaction reaches the phase during which the context manager instructs mapping agents to map the context data (i.e., context changes are pending), the context manager shall do so in a sequential manner. Each mapping agent shall be informed only once per transaction that context changes are pending.

The order in which a mapping agent is informed that context changes are pending is not specified. A mapping agent shall not assume the existence of other mapping agents and shall not assume that any subject other than the one it is responsible for mapping has been mapped.

9.10.6 Application Treatment of Multiple Subjects

An application can change either or both the patient and user subjects in a single context change transaction. However, unless the user expects multiple subjects to change as a result of a gesture, it is recommended that an application generally change only one subject at a time. This enables the user to relate changes in the common context to gestures that they have explicitly performed. Cause-and-effect between a user's gesture and a change in application state is an important element in creating systems that are easy for people to use.

9.11 Access Control Lists

Access control lists (ACL), which determine the privileges and capabilities a particular user has, are presumed to be maintained by each application. While it is desirable that there be only one centrally administered ACL, achieving this is beyond the scope of the CMA. However, before central or distributed ACL's can be properly used it is essential that the user be authenticated. This is precisely the capability that User Link supports.

9.12 Empty Contexts

With multiple independent subjects, applications need a way to explicitly indicate that the user context, patient context, or both are empty. The reasons include:

- Enabling applications to change the user context without necessarily carrying over the existing patient context.
- Enabling applications to log-off users by indicating that there is no user context.

The capability to explicitly indicate that a context is empty is already defined in Section 5.6.8, Representing an Empty Context Subject. The stated rules are extended to apply to User Link. This means that the context can identify both a user and a patient, just a user, just a patient, or neither.

When one or both context subjects are empty, all of the applications in the context system shall clearly indicate to the user that this is the case. The appearance of this indication is specified in each of the HL7 context management technology-specific user interface specification documents.

9.13 Changing Users

With User Link, it is advantageous for applications to support a change-user capability. This capability enables a new user to sign on without explicitly requiring that the current user first log off. There are two ways in which this can be implemented by an application:

- 1 • The application performs a single user context change transaction to establish the new
2 user as the current user.
- 3 • The application performs a two-step process. In the first step, the current user is
4 logged off and the user context is set to empty (to indicate that there is no user). In the
5 second step, the new user is signed on, and the user context is set to indicate who the
6 new user is.

7 The first approach is recommended because it is the simplest and the most efficient from the
8 perspective of the context system (e.g., only one context change transaction per user change).
9 The second approach is acceptable, however the two step process should be invisible to users.

10 The gestures needed to change the user, and the appearance of the application as it pertains to
11 this capability, are not specified by the CMA.

12 **9.14 Logging-Off and Application Termination**

13 User Link provides applications with an easy way to enable users to:

- 14 • Terminate a specific User Linked application on the clinical desktop⁶.
- 15 • Log off from a specific User Linked application on the clinical desktop.
- 16 • Log off from all of the User Linked applications on the clinical desktop.

17 There are many possible ways in which these capabilities can be realized in a common context
18 system. The approach described in Table 1: User Link-Enabled Application Behavior for
19 Termination and Log-Off is defined because it is simple for users to understand, yet enables
20 design flexibility for application developers.

21 The basic idea is that each User Link-enabled application optionally supports gestures that
22 enable the user to terminate the application, log off from just the application, or log off from
23 all of the User Linked applications that are resident on the same desktop.

24

⁶ Terminating all of the applications on a desktop is not supported because there is no way to indicate this event via a change to the user context subject.

User Action	Effect on Application That User's Action Is Directed At	Effect on the Common Context	Effect on Other User Linked Applications on the Desktop
Terminate a specific User Linked application.	Application leaves the common context, ceases execution, and exits	None.	None.
Log-off from a specific User Linked application. See Interaction Diagram 15: User Logs Off From One Application.	Application: <ul style="list-style-type: none"> continues to run, logs the user off, visually indicates that it has no user, leaves common context (i.e., breaks link) 	None.	None.
Log-off from all of the User Linked applications that are resident on the same desktop. See Interaction Diagram 16: User Logs-Off From Desktop.	Application: <ul style="list-style-type: none"> continues to run, instigates a context change transaction to set the user context to empty, visually indicates that it has no user, continues to be a context participant. 	User subject changed to empty.	When the context change is completed, each application: <ul style="list-style-type: none"> continues to run, logs the user off, visually indicates that it has no user, continues to be a context participant.

Table 1: User Link-Enabled Application Behavior for Termination and Log-Off

All User Link-enabled applications must behave properly as participants in a context change transaction, as described in Table 1. All User Link-enabled applications must be able to properly deal with the context when the user context is empty.

However, the CMA does not specify the user gestures that are needed to initiate the actions described in Table 1. The gestures may be different among applications. Further, an application may chose which action gestures, if any, it will support. For example, a particular application might not enable the user to terminate it, log off from it, or log off from the User Linked desktop.

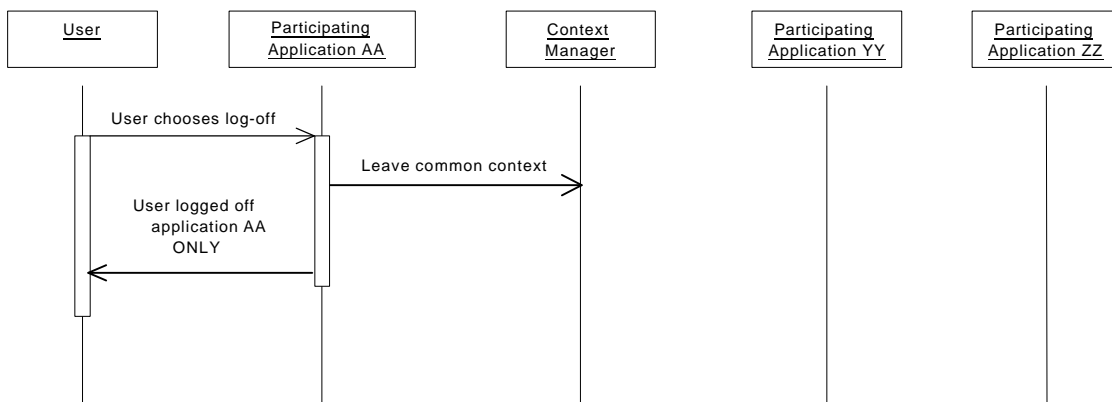
An application that enables the user to log off shall clearly indicate that in doing so, the user will cause the application to break its link with the common context system.

There are several subtleties involved with the behaviors described in Table 1:

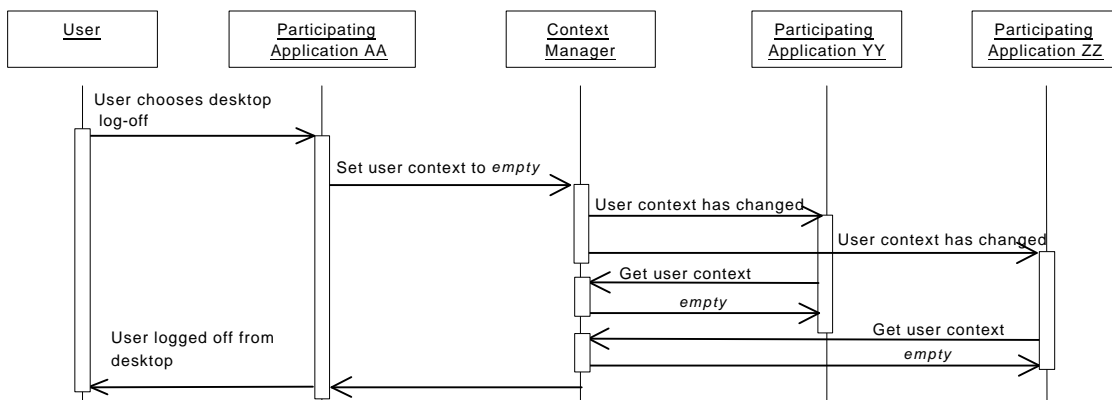
- Any application can set the user context to empty, including applications that have not been designated for authenticating users. This enables any application to be used for logging off from all of the User Linked applications on a desktop.

- A user might terminate the application(s) designated for authenticating users. The next user will need to relaunch one of the designated applications before being able to sign on to the User Linked desktop.
- It is conceivable that the collective capabilities of a particular set of User Link-enabled applications on a desktop result in a system that does not provide any way for the user to log off from the desktop. A site must be mindful in its choice of applications in order to prevent this from happening.

One issue with desktop log off is the treatment of “busy” applications. Busy applications affect single sign on as well as desktop log-off, and is dealt with in Section 9.17, Busy Applications.



Interaction Diagram 15: User Logs Off From One Application



Interaction Diagram 16: User Logs-Off From Desktop

9.15 Automatic Log-Off

An automatic log-off logs the current user off of the User Linked applications on a desktop when the user has not interacted with the applications for an appreciable period of time.

Any application can initiate an automatic log-off by performing a context change transaction that sets the user context to empty. This will have the effect of causing all of the other User Linked applications on the desktop to also log the user off. Once an automatic log-off has completed, the next user signs-on via one of the designated applications.

In contrast to a user-initiated log-off, an automatic log-off is initiated automatically by an application. The CMA does not specify an automatic log-off policy or implementation. It is an application decision as to how and when to initiate an automatic log-off.

For example, an application might monitor user interactions with the mouse and keyboard to determine whether or not the user is actually engaged in using any of the applications on the desktop. The capability to do this depends upon the application's implementation and the underlying desktop technology.

An application that initiates a context change transaction to affect an automatic log-off must be prepared to handle the condition in which surveyed applications are busy, or have responded with a conditional accept of the transaction. In this case the instigating application shall cancel the context change transaction. It shall not present a dialog to the user, as this could be disruptive or confusing to the user. The application may elect to initiate an automatic log-off again in the future.

It is necessary that the administrator is able to configure the behavior of automatic log-off as it pertains to a clinical desktop. Otherwise, the administrator has no control over an application whose policy for initiating an automatic log-off interferes with the users' work.

Therefore, any application that initiates an automatic log-off shall provide a means for controlling this capability. Specifically, it shall be possible to configure that application in terms of whether the log-off it initiates is desktop-wide (and therefore affects all of the context participants), or is limited to just the application. If the automatic log-off is limited to just the application, then the application shall not perform a context change transaction when the automatic log-off interval transpires. Instead, it shall just log the user off from itself.

9.16 Reauthentication Time-out

A reauthentication time-out requires the currently signed-on user to reauthenticate herself before being allowed to continue using the applications on a clinical desktop. The time-out occurs when the user has not interacted with the desktop for an appreciable period of time. Applications maintain their internal state as the user left it prior to the time-out, but interaction with the applications cannot resume until the user has been reauthenticated.

The time-out often manifests as a screen that overlays the entire display and that provides a mechanism with which the user can reauthenticate herself. However, the CMA does not specify a reauthentication time-out policy, visual appearance, or implementation.

Any application can initiate a reauthentication time-out. However, a User Link-enabled application that does so shall be:

- responsible for enabling the user to re-authenticate herself
- configurable such that a systems administrator can enable or disable the time-out capability.

These requirements enable sites to practice the following CMA recommendation: only a User Link-enabled application that has been designated for authenticating users should be allowed to initiate a reauthentication time-out. This enables the user to reauthenticate herself using an application that is also normally used for signing on to the clinical desktop.

This recommendation avoids the problem of forcing the user to be reauthenticated by an application not normally used for signing on, and therefore having to remember their logon name and password for the application.

Once the current user is reauthenticated, then the User Link-enabled applications resume as they were. If a different user signs on, then the User Link-enabled applications handle this as they do whenever there is a change of user.

9.17 Busy Applications

When a context change transaction is conducted, it is possible that an application is unable to participate because it is busy. For example, a single-threaded application that has a modal dialog open will not be able to respond until the dialog is closed.

User Link deals with busy applications the same way as for Patient Link. Specifically, a busy application effectively prevents a context change transaction from occurring. The only option for the application that instigated the transaction is to ask the user if they want to break the link.

Breaking the link has the potential to compromise user security. With a broken link, multiple users would effectively be logged on to different applications on the same desktop.

However, this situation is not substantially different from breaking the Patient Link, which results in different applications on the same desktop being tuned to different patients. Further, without the option to break the link, CMA support for some important use cases, such as “stat” admissions (see Section 7.12, Stat Admissions), would be lost.

9.18 Co-Existence with Applications Not User Link-Enabled

User Link-enabled applications will co-exist with applications that are not User Link-enabled. Users will still need to manually sign on to and log-off from each of the applications that are not User Link-enabled.

Co-existence can create confusion among users, as they might assume that all of the applications on a desktop are User Link-enabled. Training, plus visual cues documented in the HL7 context management technology-specific user interface specification documents are partial solutions. Ultimately, users will come to learn which applications are User Link-enabled, and which are not, and will adjust their use of these applications accordingly.

9.19 Populating the User Mapping Agent

The user mapping agent is conceptually similar to the patient mapping agent defined for a Patient Link common context system. For example, both types of mapping agents implement the same interface specification, MappingAgent. However, the behavior and management of the user mapping agent is substantially influenced by security considerations. Several of these considerations are described in this section. The role of the user mapping agent is illustrated in Figure 19: User Subject Context Data Mapped for Different Applications.

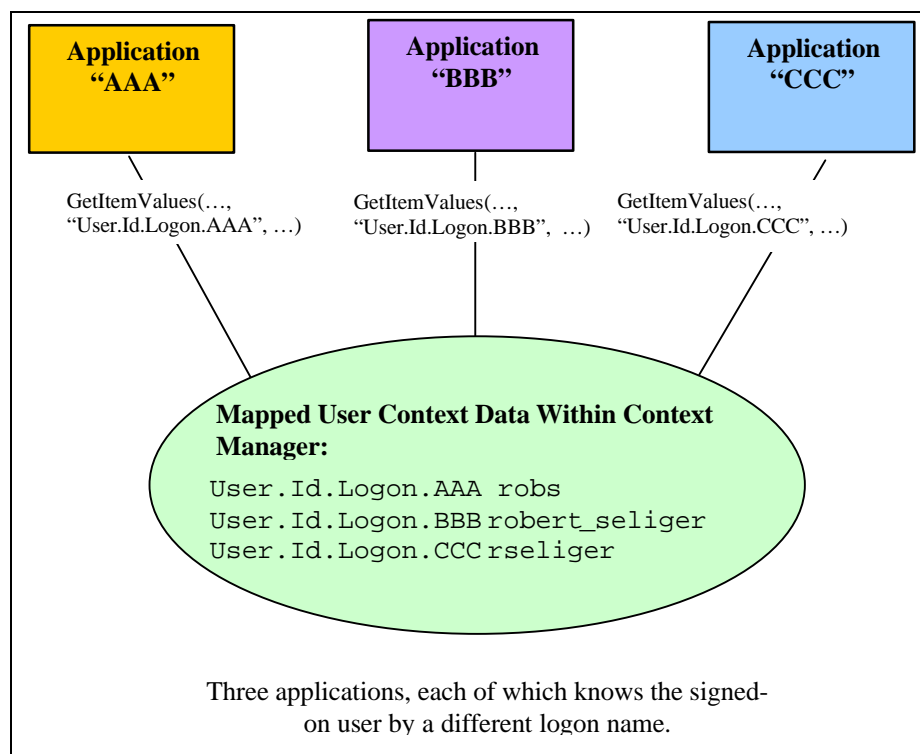


Figure 19: User Subject Context Data Mapped for Different Applications

1 In order for the user mapping agent to be able to provide additional logon names for users, it
2 must be populated with the necessary logon names. However, unlike the patient mapping
3 agent, for which there exists healthcare standards that can be used to obtain the necessary
4 patient data (e.g., HL7's Admission/Discharge/Transfer messages), an equivalent means does
5 not exist for user data. In the absence of applicable standards, the means by which a user
6 mapping agent is populated depends upon the user mapping agent implementation.

7 **9.20 Authentication Repository**

8 The chain of trust has the potential to maximize the overall security of a common context
9 system because the data used to authenticate a user is never passed between applications and
10 therefore cannot be easily intercepted or spoofed. However, not passing around this data
11 creates a problem when there are applications that require user authentication data to perform
12 a user sign on. For example, many existing healthcare applications require the user's password
13 to establish sessions with their underlying databases.

14 The common context system therefore includes a user authentication data repository as an
15 additional context management component. This repository enables applications to securely
16 maintain application-specific user authentication data. The repository is used by applications
17 that do not have a built-in means to easily sign on a user given only a logon name. The
18 repository may be implemented as a distributed or centralized service.

19 For example, some applications obtain the user's password from the user and then hand it off
20 to an underlying database. The database does the actual authentication. The security
21 capabilities of the database prevent these applications from retrieving user passwords.
22 Therefore, it is not possible for these applications to sign on a user knowing only the user's
23 logon name. For these applications, an external means of maintaining user logon names and
24 associated authentication data is required.

25 The authentication repository provides a way of doing this that is minimally invasive to the
26 application. The repository is not used for authenticating users. Rather, it enables existing
27 applications that need user authentication data to sign on the user to have a means for
28 obtaining this data when participating in a User Link common context system.

29 The User Link user authentication data repository provides the capability to securely store the
30 data that an application uses to authenticate its users. The application can use a user's logon
31 name to retrieve the user's authentication data from the repository. The application can then
32 use the authentication data to establish a user session with a database or other underlying
33 application services.

34 In keeping with the spirit of the CMA, the interfaces to the authentication repository, but not
35 its implementation, are defined. These interfaces enable an application to securely retrieve a

1 user's authentication data and to update this data when necessary (for example, if the
2 application periodically requires that users change their passwords).

3 **9.20.1 Repository Implementation Considerations**

4 The repository can be implemented as a central or distributed service that services multiple
5 applications. However, the repository shall always appear as a private service to each
6 application. This means that an application should never be aware that there are other
7 applications using the repository.

8 The user authentication data stored in the repository on behalf of an application shall be
9 encrypted by the application prior to being communicated to the repository. The encryption
10 technique that is used is determined by the application. The authentication data shall remain
11 encrypted within the repository, as the repository never has the need to interpret or use this
12 data.

13 The interface AuthenticationRepository enables an application to put tuples comprised of a
14 logon name and a corresponding bit stream (representing the user's authentication data) into
15 the repository. This interface also enables an application to retrieve a user's authentication
16 data using the user's logon name.

17 The means by which the repository maintains its data must be secure and shall guard against
18 security attacks. However, the security mechanisms that are employed to achieve these
19 objectives are an authentication repository implementation decision.

20 **9.20.2 Populating the Repository**

21 The authentication repository needs to be populated with the authentication data for each user
22 for each application that it services. One way to do this is to create a batch process that loads
23 the necessary data. However, in many cases the necessary data is inaccessible. For example,
24 most database management systems do not provide a means for accessing the user passwords
25 that they store.

26 A simpler alternative is to incrementally populate the repository. This can be accomplished by
27 involving each of the applications that use the repository in the process of populating the
28 repository, as follows:

- 29 • When the context manager informs the application that the user context has changed,
30 the application obtains the logon name for the new user from the context manager.
- 31 • The application then accesses the repository to securely retrieve the user's
32 authentication data. The user's logon name is supplied as the search parameter.

- 1 • If the repository cannot find the user logon name, which will be the case if the
2 repository has not yet been populated with data for the user, then it informs the
3 application that the logon is not known.
- 4 • The application then prompts the user to enter his/her authentication data by whatever
5 means the application normally uses (e.g., a password dialog box).
- 6 • The application attempts to sign-on the user using whatever underlying mechanism
7 (e.g., database) it normally uses to do this.
- 8 • If the user is successfully signed on, then the application updates the authentication
9 repository with the user's authentication data, using the user's logon as the update key.
10 The application shall encrypt the user's authentication data prior to putting the data in
11 the repository.

12 This scheme is relatively easy to implement for almost any application. It is essential, though,
13 that the repository and its interfaces are secure, as detailed in Chapter 11.

10 Chain of Trust

This chapter defines the behaviors, algorithms, policies, and protocols that User Link-enabled applications and components must adhere to in order to properly realize the chain of trust.

10.1 User Context Change Transactions and the Chain of Trust

The major difference between a context change transaction that involves the user subject and a transaction that involves only the patient subject is support in the former for the chain of trust. Additional application and component behaviors are defined to prevent the chain of trust from being violated.

Two types of defenses are required:

- The applications and components that participate in the chain of trust must be able to authenticate each other's identity. The objective is to prevent rogue applications or components from impersonating a real application or component as a means to manipulate the user context. Such manipulations could result in an unauthorized user gaining access to the User Link-enabled applications.
- The applications and components that participate in the chain of trust must be able to validate the integrity of user context data that they communicate to each other. The objective is to prevent a rogue program from modifying the data as it is passed between applications and components as a means to manipulate the user context. Such manipulation could result in an unauthorized user gaining access to the User Link-enabled applications.

Techniques for creating the chain of trust using passcodes, message authentication codes, and digital signatures are described next.

10.2 Creating the Chain of Trust

There are three general sources of mechanisms for creating the chain of trust:

- Mechanisms incorporated into existing commercially available object infrastructures, such as those based upon CORBA or COM.
- Mechanisms based upon existing commercially available secure communications infrastructures, such as the Secure Socket Layer service (SSL) or the Secure Hyper-Text Transfer Protocol (S-HTTP).

- Mechanisms based upon existing widely available security building blocks, such as public key / private key encryption.

These alternatives are discussed next.

10.2.1 Object Infrastructures

It is conceivable that the chain of trust could be realized using the security mechanisms built into commercially available object infrastructures such as those based upon CORBA or COM. Unfortunately, these infrastructures currently employ security models that are fundamentally different from what is needed for User Link:

- Security for these infrastructures is based upon keeping track of who the user is and their respective access privileges.
- To do this requires that the user has signed on to the underlying operating system.
- However, signing on at the operating system level takes too much time. This is the very problem that User Link is trying to solve.

For example, security in Microsoft's COM-based infrastructure is based upon tracking who the user is and what their permissions are. This means that when security is enabled for a COM interface, a COM server accepts or rejects a COM client's access attempts based upon the privileges of the user on whose behalf the COM client is working. This does not work for User Link because a COM server (specifically, the context manager) needs to accept or reject accesses based upon which application is the COM client. The user is not relevant in this case.

It may be possible to establish a stylized approach for adapting object infrastructure security mechanisms to realize the chain of trust. However, this could make it particularly difficult to define a technology-neutral specification for the chain of trust. This could result in different User Link architectures for different technologies. This is counter to the overall CMA objective of technology-neutrality.

10.2.2 Secure Communications Protocols

User Link-enabled applications and the various CMA components could communicate using a secure communications protocol, such as the Secure Sockets Layer (SSL) service. SSL enables secure (i.e., encrypted) transmission of data between a client and a server. It also enables a client to authenticate a server (and a server to authenticate a client).

SSL uses the RSA public key encryption system for authentication and for data integrity and confidentiality. Of interest for the chain of trust is the SSL capability for clients and servers to authenticate each other. An SSL server uses its private key to create a digital signature. Public keys are issued to prospective clients. The public key is used by the client to authenticate the

server by decoding the server's signature. Only a signature that has been encoded using the server's private key can be (easily) decoded via the server's public key.

For example, in the chain of trust, an SSL connection would be established between an application that has been designated for authenticating users and the context manager. In this scenario, the application is an SSL *server*, while the context manager is an SSL *client*.

SSL and its secure communications counterparts, such as S-HTTP, provide off-the-shelf mechanisms for implementing the chain of trust. However, this technology has not been integrated with popular object infrastructures, such as those based upon COM or CORBA.

While secure communication services could provide a means for implementing the chain of trust, the practical implications of using multiple communications technologies within the User Link architecture are a cause for concern. For example, it could become overly complicated to have some communications be via COM or CORBA interfaces, while other communications use SSL or S-HTTP.

Further, the chain of trust generally does not require confidentiality. For example, the User Link architecture does not require that sensitive data, such as a user's password, be communicated between applications. Secure communication channels are overkill and are not a good fit for User Link.

10.2.3 Security Building Blocks

The security building blocks that are available on most popular operating systems can form the basis for realizing the chain of trust. The two building blocks of particular interest are:

- Digital signatures.
- Secure (or one-way) hashing.

Digital signatures, which cannot be easily forged, are typically used by people as a means to authenticate each other's identity whenever they communicate electronically. However, a digital signature also enables an application or component to identify itself in a way that can be authenticated whenever it communicates with another application or component.

Digital signatures are formed using public key / private key encryption techniques. While these techniques enable encryption, they also enable the formulation of digital signatures. An application or component formulates its digital signature using its private key and sends the signature along with the data that it wants to share. The recipient of a signed message applies the sender's public key to the signature to authenticate the sender and to verify the integrity of the data that was sent.

1 There are several public key / private key algorithms and related standards. Commercial
2 implementations of many of these algorithms are available in a variety of technologies. RSA is
3 an example of an algorithm that has been widely implemented.

4 A secure hash function is used for producing a unique numeric surrogate from an arbitrary
5 data stream. It is improbable that two different data streams will yield the same hash value. A
6 secure hash function is an essential part of the infrastructure needed to support the use of
7 digital signatures.

8 Specifically, a secure hash function enables the efficient computation of a digital signature. A
9 secure hash function also plays a role in enabling public keys to be reliably distributed. It is
10 essential that the holder of a public key is able to determine who (or what) the key belongs to.
11 Otherwise an impostor could present its own public key while claiming to be someone or
12 something that it is not. The holder of the public key would mistake subsequent
13 communications as coming from a valid source when in fact it came from an impostor.

14 There are several secure hashing algorithms and related standards. Commercial
15 implementations of many of these algorithms are available in a variety of technologies. MD5 is
16 an example of an algorithm that has been widely implemented.

17 Taken together, digital signatures and secure hashing could be used in the chain of trust as the
18 means for User Link-enabled applications and User Link components to authenticate each
19 others' identity each time they communicate. This capability is fundamental to the
20 establishment and maintenance of the chain of trust.

21 To accomplish this, a digital signature would be explicitly included as a method parameter for
22 each CMA-specified interface that requires this level of security. The use of digital signatures
23 enables the specification of a system that has the desired User Link semantics and that can be
24 readily implemented using existing security standards and technology.

25 Creating a system that employs digital signatures for applications and components is simpler
26 than creating a signature-based system for users. This is because the population of applications
27 and User Link components that requires signatures is small compared to the number of users of
28 the system. Further, the population of applications and User Link components does not change
29 nearly as often as the user population. The result is that the work required to create and
30 maintain the chain of trust is substantially less than would be the case if user signatures were
31 required.

32 Another advantage of digital signatures is that they can be used to ensure the integrity of any
33 data communicated during interactions among and between User Link components and User
34 Link-enabled applications. The recipient of the data can use the signature to determine if the
35 data has been tampered with between the time it was sent and the time it was received.

1 Method-based digital signatures fit well with the component-based Context Management
2 Architecture. For example, realizing the chain of trust in this manner enables a technology-
3 neutral specification for the chain of trust. This is because the approach can exploit
4 capabilities common to public key / private key implementations that are commercially
5 available in multiple technologies. Further, the ways in which digital signatures are used can be
6 arranged to achieve the desired security behaviors needed for User Link.

7 The trade-off is that more effort is required to architect the chain of trust than would be the
8 case if a standard “off-the-shelf” component-based solution was available. This trade-off is
9 viewed as acceptable. Therefore the approach pursued in the CMA is to use method-based
10 digital signatures as the basis for the chain of trust.

11 **10.2.4 Security Attacks On the Chain Of Trust**

12 The primary challenge for realizing the chain of trust is minimizing the likelihood that an
13 intruder is able to violate the chain of trust to obtain access to a User Link-enabled application.
14 This violation could occur if a rogue program was able to set the user context to represent a
15 user who either has not been authenticated, or who is different from the user who has been
16 authenticated.

17 The chain of trust based upon the security building blocks described in Section 10.2.3,
18 Security Building Blocks, defends against the security attacks described in the table below, all
19 of which are directed at manipulating the user context. Refer to Figure 18: User Link Sign-On
20 Process for the specific trust relationships:

Attack	Defense
Attempt to impersonate an application in order to set the user context (Step #2).	An application presents its signature to the context manager in order to set the user context. The context manager uses the signature to authenticate the application to ensure that has been designated for authenticating users.
Attempt to impersonate the context manager so that the user context that the user mapping agents sees, and therefore maps, is bogus (Step #3).	The context manager presents its signature to the mapping agent when the mapping agent gets the user context data from the context manager. The mapping agent uses the signature to authenticate the context manager.
Attempt to impersonate the user mapping agent as a means to set bogus user logon names within the user context (Step #3).	The mapping agent presents its signature to the context manager when it sets user context data. The context manager uses the signature to authenticate the mapping agent.
Attempt to impersonate the context manager so that the user context that a participant application sees is bogus (Step #5).	The context manager presents its signature to the participant application when the application gets the user context data from the context manager. The application uses the signature to authenticate the context manager.
Attempts to impersonate the authentication repository as a means to obtain user authentication data from an application (Step #6b).	The application encrypts the user authentication data using the authentication repository's public key before providing the data to the repository. Only the real authentication repository can decrypt this data. Further, the application pre-encrypts the data using an application-specific encryption scheme. The data remains encrypted even when stored inside the repository.
Attempt to impersonate an application as a means to obtain user authentication data from the authentication repository (Step #6b).	An application must present its signature to the authentication repository when it gets user authentication data from the repository. The repository uses the signature to authenticate the application. Further, the application encrypts the authentication data before storing it in the repository. Only the application that encrypted the data can subsequently decrypt it.

Table 2: Chain of Trust Attacks and Defenses

The chain of trust does not necessarily need to defend against every type of attack, including attacks to gain access to the user's logon name (i.e., Step #4). A user's logon name is easy to guess or obtain, and in the absence of user authentication data (e.g., a password) a logon name does not provide a means for gaining access to a system.

1 The chain of trust also does not defend against applications that do a poor job of authenticating
2 users (i.e., Step #1). Provider institutions must ensure that the applications they designate for
3 authenticating users meet their security needs.

4 Other types of attacks that are not defended by the chain of trust can result in a denial of
5 service, which may cause a common context system to function improperly. For example, a
6 rogue program might continually invoke context manager methods, causing the context
7 manager's performance to degrade while it services these invocations.

8 These programs do not breach security in terms of enabling unauthorized access to User Link-
9 enabled applications, but they do result in inconveniences for users of the system. In general it
10 is extremely hard, and can be quite costly, to defend against denial of service attacks.

11 The most effective preventatives for denial of service attacks begin with physical security, in
12 which a malicious user is denied access to any of the computers within a system. Without
13 access to the system, a malicious user will have a much harder time installing rogue programs.
14 Physical security is strongly encouraged, but it is beyond the scope of the CMA to specify the
15 necessary measures.

16 Additional potential limitations of the chain of trust are described in Section 10.2.5, Chain of
17 Trust Implementation Limitations.

18 **10.2.5 Chain of Trust Implementation Limitations**

19 A secure implementation of the chain of trust requires that the User Link components (i.e.,
20 context manager, applications, mapping agent, authentication repository) all have a robust way
21 of authenticating each other's identity. Providing this capability requires the use of underlying
22 operating systems primitives, including file access privileges and memory protection
23 mechanisms.

24 Not all operating systems implement these security primitives to the same degree of robustness.
25 The approach for implementing the chain of trust described below is therefore fundamentally
26 limited by the capabilities (or lack thereof) of the underlying operating system upon which a
27 User Link system is deployed.

28 In particular, Windows NT and most Unix-based operating systems provide the necessary
29 primitives. User Link systems deployed on these operating systems will offer robust security
30 capabilities. In contrast, Windows 95 and Windows 98 lacks many of the necessary primitives.
31 User Link systems deployed on this operating system will offer useful capabilities, but the
32 systems will not be any more secure than native Windows 95/98.

10.3 Digital Signatures and CMA Components

Digital signatures created using a public key / private key encryption system are incorporated into the component interfaces defined for User Link-enabled applications and components. In the chain of trust these signatures (and corresponding keys) are not associated with a user, but rather with an application or component. The signatures and keys for a particular application are the same independent of who the user is.

Several of the methods defined for the existing context manager interfaces already require that applications identify themselves (e.g., `ContextData::SetItemValues`). The participant coupon, which is an integer, is assigned by the context manager to an application when it joins a common context system (via `ContextManager::JoinCommonContext`). This coupon is subsequently used by the application to identify itself when it calls a context manager method that requires application identification.

The methods requiring applications to identify themselves do so in enforce the correct behavior of a common context system. For example, only the application that instigated a context change transaction or a mapping agent can set context data. Similarly, only the instigating application can end the transaction in progress.

However, the use of a participant coupon is not intended to be a security mechanism. For example, a rogues application can impersonate a valid application by obtaining (or even guessing) the value of the valid application's coupon. Coupons are simply to enable the context manager to identify the applications it is dealing with.

An elaboration of the coupon approach is to use digital signatures as a means for applications to identify themselves in a manner that can be authenticated. It is relatively straightforward to use digital signatures in addition to coupons whenever it is necessary to authenticate an application or component.

Based on this approach, CMA interfaces are defined that enable the establishment of the necessary signature-based security relationships among and between applications and context management components. Additional CMA-defined interfaces subsequently enforce these security relationships as applications and components interact during the course of a context change transaction.

10.3.1 Public Key / Private Key Encryption as a Means for Generating Signatures

Providing applications with digital signatures requires that each application or component that is to be trusted is assigned a public key and private key based upon an algorithm such as RSA. The private key is used to create a digital signature. The corresponding public key is used to verify the signature.

For example, an application supplies its participant coupon *and* its signature to the context manager whenever it performs a context manager method that requires the context manager to

1 authenticate the identity of the application and validate the integrity of the data sent by the
2 application.

3 A digital signature is formed by applying a secure hash function (alternatively known as a one-
4 way hash function) to the data that is to be transmitted. The resulting hash value is referred to
5 as the message digest, as it is a numeric surrogate for the plain-text message. It is
6 computationally improbable that two messages will produce the same hash value⁷.

7 The message digest is then encrypted by the sender using its private key⁸. The digest can only
8 be decrypted using the sender's public key. In other words, any party holding the sender's
9 public key can authenticate that the message came from the sender and that the data sent was
10 received in tact⁹.

11 The encrypted hash value enables the sender of the data to ensure that the receiver of the data
12 can authenticate the sender's identity. The receiver uses the same secure hash function as the
13 sender to perform its own computation of a hash value using the data it received. Note that the
14 data was not encrypted. Just the hash value computed from the data was encrypted.

15 The receiver compares the hash value it computed with the value it decrypted. The encrypted
16 hash value can only be successfully decrypted using the public key that matches the sender's
17 private key. If the hash values match, then the data sender's identity has been confirmed, and
18 the integrity of the data has been validated.

19 If the hash values do not match, then either the data was tampered with between the time it was
20 sent and was received, or the sender is not who it claims to be.

21 The algorithm for creating the hash value must be compatible with the public key / private key
22 scheme that is employed. For example, if RSA is the public key / private key scheme that is
23 used, then an RSA-supported hashing algorithm (e.g., MD5, SHA-1) must be employed to
24 create the hash value. When the signature is computed in this manner, authenticity and data
25 integrity can be verified.

26 The specific secure hash algorithm and the public key / private key scheme that is employed is
27 technology-specific. Each of the HL7 Context Management Technology Mapping

⁷ When a secure hash function is used, it is also computationally infeasible to invert the computed hash value. Specifically, given the secure hash function f and input value x , $f(x)$ is relatively easy to compute. However, even knowing f it is infeasible to compute x given $f(x)$.

⁸ The signing of a message digest rather than of the plain-text message is a performance expediency. A digest is typically several bytes in size, whereas the message represented by a digest can be of arbitrary size. It is generally faster to encrypt the digest rather than the entire message.

⁹ This is the inverse of the process used to send a secret message, in which the sender encrypts data with the intended recipient's public key. Only the holder of the private key can decrypt the data.

Specifications indicates the secure hash algorithm public key / private key scheme that is needed for a particular technology-specific implementation.

The overall process for signing a message is illustrated Figure 20: Signing A Message.

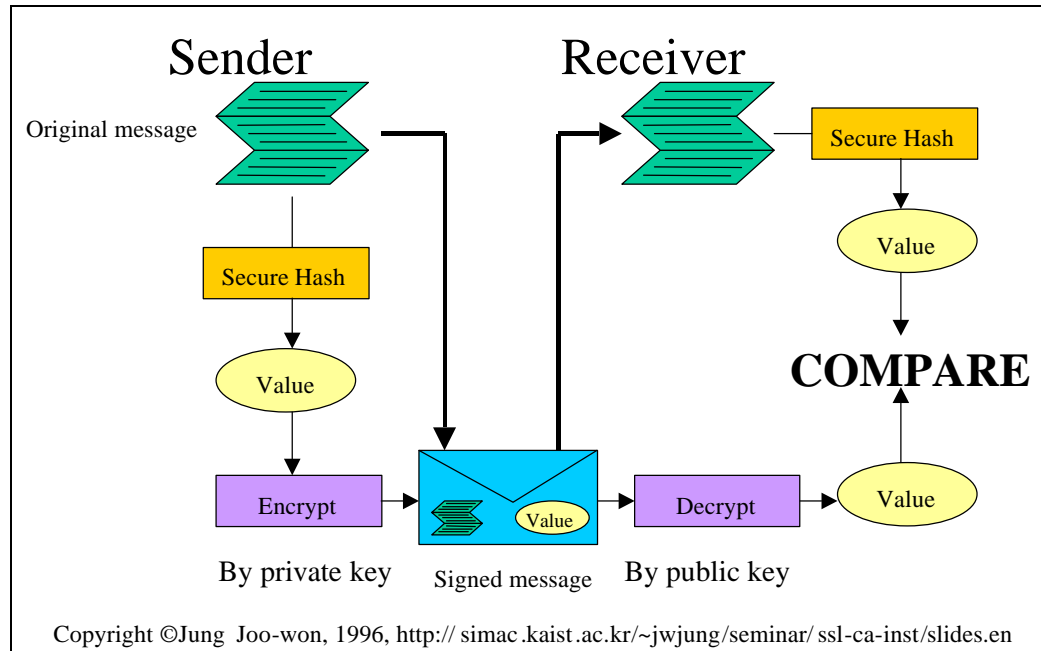


Figure 20: Signing A Message

10.3.2 Incorporation of Signatures into the Context Management Architecture

Digital signatures are incorporated in the Context Management Architecture to enable authentication between User Link-enabled applications and User Link components. For example, digital signatures enable the context manager to authenticate the identity of any application that performs a context manager method. The context manager can also ensure the integrity of the parameter values that it received from the application.

The context manager accomplishes this by computing a hash value from the input parameters it receives from the application. To obtain the application-computed hash value from the signature the context manager must use the same public key / private key scheme as the application. The context manager must also use the same hash algorithm as the application.

The context manager compares the hash value it computes to the hash value it has obtained by decrypting the application's digital signature. If the two hash values match, then the method invocation is authentic and data integrity is ensured.

Otherwise, there has been a breach of security: either the method was invoked by an impostor of the application, and/or the parameter values provided by the application were tampered with

after they were sent but before they were received by the context manager. The context manager rejects the method invocation.

To be more specific, for the context manager method `SecureContextData::SetItemValues`, the hash value would be computed using the value of the participant application's coupon (i.e., input parameter *participantCoupon*), current context change transaction coupon¹⁰ (i.e., input parameter *contextCoupon*), the names of the items whose values are to be set (i.e., input parameter *itemNames*), and the values for these items (i.e., input parameter *itemValues*).

The use of a hash in forming a signature is illustrated Figure 21: Forming Signature Using Method Parameters.

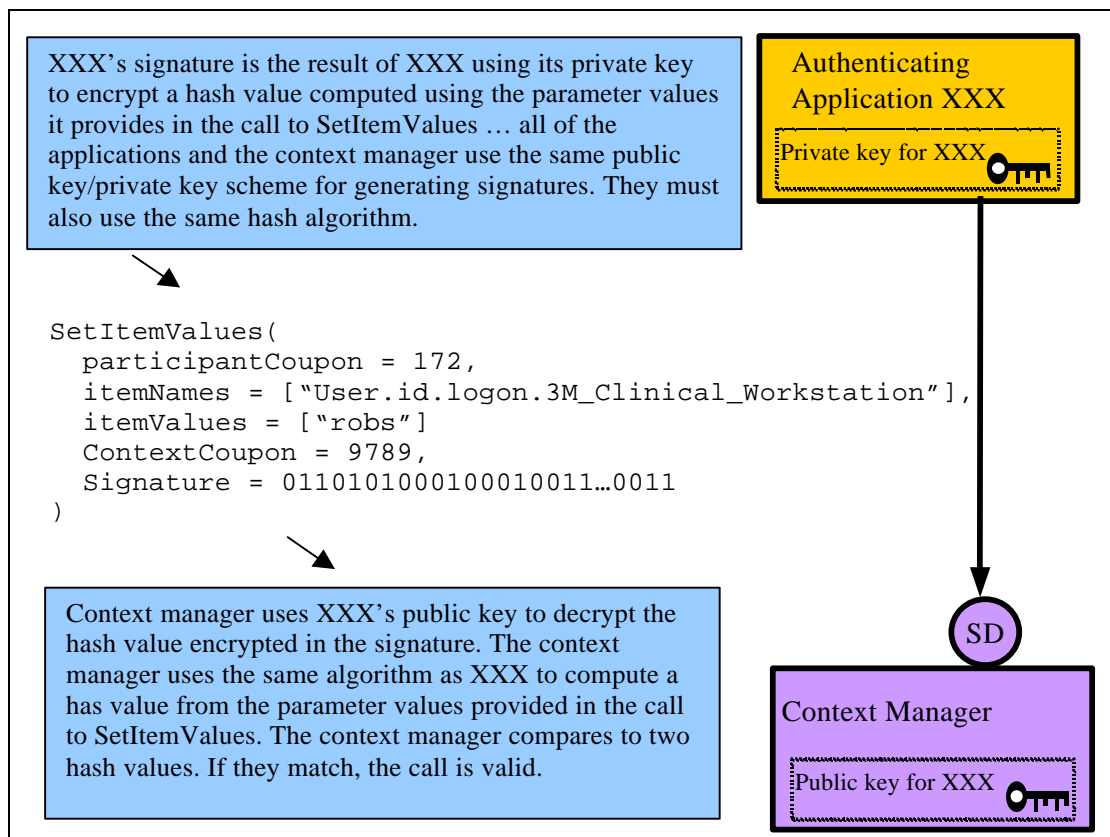


Figure 21: Forming Signature Using Method Parameters

¹⁰ This coupon denotes the current context change transaction, not the application. Each context change coupon is unique over the execution lifetime of a particular context manager.

10.3.3 Computing a Digital Signature

Secure hash algorithms use a character string as the representation of the data value upon which a hash value is to be computed. Therefore, parameter values that are to be protected from tampering during a method invocation must be converted to character strings. These strings must then be concatenated to form a single string. It is the concatenated string that is used to compute the hash value.

The rules for concatenation are as follows. These rules take into account the fact that the mapping of CMA interfaces to specific technologies may alter the order in which method parameters are declared and/or may require additional technology-specific parameters. The rules ensure that the process for creating signatures is invariant across technologies:

- The architectural specification for each method that is to be signed will define which method parameters must be protected from tampering, and are therefore to be used in formulating the signature.
- The architectural specification for each method that is to be signed will define the order in which the string representations of the parameters are to be concatenated.
- The string representation of an array parameter starts with the first element in the array and ends with the last element in the array.
- A parameter or array element whose value is *null* or *empty* is omitted from the string.
- An array that does not contain any elements (i.e., the array length is zero) is omitted from the string.
- Delimiters are not required because there is no need to parse the string.

For example, the concatenated string that might be produced based upon the example in Figure 21: Forming Signature Using Method Parameters would look like:

```
172User.id.logon.3M_Clinical_Workstationrobs9789
```

In another example, where the value of the context item “logon” is null, the concatenated string would look like:

```
172User.id.logon.3M_Clinical_Workstation9789
```

In a final example, where the context items are:

- User.id.logon.3M_Clinical_Workstation = “robs”
- User.co.GivenName = “Robert Seliger”

1 The concatenated string would look like:

2 172User.id.logon.3M_Clinical_WorkstationUser.co.GivenNameRobt Seliger9789
3

4

5 The rules for representing various data types as character strings are specified in Section
6 11.2.9, Representing Basic Data Types as Strings.

7 Finally, once the hash value has been computed, encrypting the hash value with the sender's
8 private key generates the digital signature.

9 **10.3.4 Public Key Distribution**

10 Public key distribution is the process by which an entity, such as the context manager, makes
11 its public key available to the other entities, such as an application, that need to use the key.
12 This process must ensure that a receiving entity can reliably establish the identity of the entity
13 that created the key. If this is not accomplished then it is possible for a rogue entity to
14 impersonate a valid entity by representing the valid entity's public key as its own.

15 In contrast, private keys are not distributed, but remain the secret of the owner of the
16 corresponding public key. A discussion about protecting private keys appears in Section
17 10.3.4.3, Protecting Private Keys.

18 There are a variety of ways that keys can be distributed, including via a certificate authority.
19 However, the approach chosen for the CMA minimizes the amount of infrastructure that is
20 required to create a User Link solution, yet is upwards compatible with more elaborate
21 approaches.

22 Specifically, public keys are exchanged as part of a dynamic process that occurs each time a
23 User Link-enabled application¹¹ or User Link component is launched. This approach enables a
24 high-degree of security while minimizing the effort and cost to develop and deploy User Link
25 solutions.

26 A two-step binding process is used to dynamically distribute an application's public key. The
27 process depends upon the use of secret passcodes that are assigned to user Link-enabled
28 applications (specifically, applications that are capable of being designated for authenticating
29 users) and User Link components. An application or component uses its passcode to prove its
30 identity when it presents its public key. A passcode is a complex, arbitrary alphanumeric
31 string.

¹¹ Not all applications need a public key. Applications that need public keys are those that are designated for authenticating users, and those that use the authentication repository.

A passcode is not actually transmitted when a secure binding is established. Instead, a secure hash function is used to produce a message authentication code. A message authentication code is a secure hash value produced from a data stream that consists of data that is openly communicated between two parties, and “secret” data that they both know but do not openly communicate. In the CMA, a passcode serves as the shared secret.

The binding process involves a “bindee” and a “binder.” In order to bind, a bindee must have a passcode. Both the bindee and the binder must have knowledge of the passcode. The means for providing the bindee and binder with a passcode are not specified in the CMA. However, requirements and guidelines are described in Section 10.3.4.1, Passcode Generation Requirements.

The following table describes the relationships between User Link-enabled applications and User Link components in terms of the secure binding process:

Bindee	Binder
Context Participant Application	Context Manager
Context Participant Application	Authentication Repository
Mapping Agent	Context Manager

The bindee initiates the binding process with the binder. The bindee assumes it knows the identity of the binder, but will prove the binder’s identity as part of the binding process. Similarly, the binder will establish the identity of the bindee as part of the binding process.

The following interactions then occur:

1. The bindee symbolically identifies itself to the binder. The binder uses this information to locate the binder’s copy of the bindee’s passcode. The passcode is not transmitted by the bindee.
2. The binder sends back its public key, and a message authentication code. This code is a secure hash value computed from a data stream formulated from the binder’s public key and the binder’s copy of the bindee’s passcode.
3. The bindee uses the public key it has received and its copy of its passcode to formulate a data stream from which it also computes a secure hash value. (The hash algorithm it uses must be the same as the one that the binder used.) The bindee compares the resulting hash value to the message authentication code. If the two match, then the

binder is who it claims to be and the public key received by the bindee indeed belongs to the binder.

4. The bindee again identifies itself to the binder and sends its public key, along with a new message authentication code. This code is a secure hash value computed from a data stream formulated from the bindee's public key and the bindee's copy of its passcode.
5. The binder uses the public key it has received and its copy of the bindee's passcode to formulate a data stream from which it also computes a secure hash value. (The hash algorithm it uses must be the same as the one that the bindee used.) The binder compares the resulting hash value to the message authentication code. If the two match, then the bindee is who it claims to be and the public key received by the binder indeed belongs to the bindee.

An application requires a passcode for binding with the context manager. This passcode is a secret known only to the application and the context manager.

An application also requires a passcode for binding with the authentication repository. This passcode is a secret known only to the application and the authentication repository. An application that binds to both the context manager and the authentication repository shall use different passcodes for each binding.

10.3.4.1 Passcode Generation Requirements

Passcodes are similar to passwords used by people. However, because passcodes are only used by computer programs, they can be much longer and complex than passwords typically are. This makes passcodes extremely hard to guess, even when brute force techniques are employed.

An application passcode shall be a character string comprised of no less than one hundred (128) characters and no greater than two-hundred fifty-six (256) characters. A passcode shall only be comprised of alphanumeric characters, as well as the underscore (_) and dash (-) characters. A passcode shall not contain white space (e.g., tabs, spaces). A passcode shall be arbitrary but shall not contain any words or phrases.

An application's passcode may be generated such that the same passcode is used for every instance of the application everywhere. This is the least secure means of generating passcodes, because a security breach affects every instance of the application.

An application's passcode may be generated such that the same passcode is used for every instance of the application at a particular site. This is a moderately secure means of generating passcodes, because a security breach is at least limited to a particular site.

1 An application's passcode may be generated such that a unique passcode is used for each
2 desktop upon which the application is used. This is the most secure means of generating
3 passcodes because a security breach is limited to a single desktop. This is the recommended
4 approach.

5 ***10.3.4.2 Protecting Passcodes***

6 Passcodes must remain secret. There are numerous ways in which this can be achieved. The
7 specific approach is left as an implementation decision for applications and the various context
8 management components.

9 However, the following approach is recommended for applications. The assumption is that any
10 application that is used to authenticate users probably uses a server to maintain user account
11 and authorization information. The application might be organized using a client/server
12 architecture, or a web server architecture.

13 The principle challenge is how to create an application such that the portion of the application
14 that serves as a context participant has a secure means to store and retrieve its passcode. In the
15 case of client/server systems, an approach could be to store the passcode on each clinical
16 desktop upon which the client has been loaded. In web systems, an approach could be to
17 transmit the passcode from the web server to the desktop. Both of these approaches introduce
18 substantial security risks that would require great effort to defend against.

19 An alternative is for an application to store its passcode in a server, where it can be more
20 readily protected (including literally placed under lock and key). This could be the
21 application's database server, or it could be a separate server whose specific role is to securely
22 maintain passcodes.

23 The server would never actually transmit the passcode. Rather, it would be responsible for
24 verifying message authentication codes received by the application. It would also be
25 responsible for computing the application's message authentication code.

26 In this approach, the server must be able to authenticate the identity of the application. The
27 server must also be sure that the data it sends and receives from the application is not tampered
28 with while it is in transit. This implies that the application must have the means for
29 establishing a trusted relationship with the server in a manner somewhat akin to the
30 relationship the application establishes with the context manager or authentication repository.

31 There are many ways in which the necessary relationship can be implemented. However,
32 because this relationship does not involve interoperation between applications, and because the
33 optimal approach depends heavily upon the architecture and design of the application, a single
34 approach is not specified. Instead, the approach for the server-based maintenance of an
35 application's passcode is left as an application design exercise.

10.3.4.3 *Protecting Private Keys*

The key distribution process described in Section 10.3.4, Public Key Distribution, does not prescribe when keys are created. However, once created, a private key must remain the secret of its owner for as long as it is in use.

It is possible to statically create a public key / private key pair for an application or component. However, this approach requires the use of a persistent store within which the public key / private key pair are housed when the application or component is not executing. If such a store were used, it would need to be defended against security attacks. This can be accomplished, but at the cost of adding complexity to applications or components.

The recommended alternative approach is for an application or component to dynamically create its key pair when launched. This enables the keys to be kept in memory, and avoids the complexity of using a persistent store. While it is conceivable that an in-memory private key could be accessed by an intruder, most contemporary operating systems enable a process to prevent other processes from reading its memory.

10.3.5 System Configuration Requirements

The system configuration capabilities necessary in order to deploy a User Link system are summarized as follows:

- The context manager shall provide a means for entering the symbolic names of the applications that have been designated for authenticating users. It shall be possible to establish these names on a per-desktop basis for each site. It shall not be possible for anyone but the site's system administrator to modify the names known to a context manager.
- The context manager shall provide a means for entering the symbolic name and corresponding passcode for each application that has been designated for authenticating users at a particular site. This process shall be performed such that the passcode remains a secret known only to the application, the context manager, and perhaps the system administrator who conveys the information from the application to the context manager.
- The context manager shall provide a means for entering the symbolic name and corresponding passcode for the user mapping agent used at a particular site. This process shall be performed such that the passcode remains a secret known only to the user mapping agent, the context manager, and perhaps the system administrator who conveys the information from the application to the context manager.
- The authentication repository shall provide a means for entering the symbolic name and corresponding passcode for each application that uses the authentication repository at a particular site. This process shall be performed such that the passcode

remains a secret known only to the application, the authentication repository, and perhaps the system administrator who conveys the information from the application to the authentication repository.

- Applications capable of being designated for authenticating users, and the user mapping agent, shall provide a means of either obtaining a passcode or for entering a passcode. This process shall be performed such that the secret passcode remains a secret known only to the application or user mapping agent, the context manager, and perhaps the system administrator who conveys the information from the application or user mapping agent to the context manager.

There are numerous ways in which these capabilities can be implemented. It is beyond the scope of the CMA to specify these capabilities. The specific approaches are left as an implementation decision for applications and the various context management components.

10.3.6 Defending Against Replay Attacks

In a replay attack, an intruder captures valid messages that have been previously communicated and retransmits them at a later time in the hope of violating a system.

For example, an intruder might capture a message that enables a user to log on. Even though the intruder might not be able to read the message (it might be encrypted), the intruder might be able to “replay” the message at later time in order to gain access to the system. In this case, the intruder would be able to log on as the user whose actions resulted in the transmission of the original message.

The general approach for defending against replay attacks is to include a “nonce” in each message. The nonce is simply a number that is different each time a message is sent, and is used in computing the hash value for a message. The recipient of a message can keep track of nonces it has seen, and simply reject messages that contain previously seen nonces.

In the CMA, context change coupons in conjunction with the recommend approach of dynamically-generated public key/private key pairs (see Section 10.3.4.3, Protecting Private Keys) defend against replay attacks.

A context change coupon serves as a nonce whose uniqueness is ensured while a context management system is active (i.e., from the time the first participant joins to the time the last participant leaves). Dynamically-generated keys ensure that signed messages can only be authenticated while a context management system is active. Signed messages from earlier activations of the system are meaningless. Together, the use of context change coupons as nonces and dynamically generated keys provide a strong defense against replay attacks.

10.4 Trust Relationships

This section specifies application and component behaviors for realizing the chain of trust.

10.4.1 Trust Between Applications and Context Manager

A User Link-enabled application shall obtain a reference to the context manager's principal interface from the interface reference registry. The application shall interrogate this interface to obtain a reference to the context manager's SecureBinding interface.

A User Link-enabled application shall establish a secure binding with the context manager, per Section 10.3.4, Public Key Distribution, after it has joined the common context system but before it instigates any user context change transactions. This ensures that the application:

- is communicating with the real context manager,
- has obtained the real context manager's public key,
- has provided the context manager with its public key.

A User Link-enabled application shall create a digital signature to sign the context manager methods it invokes in order to set context data that includes user subject context items. This enables the context manager to authenticate the application, and to ensure the integrity of the communicated context data items.

The context manager shall create a digital signature to sign return values it communicates to an application whenever these values include user subject context items. This enables the application to authenticate the context manager, and to ensure the integrity of the communicated context data items.

All other interactions between applications and the context manger do not need to follow these rules.

10.4.2 Trust Between Context Manager and User Mapping Agent

The user mapping agent shall obtain a reference to the context manager's principal interface from the interface reference registry. The user mapping agent shall interrogate this interface to obtain a reference to the context manager's SecureBinding interface.

The user mapping agent shall establish a secure binding with the context manager, per Section 10.3.4, Public Key Distribution, before it maps any user context data. This ensures that the user mapping:

- is communicating with the real context manager,
- has obtained the real context manager's public key,

- has provided the context manager with its public key.

The user mapping agent shall create a digital signature to sign the context manager methods it invokes in order to set context data that includes user subject context items. This enables the context manager to authenticate the user mapping agent, and to ensure the integrity of the communicated context data items.

The context manager shall create a digital signature to sign return values it communicates to the user mapping agent whenever these values includes user subject context items. This enables the user mapping agent to authenticate the context manager, and to ensure the integrity of the communicated context data items.

All other interactions between the context manager and the user mapping agent do not need to follow these rules.

10.4.3 Trust Between Applications and Authentication Repository

A User Link-enabled application shall obtain a reference to the authentication repository's principal interface from the secure registry. The application shall interrogate this interface to obtain a reference to the authentication repository's SecureBinding interface.

A User Link-enabled application shall establish a secure binding, with the authentication repository, per Section 10.3.4, Public Key Distribution, after it has joined the common context system but before it instigates any user context change transactions. This ensures that the application:

- is communicating with the real authentication repository,
- has obtained the real authentication repository's public key,
- has provided the authentication repository with its public key.

A User Link-enabled application shall create a digital signature to sign the authentication repository methods it invokes in order to set user authentication data. This data shall also be encrypted by a means chosen by the application, and then encrypted again upon communication using the authentication repository's public key. The repository shall decrypt the data using its private key only when it needs to service a valid application request to retrieve the data. The repository shall never decrypt the data from its application-specific encrypted form.

This enables the authentication repository to authenticate the application, to ensure the integrity of the communicated authentication data, to keep the authentication data confidential when it is communicated, and to defend against intrusions into the repository to obtain user authentication data.

The authentication repository shall create a digital signature to sign user authentication data it communicates to an application. User authentication data that is communicated back to an application shall remain encrypted as it was when provided by the application. This data shall be encrypted again upon communication using the application's public key.

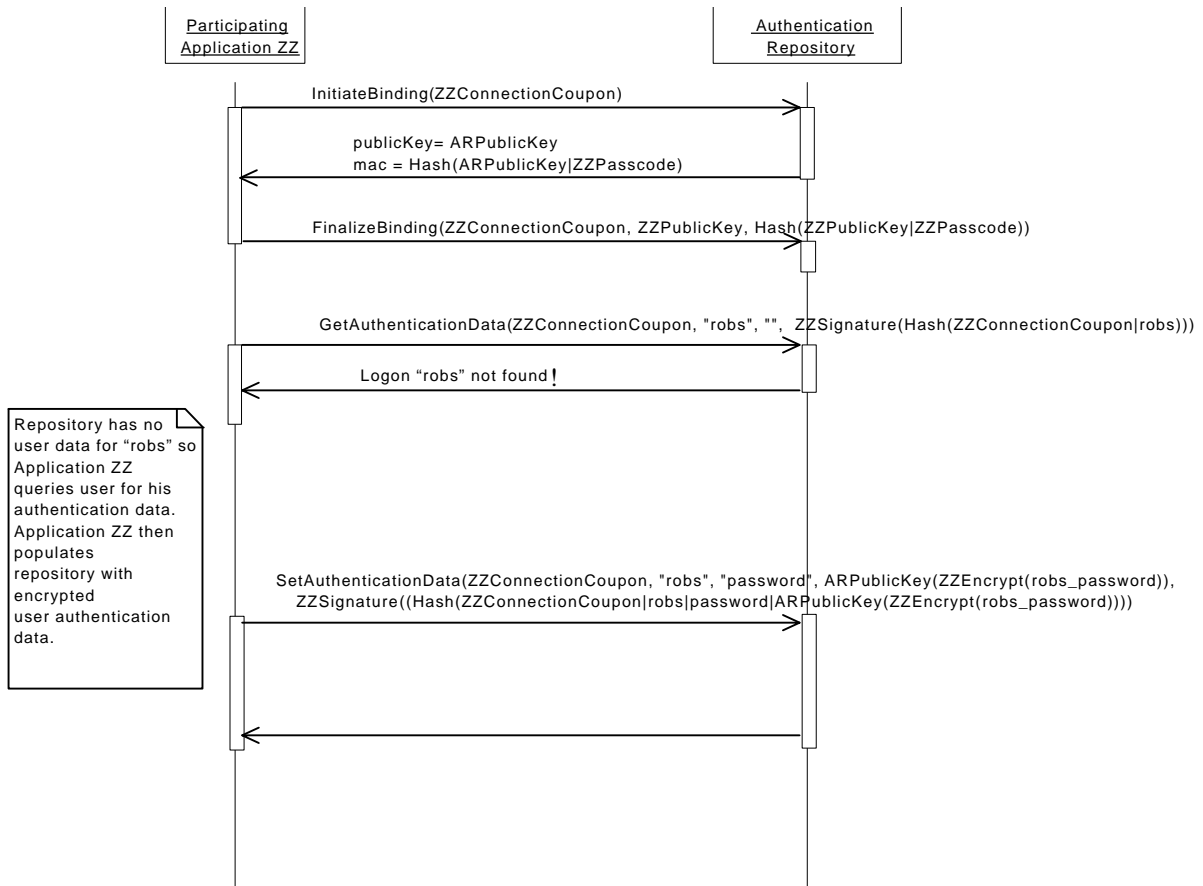
This enables the application to authenticate the authentication repository, to keep the authentication data confidential when it is communicated, and to ensure the integrity of the communicated user authentication data.

All other interactions between applications and the authentication repository do not need to follow these rules.

10.5 Chain of Trust Interactions

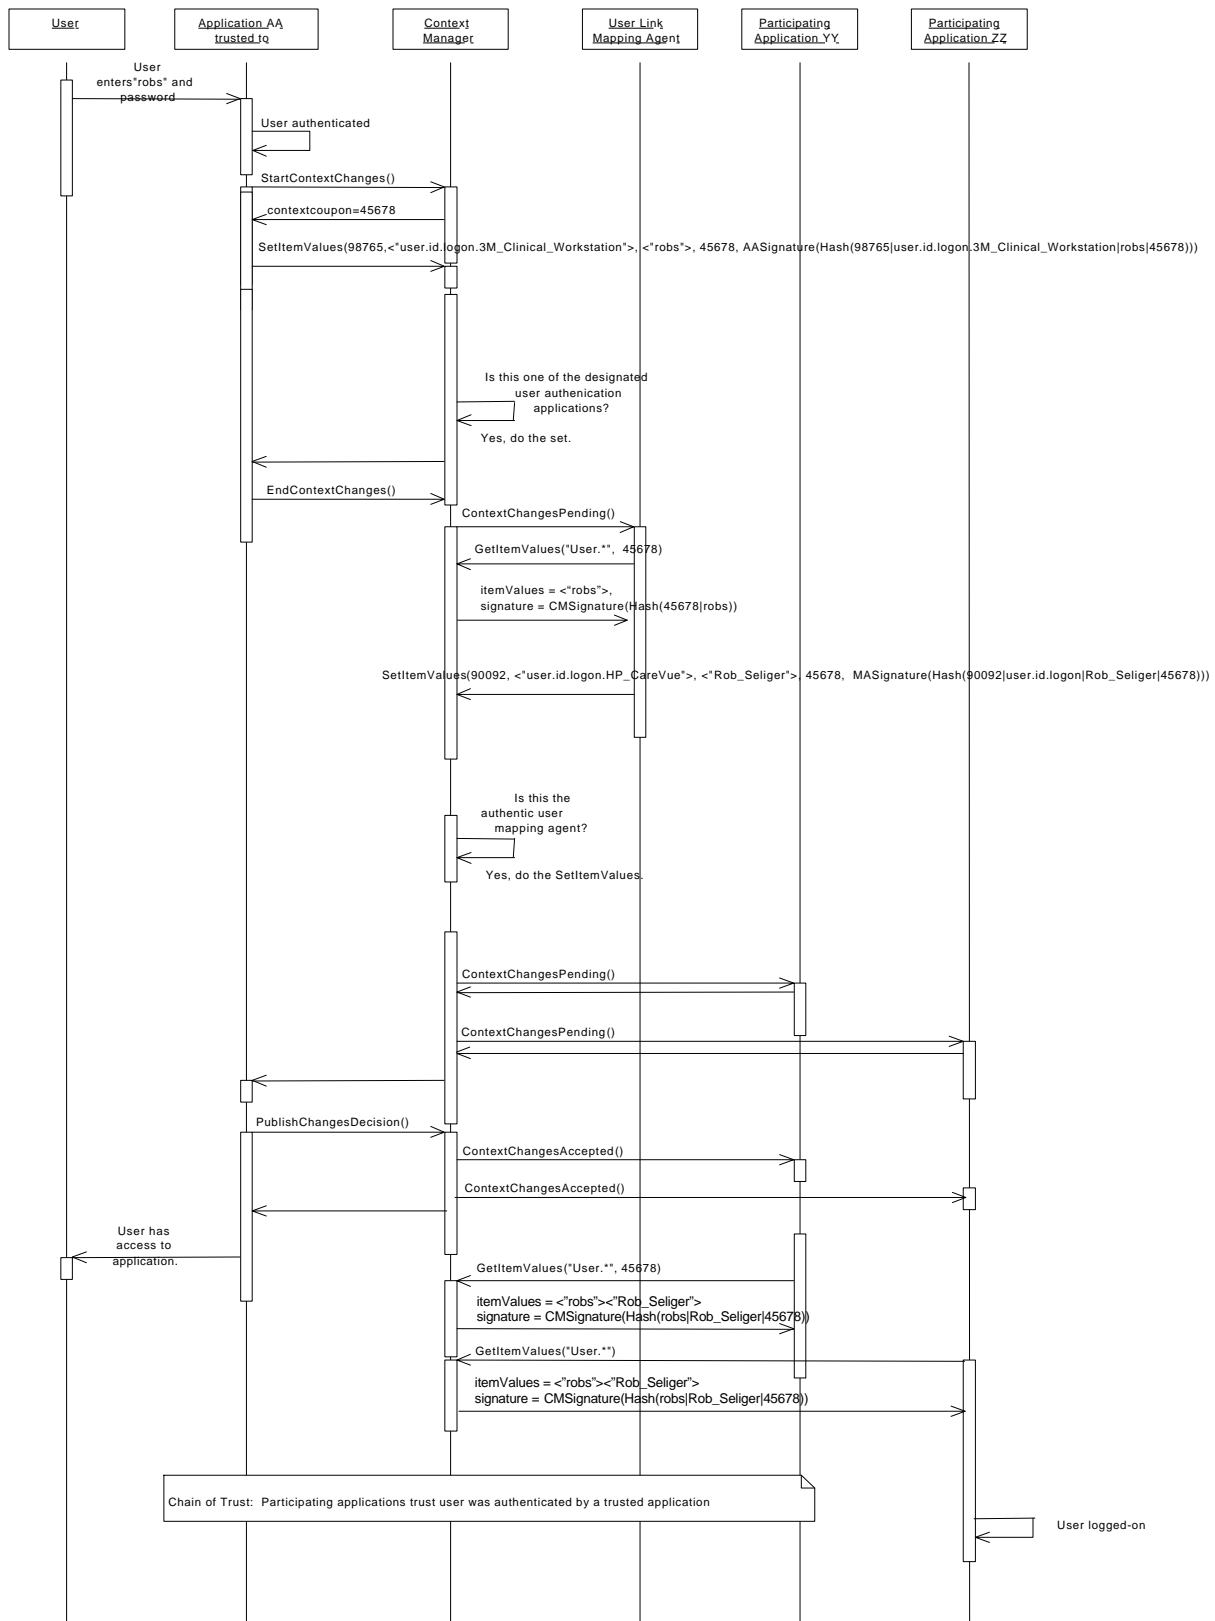
The detailed interactions for several use cases involving the chain of trust are illustrated below. A description for how to interpret the notation used in these diagrams appears in Appendix I. The following additional notation is used:

- The character “|” indicates the concatenation of two strings, for example, “**qrs|xyz**” to form “**qrsxyz**”.
- **XXSignature(a|b|c)** indicates the digital signature for XX. The signature is formed by applying a one-way hash function to the parameter values **a**, **b**, and **c**, and then encrypting the resulting hash value using XX's private key.
- **XXPublicKey(abcd)** indicates that the data “**abcd**” is encrypted using the public key for XX.
- **XXEncrypt(abcd)** indicates that the data “**abcd**” is encrypted using an encryption scheme chosen by XX.
- **Hash(abcd)** indicates a value produced by applying a one-way hash function to the data “**abcd**”.
- The abbreviation **ZZ** represents application ZZ, **CM** represents the context manager, **AR** represents the authentication repository, and **MA** represents the user mapping agent.



1
2

3 **Interaction Diagram 17: Populating Authentication Repository with User Authentication Data**



1
2
3

Interaction Diagram 18: User Link Context Change Transaction

11 Interface Definitions

It is assumed that an underlying technology infrastructure that supports distributed objects is used to implement a common context system, although a specific technology is not assumed. However, the capabilities of Microsoft's COM-based Automation technology are considered as a baseline. This implies that the architecture must work well within the constraints of Microsoft Automation, including issues that pertain to performance and supported data types.

An abstract set of CMA component interface definitions is described below. These interfaces are defined using a precise and concise interface definition language (IDL) created for specifying the CMA. This IDL is not meant to be a comprehensive interface specification language. Only the capabilities that are required for specifying CMA component interfaces are included in the IDL.

A CMA-specific IDL is used because existing interface specification languages have direct or indirect ties to specific technologies. For example, OMG's IDL implies that the interfaces are implemented using CORBA-based technology. Microsoft's MIDL requires that the interfaces are implemented using COM/DCOM technology. The use of these specification languages confuses and possibly compromises the technology-neutrality of the CMA specification.

Experience has shown that the interface constructs represented in IDL defined below can be easily mapped to interfaces that can be implemented using a specific technology such as ActiveX, CORBA, Java, or HTTP. The mapping for each specific technology appears in a separate Context Management specification document.

11.1 Interface Definition Language

The interface definition language (IDL) used in this document enables specifying the following facts about a component interface:

- The interface's symbolic name.
- The set of component properties and methods that can be accessed via the interface.
- The name and data type of each property, and optional restrictions (e.g., read-only).
- The names and data types for each method's input and outputs.
- The names and data content for each method's exceptions.

The IDL also defines a set of simple data types and the capability to represent sequences of these types.

In the following sections, IDL reserved words are shown in bold font. Identifiers are shown in italics. An identifier is an alphanumeric string that starts with an alphabetic character.

11.1.1 Interface Definition Body

The body of an interface definition creates a lexical scope distinct from all other interface definitions. The body of an interface is specified as:

```
interface interfacename { ... }
```

Interfacename is the symbolic name of the interface. The curly brackets delimit the scope of the interface's body.

The body of an interface begins with the declaration of any exceptions that can be raised by methods defined for the interface. The details of declaring exceptions are discussed later.

The properties that can be accessed through the interface are listed next. A property is a data value that can be read or set via the interface:

```
datatype propertyname
```

Datatype is the data type for the property. The type is one of the simple types defined below, as denoted by the appropriate IDL reserved word.

Propertyname is the symbolic name of the property. A property's name must be distinct as compared to the names of other properties, methods, and exceptions defined within the same lexical scope.

Properties can also be sequences. Sequences are described below.

Properties can be restricted to read-only:

```
readonly datatype propertyname
```

The value of a read-only property can be read, but not set, via the interface.

Finally, the methods are listed:

```
methodname inputs ( ... ) outputs ( .... ) exceptions ( ... )
```

Methodname is the symbolic name of the method. A method's name must be distinct as compared to the names of other properties, methods, and exceptions defined within the same lexical scope.

The method's inputs, outputs, and exceptions follow the method's name. If a method does not have any inputs, outputs, or exceptions, then only white space should appear between the appropriate set of parentheses.

1 Each input and output is defined as:

2 *datatype name*

3

4 *Datatype* is the data type for the input or output. The type is one of the simple types defined
5 below, as denoted by the appropriate IDL reserved word. In an actual interface definition, the
6 appropriate IDL reserved word is used to indicate the type. Inputs and outputs can also be
7 sequences. Sequences are described below.

8 *Name* is the symbolic name of the input or output. The name of inputs for a method must be
9 distinct for the method. The name of each output for a method must be distinct for the method.

10 Multiple inputs and outputs are separated by a comma.

11 Exceptions are listed only by their name. Multiple exceptions are separated by a comma.

12 11.1.2 Simple Data Types

13 The following simple data types are supported. The reserved words used to indicate each type
14 are shown:

byte	Eight uninterpreted bits
short	16-bit signed integer
long	32-bit signed integer
float	32-bit floating point number
double	64-bit floating point number
boolean	Indicates true, or false
string	A string of characters
date	A specific year/month/day/time, with a precision of one second, and including the time zone
type	An enumeration that denotes each of these data types (except <i>type</i>) as well as the special types <i>null</i> (valid value not known) and <i>empty</i> (data type not known)
variant	A tagged union of all of these data types (including <i>type</i> and <i>variant</i>)

15

16 The concrete representations of these data types are not defined. They depend upon the
17 interface implementation technology.

11.1.3 Exception Declaration

An exception declaration introduces an exception that can be raised by one or more of the methods defined for the interface within whose lexical scope the exception declaration appears. Each exception declaration indicates the exception name and an optional set of data values. The name denotes the exception and the data values provide additional run-time information about the reason for the exception.

An exception declaration is specified as:

```
exception name { ... }
```

Name is the symbolic name of the exception. An exception's name must be distinct as compared to the names of other properties, methods, and exceptions defined within the same lexical scope.

Exception data values are specified as:

```
datatype name ;
```

Datatype is the data type for the exception value. The type is one of the simple types defined above, as denoted by the appropriate IDL reserved word. In an actual interface definition, the appropriate IDL reserved word is used to indicate the type. Exception values can also be sequences. Sequences are described below.

Name is the symbolic name of the exception value. The name of each value for an exception must be distinct for the exception.

11.1.4 Sequences

A sequence is a single-dimensional vector of sequential data values. Each data value is denoted by an index whose type is **long**. The values for these indices are sequential. The value of the first index is not specified; this value depends upon the interface implementation technology.

A sequence with no restrictions on the quantity of values it can contain is specified as:

```
datatype[ ]name
```

Datatype is the data type of the values in the sequence. The type is one of the simple types defined above, as denoted by the appropriate IDL reserved word. *Name* is the name of the property, input or output, or exception data value.

A sequence with restrictions on the quantity of values it can contain is specified as:

```
datatype[quantity] name
```

Quantity is a numeric value that indicates the maximum quantity of values that the sequence can contain. A sequence may contain less than this quantity. The means by which the quantity of values in a sequence is determined depends upon the interface implementation technology.

11.1.5 Interface References

An interface reference enables access to a specific interface to a specific instance of a component that implements the interface. The interface reference data type represents an interface reference. The type of a property, method input, method output, and exception data value can be an interface reference:

```
interfacename name
```

Interfacename is the name of the interface that the reference represents. *Name* is the name of the property, input or output, or exception data value.

11.1.6 Principal Interface

The reserved word **Principal** is the interface name for a component's principal interface. The role of a component's principal interface is discussed in Section 6.1, **Component and Interface Concepts**. The type of a property, method input, method output, and exception data value can be an interface reference to a principal interface:

```
Principal name
```

Name is the name of the property, input or output, or exception data value.

11.1.7 Qualifying Names

In the IDL there is never a case in which the names of properties, methods, and exceptions defined in one lexical scope are referenced in another lexical scope. However, when documenting the interfaces it can be useful to indicate the scope within which a particular property, method, or exception name has been defined.

The convention for doing so is to formulate a qualified name comprised of the name of the interface within whose scope the property, method, or exception of interest was defined, followed by a pair of colons (::) followed by the name of the property, method, or exception, for example:

```
ContextManager::JoinCommonContext
```

denotes the method JoinCommonContext as defined for the interface ContextManager.

11.2 Interface Implementation Issues

This section describes requirements that all CMA interface implementations must respect.

11.2.1 NotImplemented Exception

In the event that a method is not implemented, the exception `NotImplemented` shall be raised. This exception can be raised, for example, when a method has been deprecated and is no longer implemented by a CMA component. This exception can implicitly be raised by any method defined using CMA IDL and need not be explicitly declared.

11.2.2 GeneralFailure Exception

In the event that a method cannot be properly performed due to an error or failure condition, and an explicitly defined exception does not appropriately represent the situation, then the exception `GeneralFailure` shall be raised. This exception might be raised, for example, when a CMA component is unable to complete a computation due to an internal error. This exception can implicitly be raised by any method defined using CMA IDL and need not be explicitly declared.

11.2.3 Coupon Representation

A participant coupon is a 32-bit integer, represented as the CMA IDL data type **long**, that is assigned by a common context manager to denote each application that joins a common context system. An application is assigned a participant coupon when it joins a common context system. It subsequently uses the coupon to identify itself when performing methods on the context manager.

A context coupon is a 32-bit integer that is assigned by a common context manager to denote each context change transaction. Each time a new transaction is started a new coupon is assigned by the context manager to denote the transaction. Applications use a context coupon to denote the transaction of interest.

Participant coupons shall have unique values for the duration of a common context session (i.e., from the time the first application joins to the time the last application leaves). Context coupons shall also have unique values for the duration of a common context system.

The distinguished value of 0 shall never be assigned as a participant coupon value or as a context coupon value.

11.2.4 Format for Application Names

Several interfaces require that an application provide a CMA IDL **string** that contains a symbolic name for the application. This string is generally used to distinguish one application from another.

1 This string shall only be comprised of alphanumeric characters, blank spaces (no tabs), and the
2 underscore (`_`) character. The string shall neither begin nor end with a blank space.

3 Additionally, an application that is capable of allowing multiple instances of itself to execute
4 on the same desktop shall append to the end of its symbolic name the number-score character
5 (`#`) followed by a string that distinguishes one instance of the application from another.

6 The composition of the appended string is not specified, as long as no two running instances of
7 the application running on a particular desktop use the same appended string at the same time.
8 The appended string shall only be comprised of alphanumeric characters, blank spaces (no
9 tabs), as well as the underscore (`_`) character. The appended string shall neither begin nor end
10 with a blank space.

11 Character case is not considered when comparing application names.

12 An example of this convention is:

```
13     "3M Clinical Workstation#0"  
14     "3M Clinical Workstation#1"  
15     "3M Clinical Workstation#2"
```

16
17 Application names formed as such shall be interpreted as representing the same logical
18 application (e.g., "3M Clinical Workstation") while also representing distinct running
19 instances of the application (i.e., three instances of "3M Clinical Workstation").

20 **11.2.5 Extraneous Context Items**

21 Context participants shall robustly deal with the situation in which context data items that they
22 do not recognize are nevertheless part of the common context. This might occur, for example,
23 in a system comprised of context participants that have been implemented using different
24 versions of the CMA data definition specifications. A participant implemented using an earlier
25 version of these specifications might not recognize context items defined in subsequent versions
26 of the specifications. Context participants shall simply ignore context data items whose names
27 they do not recognize.

28 Similarly, context managers shall allow any context data item for any CMA-defined subject to
29 be part of the context, as long as the name for the item is properly formatted.

30 **11.2.6 Forcing the Termination of a Context Change Transaction**

31 The context manager may need to force the termination of a context change transaction when it
32 appears that the instigator of the transaction has failed before completing the transaction.
33 Specifically, it is recommended that any context manager method that can result in the
34 `ContextManager::TransactionInProgress` exception being thrown should first explicitly confirm
35 that the transaction instigator is still alive.

Most context manager implementations will employ a timer to monitor the activity of a transaction instigator. If the instigator does not perform the necessary operations on the context manager's interfaces in a timely manner, it can be inferred that the instigator has failed. The method `ContextParticipant::Ping` is defined to enable the context manager to probe a context participant to determine its liveness. The context manager may additionally confirm the liveness of a context participant using technology-specific mechanisms.

The duration of these timers, and the use of confirmation techniques, are implementation-dependent.

The context manager shall clean up after the failure of the instigator by performing the following actions:

1. The coupon assigned by the manager for the transaction is invalidated.
2. The transaction-specific version of the context data is discarded.
3. The coupon and context data associated with the most recently committed transaction are unaffected.
4. The context manager's internal state is set to indicate that there is no longer a transaction in progress.

Additional actions depend upon when the context manager determines that the instigator has failed, as described in Table 3: Handling Transaction Instigator Failure.

Instigator fails ...	Leaving systems in the following state ...	Context manager cleans-up by ...
before ending the transaction (see <code>ContextManager::EndContextChanges</code>)	a context change transaction is in progress, although surveying has not yet been performed	performing the actions described above
after ending the transaction but before publishing its decision to accept or cancel the changes (see <code>ContextManager::PublishChangesDecision</code>)	a context change is in progress and the surveyed participants are waiting for the survey decision	publishing the fact that the context changes have been canceled and then performing the actions described above

Table 3: Handling Transaction Instigator Failure

11.2.7 Character-Encoded Binary Data

Several of the CMA component interfaces use CMA IDL **string** parameters that contain character-encoded binary data. The following representation of character-encoded binary data shall be applied for all such parameters¹².

Each byte of data shall be represented by two printable characters. The four high bits of the byte (i.e., the high octet) shall be represented by the left character. The four low bits of the byte (i.e., the low octet) shall be represented by the right character.

An array of bytes shall be represented by character-encodings such that the left most character-encoded byte in the string represents the data byte at lowest array index. The encoding follows sequentially, such that the right most character-encoded byte in the string represents the data byte at the highest array index.

Each four bits of data (i.e., an octet) is represented by an alphanumeric character as follows:

Data (Octet)	Character
0000	0
0001	1
0010	2
0011	3
0100	4
0101	5
0110	6
0111	7
1000	8
1001	9
1010	A or a
1011	B or b
1100	C or c
1101	D or d
1110	E or e
1111	F or f

The actual character set that is employed is technology-specific. Each of the HL7 context management technology mapping specification documents indicates the character set that is used for a particular technology-specific implementation.

¹² Base64 encoding was not selected as a character-encoding scheme for binary data, as the added compression offered by the scheme is of minimal advantage for the CMA, wherein only relatively small quantities of binary data are transmitted.

Binary data that is character-encoded as a string shall not include white space or any other characters other than the ones shown in the table above. The character-encoded string is not case sensitive. An example of binary data character-encoded per these conventions is:

Binary Data: 00000001 11101001 11000111 1000010

Character-Encoded String: 01E9C782

11.2.8 Representing Message Authentication Codes, Signatures and Public Keys

Message authentication codes, digital signatures, public keys are used as input or output parameters for several of the methods defined for CMA component interfaces. The CMA IDL data type for each of these parameters is **string**. Each string contains character-encoded binary data, encoded per Section 11.2.7, Character-Encoded Binary Data.

The binary data that is encoded is technology-specific. Each of the HL7 context management technology mapping specification documents indicates the binary data types needed for a particular technology-specific implementation. It is necessary that both the sender and receiver of a message authentication code, digital signature, or public key agree upon the format of the underlying binary data type, and the algorithms used to create the data. The method `SecureBinding::InitiateBinding`, defined in 11.3.7.1, enables this agreement to be established.

11.2.9 Representing Basic Data Types as Strings

Several of the CMA component interfaces use input or output parameters whose values are computed from the string representations of data values of various types. For example, digital signatures are computed from a one-way hash value, which is, in turn, computed from a string formed by concatenating a list of data values, each of which is represented as a string.

The following data types shall be represented as character strings using the formats described in Table 4: Character Representations for Basic Data Types.

1

Type	String Representation	Comments
boolean	0, if false 1, if true	
short	dddd, where d is a numeric character representing a decimal digit and the number of characters depends upon the value of the number.	Leading minus sign (-dddd) if number is negative. No plus sign if positive.
long	Same as for short .	
date	yyyy/mm/dd hh:mm:ss	
string	As is.	Case is preserved.
float	dddd.dddd, where d is a numeric character representing a decimal digit. The number of digits before the decimal point depends on the magnitude of the number, and the number of digits after the decimal point depends on the precision.	Leading minus sign (-dddd.dddd) if number is negative. No plus sign if positive.
double	Same as float, except that there can be more digits.	
byte	bb, where b is a hexadecimal digit. The byte is represented as unsigned.	Lower case for alphabetic characters that represent hex digits (i.e., a, b, c, d, e, f).

2 **Table 4: Character Representations for Basic Data Types**

3

4 The actual character set that is employed is technology-specific. Each of the HL7 context
5 management technology mapping specification documents indicates the character set that is
6 used for a particular technology-specific implementation.

7 **11.2.10 Pre-Defined Mapping Agent Coupons**

8 A participant coupon value is pre-defined for each type of mapping agent. In general, a
9 negative coupon value denotes a mapping agent, as opposed to a context participant
10 application. The following values are currently allocated:

Mapping Agent	Coupon Value
Patient	-1
User	-2
Reserved for future	-3 through -500

1

2 Pre-defined coupon values are used for mapping agents because they do not explicitly join the
3 context system. Instead, a mapping agent is implicitly “pulled” into the context system each
4 time a context change transaction occurs, when the context manager performs the mapping
5 agent method `MappingAgent::ContextChangesPending`. See Section 11.3.6.1,
6 `ContextChangesPending`.

7 However, agents such as the user mapping agent need to know their participant coupon values
8 prior to the first context change transaction. For example, the user mapping agent needs to
9 establish a secure binding with the context manager before it can set user context items. In
10 order to establish this binding, the user mapping agent must present the context manager with
11 its coupon (see Section 11.3.7.1, `InitiateBinding`). By having a priori knowledge of its coupon
12 value, the user mapping agent can establish its secure binding whenever it decides to, up until
13 the time it actually attempts to set the context.

11.3 Interfaces

This section specifies the methods for each of the CMA interfaces.

11.3.1 AuthenticationRepository (AR)

```

interface AuthenticationRepository {
    exception AuthenticationFailed { string reason; }
    exception UnknownApplication {}
    exception UnknownConnection {}
    exception LogonNotFound { string logonName; }
    exception UnknownDataFormat { string dataFormat; }

    Connect
    inputs(string applicationName)
    outputs(long connectionCoupon)
    raises()

    Disconnect
    inputs(long connectionCoupon)
    outputs()
    raises(UnknownConnection)

    SetAuthenticationData
    inputs(long connectionCoupon, string logonName, string dataFormat,
           string userData, string appSignature)
    outputs()
    raises(UnknownConnection, AuthenticationFailed)

    DeleteAuthenticationData
    inputs(long connectionCoupon, string logonName, string dataFormat,
           string appSignature)
    outputs()
    raises(UnknownConnection, AuthenticationFailed, LogonNotFound,
           UnknownDataFormat)

    GetAuthenticationData
    inputs(long connectionCoupon, string logonName, string dataFormat,
           string appSignature)
    outputs(string userData, string repositorySignature)
    raises(UnknownConnection, AuthenticationFailed, LogonNotFound,
           UnknownDataFormat)
}

```

11.3.1.1 Connect

This method enables an application to establish a connection with the authentication repository. An application must have a connection before it can set or get user authentication data.

The value of the input *applicationName* is a succinct string that contains the application's symbolic name. The output *connectionCoupon* is the value of a connection coupon that the application can subsequently use to denote itself when performing other authentication repository methods.

The value of input *applicationName* is used by the authentication repository to determine the passcode for an application. The passcode is needed when an application establishes a secure binding with the authentication repository (see Section 11.3.7 SecureBinding (SB)). Multiple instances of an application can connect to the authentication repository using the same name. Each instance of the application will be assigned a unique connection coupon. Each instance of the application will need to establish a secure binding with the repository.

The value of the input *applicationName* is also used by the authentication repository to store/retrieve the user authentication data within the repository.

The exception *UnknownApplication* is raised if the input *applicationName* does not represent an application currently known to the authentication repository.

11.3.1.2 Disconnect

This method enables an application to disconnect from the authentication repository. An application shall disconnect before it terminates. The value of the input *connectionCoupon* denotes the application.

The exception *UnknownConnection* is raised if the input *connectionCoupon* does not denote an application currently connected to the authentication repository.

11.3.1.3 SetAuthenticationData

This method enables an application to store authentication data for a particular user's logon name within the authentication repository. This method also enables an application to update authentication data for a particular user's logon name that it has already stored in the repository.

The value of the input *connectionCoupon* denotes the application, the value of the input *logonName* is a user's logon name, the value of the input *userData* is the application-specific data used to authenticate the user, and the value of the input *appSignature* is the application's digital signature. This signature enables the authentication repository to authenticate that the request to set the authentication data came from the application denoted by the value of *connectionCoupon*, and that the values of *connectionCoupon*, *logonName*, *dataFormat*, and *userData*, were not tampered with between the time they were sent and were received.

Concatenating the string representations of the following inputs in the order listed shall form the data from which a message digest is computed by the application:

- 1 • *connectionCoupon*
- 2 • *logonName*
- 3 • *dataFormat*
- 4 • *userData*

5 An application shall compute its digital signature by encrypting the message digest with its
6 private key.

7 The value of the input *dataFormat* is an application-defined string that is used when an
8 application needs to maintain multiple forms of authentication data for a user (e.g., password,
9 thumbprint image, etc.). If only one form of authentication data is needed, this string can be
10 empty. Multiple calls of *SetAuthenticationData* are required to set different forms of
11 authentication data for a particular user. The value of *dataFormat* for each call should indicate
12 the form of authentication data to be stored.

13 The value of the input *userData* contains user authentication data that has been encrypted by
14 the application using an encryption technique chosen by the application. This data is character-
15 encoded per Section 11.2.7, Character-Encoded Binary Data. The structure of the encoded
16 binary data is application-dependent and is not specified.

17 The exception *UnknownConnection* is raised if the input *connectionCoupon* does not denote an
18 application that is currently connected to the repository.

19 The exception *AuthenticationFailed* is raised if the process of authentication determines that
20 the signature is not the signature for the application denoted by the input *connectionCoupon* or
21 that the input parameter's values have been tampered with.

22 **11.3.1.4 DeleteAuthenticationData**

23 This method enables an application to delete from the authentication repository some or all of
24 the authentication data that it previously stored for a particular logon name. Both the logon
25 name and the associated authentication data are deleted.

26 The value of the input *connectionCoupon* denotes the application and the value of the input
27 *logonName* is the logon name to be deleted.

28 The value of the input *dataFormat* is an application-defined string that is used when an
29 application maintains multiple forms of authentication data for a user (e.g., password,
30 thumbprint image, etc.) within the repository. If this string is empty, then all of the forms of
31 authentication data stored for the user are deleted. If this string is not empty, then just the
32 denoted form of authentication data is deleted.

1 The value of the input *appSignature* is the application's digital signature.

2 Concatenating the string representations of the following inputs in the order listed shall form
3 the data from which a message digest is computed by the application:

4 *connectionCoupon*

5 *logonName*

6 *dataFormat*

7 An application shall compute its digital signature by encrypting the message digest with its
8 private key.

9 This signature enables the authentication repository to authenticate that the request to delete
10 the authentication data came from the application denoted by the value of *connectionCoupon*,
11 and that the values of *coupon*, *logonName*, and *dataFormat* were not tampered with between
12 the time they were sent and were received.

13 The exception UnknownConnection is raised if the input *connectionCoupon* does not denote an
14 application that is currently connected to the repository.

15 The exception AuthenticationFailed is raised if the process of authentication determines that
16 the signature is not the signature for the application denoted by the input *connectionCoupon* or
17 that the input parameter values have been tampered with.

18 The exception LogonNotFound is raised if user authentication data corresponding to the logon
19 name denoted by the input *logonName* does not reside in the repository.

20 The exception UnknownDataFormat is raised if the form of authentication data denoted by the
21 input *dataFormat* is not found in the repository.

22 **11.3.1.5 GetAuthenticationData**

23 This method enables an application to retrieve from the authentication repository the
24 authentication data previously stored for a particular user's logon name. The value of the input
25 *connectionCoupon* denotes the application, the value of the input *logonName* is a user's logon
26 name, and the value of the input *appSignature* is the application's digital signature.

27 This signature enables the authentication repository to authenticate that the request to get the
28 authentication data came from the application denoted by the value of *connectionCoupon*, and
29 that the values of *coupon*, *logonName*, and *dataFormat* were not tampered with between the
30 time they were sent and were received.

1 Concatenating the string representations of the following inputs in the order listed shall form
2 the data from which a message digest is computed by the application:

- 3 • *connectionCoupon*
- 4 • *logonName*
- 5 • *dataFormat*

6 An application shall compute its digital signature by encrypting the message digest with its
7 private key.

8 The value of the input *dataFormat* is an application-defined string that is used when an
9 application needs to maintain multiple forms of authentication data for a user (e.g., password,
10 thumb-print image, etc.). If only one form of data is used, this string can be empty. Multiple
11 calls of *GetAuthenticationData* are required to get different forms of authentication data for a
12 particular user. The value of *dataFormat* for each call should indicate the form of
13 authentication data to be retrieved.

14 The value of the output *userData* is the application-specific data used to authenticate the user.
15 The output *userData* remains encrypted, as it was when it was stored by the application using
16 *SetAuthenticationData*.

17 The output *userData* shall be used as the data from which a message digest is computed by the
18 application. The authentication repository shall compute its digital signature by encrypting the
19 message digest with its private key.

20 This signature enables the application to authenticate that the authentication data returned by
21 this method came from the authentication repository and that the value of *userData* was not
22 tampered with between the time it was sent and was received.

23 The exception *UnknownConnection* is raised if the input *connectionCoupon* does not denote an
24 application that is currently connected to the repository.

25 The exception *AuthenticationFailed* is raised if the process of authentication determines that
26 the signature is not the signature for the application denoted by the input *connectionCoupon* or
27 that the input parameter values have been tampered with.

28 The exception *LogonNotFound* is raised if user authentication data corresponding to the logon
29 name denoted by the input *logonName* does not reside in the repository.

30 The exception *UnknownDataFormat* is raised if the form of authentication data denoted by the
31 input *dataFormat* is not found in the repository.

11.3.2 ContextData (CD)

```

1  interface ContextData {
2
3      exception UnknownParticipant { long participantCoupon; }
4      exception UnknownItemName { string itemName; }
5      exception BadItemNameFormat { string itemName; string reason }
6      exception BadItemType { string itemName; type actual;
7          type expected; }
8      exception BadItemValue { string itemName; variant itemValue;
9          string reason; }
10     exception NameValueCountMismatch {long numNames; long numValues }
11     exception ChangesNotPossible {}
12     exception ChangesNotAllowed {}
13     exception InvalidContextCoupon {}
14
15     GetItemNames
16     inputs(long contextCoupon)
17     outputs(string[] names)
18     raises(InvalidContextCoupon)
19
20     DeleteItems
21     inputs(long participantCoupon, string[] itemNames,
22         long contextCoupon)
23     outputs()
24     raises(NotInTransaction, UnknownParticipant, InvalidContextCoupon,
25         BadItemNameFormat, UnknownItemName, ChangesNotPossible,
26         ChangesNotAllowed)
27
28     SetItemValues
29     inputs(long participantCoupon, string[] itemNames,
30         variant[] itemValues, long contextCoupon)
31     outputs()
32     raises(NotInTransaction, UnknownParticipant, InvalidContextCoupon,
33         NameValueCountMismatch, BadItemNameFormat, BadItemType,
34         BadItemValue, ChangesNotPossible, ChangesNotAllowed)
35
36     GetItemValues
37     inputs(string[] itemNames, boolean onlyChanges, long contextCoupon)
38     outputs(variant[] itemValues)
39     raises(InvalidContextCoupon, BadItemNameFormat, UnknownItemName)
40 }
41

```

11.3.2.1 GetItemNames

This method enables a participant in a common context system to obtain the names of the common context items.

This method can be performed outside the scope of a context change transaction. In this case, the value of the input *contextCoupon* must denote the most recently committed transaction.

The output *itemNames* is a sequence containing the item names that represent the state of the common context as it was when the most recently committed transaction was completed.

This method can also be performed within the scope of a context change transaction that is currently in progress. In this case, the input *contextCoupon* must denote the current transaction. The output *itemNames* contains the item names that represent the state of the common context as it has been established so far by the transaction. The output *itemNames* is empty (i.e. zero elements) until a participant explicitly sets item values via the `ContextData::SetItemValues` method within the scope of the transaction.

The exception `InvalidContextCoupon` is raised if the input *contextCoupon* does not denote the most recently committed transaction or the transaction currently in progress.

11.3.2.2 *DeleteItems*

Note: This method has been deemed extraneous and is being deprecated. In a future version of this specification context managers may chose to not implement this method even though it remains part of the ContextData interface definition.

This method enables an application in a common context system to remove an item from the set of common context items. The application or mapping agent denotes itself with its participant coupon as the value of the input *participantCoupon*. The value of the input *contextCoupon* must denote the current context change transaction, as obtained by the instigator of the transaction when it performed the `ContextManager::StartContextChanges` method.

The exception `NotInTransaction` is raised if there is no change transaction currently in progress.

The exception `UnknownParticipant` is raised if the input *participantCoupon* does not denote an application or mapping agent that is currently a participant in the common context system.

The exception `InvalidContextCoupon` is raised if the context coupon parameter does not denote the transaction currently in progress.

The exception `BadItemNameFormat` is raised if the format of an item named for deletion does not conform to the specification for the item in the relevant HL7 context management data definition specification document.

The exception `UnknownItemName` is raised if one or more of the items named for deletion is not the name of an item in the context as it stands under the current transaction.

The exception `ChangesNotPossible` is raised if the `ContextData::DeleteItems` method is invoked after the `ContextManager::EndContextChanges` method has already been invoked for the transaction currently in progress.

The exception `ChangesNotAllowed` is raised by `ContextData::DeleteItems` if a mapping agent attempts to delete context items.

11.3.2.3 *SetItemValues*

This method enables an application or mapping agent in a common context system to set the value of one or more common context items. The application or mapping agent denotes itself with its participant coupon as the value of the input *participantCoupon*. The names of the context items to be set are contained in the input sequence *itemNames*. The values for each of these items are contained in the input sequence *itemValues*. The i^{th} element in *itemValues* is the value for the item named by the i^{th} element in *itemNames*.

If an item named in *itemNames* is not currently an item in the common context, it will be added. The data type for a newly added item is the same as the data type of the element in *itemValues* that contains the item's value.

This method can only be performed within the scope of a context change transaction. The value of the input *contextCoupon* must denote the current transaction.

The exception `NotInTransaction` is raised if there is no change transaction currently in progress.

The exception `UnknownParticipant` is raised if the input *participantCoupon* does not denote an application or mapping agent that is currently a participant in the common context system.

The exception `InvalidContextCoupon` is raised if the input *contextCoupon* does not denote the transaction currently in progress.

The exception `NameValueCountMismatch` is raised if the number of items in the input *itemNames* does not match the number of items in the input *itemValues*.

The exception `BadItemNameFormat` is raised if the format of an item named for deletion does not conform with the specification for the item in the relevant HL7 Context Management Data Definition Specification.

The exception `BadItemType` is raised if the data type for one or more of the items whose value is to be set is not the same as the expected data type.

The exception `BadItemValue` is raised if the data value for one or more of the items whose value is to be set is determined to be unacceptable. This exception is used by context manager implementations that enforce semantic constraints on the common context. Not all context manager implementations will do this.

The exception `ChangesNotPossible` is raised if the `ContextData::SetItemValues` method is invoked by an application after the `ContextManager::EndContextChanges` method has already

been invoked for the transaction currently in progress. (This exception is *not* raised if a mapping agent invokes `ContextData::SetItemValues` after `ContextManager`.)

The exception `ChangesNotAllowed` is raised if a mapping agent attempts to set a value for a context item for which a value has already been set by the application that instigated the context change transaction.

11.3.2.4 *GetItemValues*

This method enables a participant in a common context system to obtain the value of one or more context items.

When the value of the input *contextCoupon* denotes the most recently committed transaction, the item values that are returned represent the state of the common context as it existed when the transaction was completed. This is true even if there is currently a new transaction in progress.

When the value of the input *contextCoupon* denotes the transaction currently in progress, the item values that are returned represent the state of the common context as it has been established so far by the transaction.

The items of interest are indicated in the input sequence *itemNames*. These names can be fully-qualified item names, which means that the all of the fields for an item's name are explicitly specified (e.g., "Patient.Id.MRN.St_Elsewhere_Hospital").

Alternatively, a wild card represented by an asterisk (*) can be used in place of a specific string for any of the item name fields except for the subject field (which is lexically the first field on the left). The wild card enables a participant to obtain one or more items without having to specify complete item names.

If a wild card is used, it must appear in only the last field specified in the item name string (which is lexically the last field on the right). Additional field names and/or wild cards must not appear after a wild card (i.e., lexically to the right of the wild card). Examples of properly formatted items names include:

“Patient.*” matches all of the identifier and corroborating items for the patient subject

“Patient.Id.*” matches all of the patient identifier items

“Patient.Id.MRN.*” matches all of the patient identifiers that are site-specific medical record numbers

Conversely, “Patient.Id.*.*” and “Patient.Id.*.St_Elsewhere_Hospital” are examples of improperly formatted item names.

The sequence output *itemValues* contains the values of all of the items whose names match the set of names specified in the input *itemNames*. A specific item's value will be included at most once in *itemValues*, even if its name matches more than one of the names specified in *itemNames*. For example, even if *itemNames* includes the names:

“Patient.Id.MRN.St_Elsewhere_Hospital”

and:

“Patient.Id.*”

the value for the item named “Patient.Id.MRN.St_Elsewhere_Hospital” will be included only once in *itemValues*.

The elements in the sequence *itemValues* alternate between the complete name of an item (represented as a string) and the corresponding item value (represented by the appropriate data type). For example, if several context data items are returned, then the first element in the list is the name of the first item, the second element in the list is the value of the first item, the third element in the list is the name of the second item, the fourth element in the list is the value of the second item, and so on.

The input *onlyChanges* enables a participant to instruct the context manager to filter which items it returns no matter what names were specified. When the value of *onlyChanges* is true, then the items that are returned are limited to only the context subjects whose items were set by the most recently committed context change transaction, or by the transaction in progress, as indicated by the value of *contextCoupon*.

For example, if *onlyChanges* is true, *contextCoupon* denotes the most recently committed context change transaction, and *itemNames* includes the name:

“Patient.Id.*”

but items in the patient subject were not set during the transaction, then the output *itemValues* will not contain any items pertaining to the patient subject.

The exception *InvalidContextCoupon* is raised if the input *contextCoupon* does not denote the most recently committed transaction or the transaction currently in progress.

The exception *BadItemNameFormat* is raised if the format of an item named for deletion does not conform with the specification for the item in the relevant HL7 context management data definition specification.

The exception *UnknownItemName* is raised if one or more of the items named is not the name of an item in the context as it stands under the current transaction.

11.3.3 ContextManager (CM)

```

1  interface ContextManager {
2
3      exception AlreadyJoined {}
4      exception UnknownParticipant { long participantCoupon; }
5      exception TransactionInProgress { string instigatorName; }
6      exception NotInTransaction {}
7      exception InvalidTransaction { string reason; }
8      exception TooManyParticipants { long howMany; }
9      exception ChangesNotEnded {}
10     exception AcceptNotPossible {}
11     exception UndoNotPossible {}
12     exception InvalidContextCoupon {}
13
14     readonly long MostRecentContextCoupon
15
16     JoinCommonContext
17     inputs(ContextParticipant contextParticipant,
18         string applicationName, boolean survey, boolean wait)
19     outputs(long participantCoupon)
20     raises(AlreadyJoined, TooManyParticipants, TransactionInProgress)
21
22     LeaveCommonContext
23     inputs(long participantCoupon)
24     outputs()
25     raises(UnknownParticipant)
26
27     StartContextChanges
28     inputs(long participantCoupon)
29     outputs(long contextCoupon)
30     raises(UnknownParticipant, TransactionInProgress,
31         InvalidTransaction)
32
33     EndContextChanges
34     inputs(long contextCoupon)
35     outputs(boolean noContinue, string[] responses)
36     raises(InvalidContextCoupon, NotInTransaction,
37         InvalidTransaction)
38
39     UndoContextChanges
40     inputs(long contextCoupon)
41     outputs()
42     raises(InvalidContextCoupon, NotInTransaction, UndoNotPossible)
43
44 
```



```

1      PublishChangesDecision
2      inputs(long contextCoupon, string decision)
3      outputs()
4      raises(NotInTransaction, InvalidContextCoupon, ChangesNotEnded,
5             AcceptNotPossible)
6
7      SuspendParticipation
8      inputs(long participantCoupon)
9      outputs()
10     raises(UnknownParticipant)
11
12     ResumeParticipation
13     inputs(long participantCoupon, boolean wait)
14     outputs()
15     raises(UnknownParticipant, TransactionInProgress)
16 }
17

```

11.3.3.1 *MostRecentContextCoupon*

This read-only property contains the value of the context coupon that represents the most recently committed changes to the common context data. Even if there is a change transaction in progress, this property's value represents the previously committed transaction. If no transactions have been committed, the value of this property is 0.

11.3.3.2 *JoinCommonContext*

This method enables an application to join a common context system. The application must provide a reference to its `ContextParticipant` interface as the value of the input *contextParticipant*.

The value of the input *applicationName* is a succinct string that can be used to easily and clearly identify the application to the user (see Section 11.2.4, Format for Application Names). This string must be unique relative to the other applications that have already joined the common context system.

If an application subsequently attempts to establish a secure binding with the context manager (see Section 11.3.7 SecureBinding (SB)), then this string is used by the context manager to determine the passcode for an application.

The application can also indicate whether it wants to participate in context change surveys (the value of the input *survey* indicates true), or that it just wants to be informed when a context change has been accepted (the value of the input *survey* indicates false).

An application can only join a common context system between context change transactions. If no transaction is in progress, the application is able to immediately join the context change system.

If a transaction is in progress and the value of the input *wait* indicates true, this method will block until the transaction completes. It is recommended that an application that is willing to wait also display a message to the user indicating that it is attempting to join a common context system. If a transaction is in progress and the value of the input *wait* indicates false, this method immediately raises the exception *TransactionInProgress*.

The output *participantCoupon* is the value of the participant coupon that the application can subsequently use to denote itself when performing other *ContextManager* methods.

The exception *AlreadyJoined* is raised if an application with the same name as the value of *applicationName* has already joined the context.

The exception *TooManyParticipants* is raised if the context manager is unable to accommodate an additional common context participant.

11.3.3.3 *LeaveCommonContext*

This method enables an application that is a participant in a common context system to leave the system. The application denotes itself using its participant coupon as the value of the input *participantCoupon*. Once this method returns, the application is free to terminate.

In order to avoid a deadlock condition, this method does not block. If this method was allowed to block, it would be possible for an application to block while the context manager was attempting to perform a method on the application's *ContextParticipant* interface. For single-threaded applications, this could cause a deadlock.

Consequently, if a context change transaction is in progress when this method is called, the application may still be notified about the context change even though it has left the common context. The application is free to ignore this notification or may not even be capable of responding. The context manager will robustly handle the failure of an application to respond.

The exception *UnknownParticipant* is raised if the input *participantCoupon* does not denote an application that is currently a participant in the common context system.

11.3.3.4 *StartContextChanges*

This method enables an application to indicate that it wants to start changing the common context. The application denotes itself with its participant coupon as the value of the input *participantCoupon*. A context change transaction is initiated. Actual changes to the context data are conducted via the *ContextData* interface. The output *contextCoupon* is the value of the context coupon that has been assigned by the context manager to denote the change transaction.

The context manager will automatically terminate context change transaction if it does not detect activity on its `ContextData` interface or if the `ContextManager::EndContextChanges` method is not performed in a timely manner. The amount of time that the manager will wait before terminating the transaction depends upon the manager's implementation.

The exception `UnknownParticipant` is raised if the input *participantCoupon* does not denote an application that is currently a participant in the common context system.

The exception `TransactionInProgress` is raised if a context change transaction is already in progress.

The exception `InvalidTransaction` is raised if a suspended application calls this method.

11.3.3.5 *EndContextChanges*

This method enables the application that instigated a context change transaction to indicate that it has completed its changes to the common context. The value of the input *contextCoupon* denotes the transaction currently in progress. This method initiates the two-step change notification process and returns after the first phase of the notification process is conducted by the context manager. During the first phase, the applications in the common context system are surveyed to determine their ability or willingness to apply the context changes. The `ContextParticipant::ContextChangesPending` method is performed on each application in the survey.

The output *responses* is a sequence of strings that is used to convey the results of the survey to the application that instigated a context change transaction.

If all of the applications surveyed indicate that they are willing to accept the context changes, then the output sequence *responses* is empty (i.e. zero elements) and the output *noContinue* is false. The sequence is empty because there is no useful information to be conveyed about the applications that have accepted, other than the fact that they all accepted. The method `ContextManager::PublishChangesDecision` with the decision *accept* shall be subsequently performed by the instigating application to communicate to the other applications the decision to accept the context changes and to complete the transaction.

If there are surveyed applications that either are unable to provide a response to the survey (e.g., because they are "busy"), or that want to inform the user that work-in-progress might be lost if the context is changed, then the return value contains a string for each such application. The application that invoked this method is expected to display the strings to the user and to obtain guidance about how to proceed.

The output *noContinue* indicates true if the mapping agent invalidated the transaction, or at least one of the surveyed applications is "busy". It is not possible for the user to continue to apply the context change transaction if the value of *noContinue* is true. The only option the

user has is to cancel the change or to disconnect the instigating application from the common context system. For either user decision, the method `ContextManager::PublishChangesDecision` with the decision *cancel* shall be performed by the instigating application.

If the mapping agent has not invalidated the transaction and there are no busy applications (i.e., *noContinue* is false), but there are applications that have conditionally accepted the context changes, the user can instruct the instigating application to apply the context changes anyway, cancel the changes, or to disconnect from the common context system.

The method `ContextManager::PublishChangesDecision` with the decision *accept* shall be subsequently performed by the instigating application to complete the transaction if the user decides to apply the context changes.

The method `ContextManager::PublishChangesDecision` with the decision *cancel* shall be subsequently performed by the instigating application to complete the transaction if the user decides to cancel the context changes or to disconnect the instigating application from the common context system.

The exception `InvalidContextCoupon` is raised if the input *contextCoupon* does not denote the transaction currently in progress.

The exception `NotInTransaction` is raised if there is no change transaction currently in progress.

The exception `InvalidTransaction` is raised if, for each subject whose context data items have been set by the transaction, the context data changes do not include at least one item that is an identifier (e.g., context data for a subject cannot be comprised of just corroborating data).

11.3.3.6 *UndoContextChanges*

This method enables an application to discard any context data changes that it has already made. The context coupon parameter denotes the transaction currently in progress. The current transaction is brought to a close and the context coupon is no longer valid.

The exception `InvalidContextCoupon` is raised if the input *contextCoupon* does not denote the transaction currently in progress.

The exception `NotInTransaction` is raised if there is no change transaction currently in progress.

The exception `UndoNotPossible` is raised if the method `ContextManager::UndoContextChanges` is attempted after the `ContextManager::EndContextChanges` method has been performed during the course of the current transaction.

11.3.3.7 *PublishChangesDecision*

This method enables the application that instigated a context change transaction to inform the other applications in a context system about whether the changes are to be applied or have been canceled. The value of the input *contextCoupon* denotes the transaction currently in progress.

The decision to accept the changes shall be published when the context changes are to be applied. The only times that context changes cannot be applied are when there were applications for which it was not possible to obtain a survey response (e.g., these applications were “busy”) or when a mapping agent invalidates the transaction.

The decision to cancel the changes shall be published when the context changes are to be discarded.

If the decision is to accept the changes, the value of the value of the output *decision* parameter is “accept”. If the decision is to cancel the changes, the value of the output *decision* is “cancel”. The context manager is shall treat these values in a case-insensitive manner.

Once the decision has been published, the change transaction is complete.

The exception *InvalidContextCoupon* is raised if the input *contextCoupon* does not denote the transaction currently in progress.

The exception *NotInTransaction* is raised if there is no change transaction currently in progress.

The exception *ChangesNotEnded* is raised if the method *EndContextChanges* has not yet been performed during the course of the current transaction.

The exception *AcceptNotPossible* is raised if the decision to be published is *accept* but there were applications for which it was not possible to obtain a survey response (e.g., these applications were blocked). The decision *accept* in this case is erroneous. This exception defends against this case should it arise due to an application programming error.

11.3.3.8 *SuspendParticipation*

This method enables an application to indicate that it wants to suspend its active participation in a common context system while remaining registered as a participant. The application denotes itself with its participant coupon as the value of the input *participantCoupon*. It should be apparent to the user that the application is not displaying context-sensitive data, for example, the application might be minimized so that no data display can be seen.

Suspending participation is not the same as leaving the common context. Instead, this method provides an optimization for applications that temporarily do not want to track context

changes. This enables an application to perform computational tasks without being interrupted by context changes.

This method also enables an application to minimize its use of computational resources if it is in a state (e.g., minimized) in which responding to context changes provides no benefit to the user. The application can subsequently resume its participation in the common context via the `ContextManager::ResumeParticipation` method. The application will not be surveyed, nor will it be informed of changes to the common context until the application invokes the `ContextManager::ResumeParticipation` method.

In order to avoid a deadlock condition, this method does not block. If this method was allowed to block, it would be possible for an application to block while the context manager was attempting to perform a method on the application's `ContextParticipant` interface. For single-threaded applications, this could cause a deadlock.

Consequently, if a context change transaction is in progress when this method is called, the application may still be notified about the context change. The application is free to ignore this notification or may not even be capable of responding. The context manager will robustly handle the failure of an application to respond.

This method has no effect if the application has already suspended its participation.

A suspended application cannot instigate a context change transaction.

Context manager implementations are encouraged to periodically confirm that suspended context participants are still running. This is to avoid the situation in which context manager continues to allocate internal resources to a suspended participant that subsequently fails without first informing the context manager that it is leaving the common context system.

This method is an alternative to leaving the common context system. Context managers can be implemented to support a maximum number of participants. If an application leaves a context system, it risks not being able to rejoin. In contrast, by suspending its participation, this possibility is avoided.

The exception `UnknownParticipant` is raised if the input *participantCoupon* does not denote an application that is currently a participant in the common context system.

11.3.3.9 *ResumeParticipation*

This method enables an application to indicate that it wants to resume active participation in a common context system. The application denotes itself with its participant coupon as the value of the input *participantCoupon*. Upon resuming, an application must automatically ensure that it has established synchrony with the current context.

1 The application denotes itself with its participant coupon. This method has no effect if the
2 application did not previously invoke the ContextManager::SuspendParticipation.

3 An application can only resume its participation a common context system between context
4 change transactions. If no transaction is in progress, the application is able to immediately
5 resume participation in the context change system.

6 If a transaction is in progress and the value of the input *wait* indicates true, this method will
7 block until the transaction completes. It is recommended that an application that is willing to
8 wait also display a message to the user indicating that it is attempting to resume participation
9 in a common context system. If a transaction is in progress and the value of the input *wait*
10 indicates false, this method immediately raises the exception TransactionInProgress.

11 The exception UnknownParticipant is raised if the input *participantCoupon* does not denote an
12 application that is currently a participant in the common context system.

13

11.3.4 ContextParticipant (CP)

```

interface ContextParticipant {
    ContextChangesPending
        inputs(long contextCoupon)
        outputs(string decision, string reason)
        raises()

    ContextChangesAccepted
        inputs(long contextCoupon)
        outputs()
        raises()

    ContextChangesCanceled
        inputs(long contextCoupon)
        outputs()
        raises()

    CommonContextTerminated
        inputs()
        outputs()
        raises()

    Ping
        inputs()
        outputs()
        raises()
}

```

11.3.4.1 ContextChangesPending

This method informs a participant in a common context system that a change to the common context data is pending. The value of the input *contextCoupon* denotes the transaction within which the context changes occurred. The participant shall respond with an indication of how it wants to deal with the change:

- Accept the change
- Conditionally accept the change (e.g., because it is in the middle of a task that would cause significant user work to be lost if a context change was allowed)

An application that accepts the changes is willing to apply the new context data if subsequently instructed to do so (by the ContextParticipant::ContextChangesAccepted or ContextParticipant::ContextChangesCanceled methods).

An application that conditionally accepts the changes is also willing to apply the changes, but only after informing the user that the application might lose work that the user is in the midst of performing. The output *reason* shall contain a succinct but informative description of the work that might be lost. (The description should not identify the application as this information

is provided by the application when it joins the common context system.) The application through which the user instigated the context changes is responsible for informing the user of the situation and obtaining the user's decision about how to proceed.

An application that cannot interpret the context data (e.g., does not know who the patient is) should accept the changes. However, the application should clearly indicate to the user (e.g., by displaying a message) that it cannot apply the current context data.

If the response is to accept the changes, the value of the output *decision* is "accept". If the decision is to conditionally accept the changes, the value of the output *decision* "accept_conditional". The context manager shall treat these values in a case-insensitive manner.

If a participant does not respond in a timely manner, it will be interpreted by the context manager as being busy. The amount of time that the manager will wait before determining that an application is busy depends upon the manager's implementation. This method is not performed upon the application that instigated the context changes. Instead, the application is blocked by the manager when it performs `ContextManager::EndContextChanges`.

11.3.4.2 ContextChangesAccepted

This method informs a participant in a common context system that the result of the most recent context change survey was to accept the changes and that the common context data has indeed been changed. The participant can access the context data via the context manager's `ContextData` interface to obtain the changes. The value of the input *contextCoupon* denotes the transaction within which the context changes occurred. This coupon is needed in order to access the context data.

If it is not possible to perform this method on an application because it is busy, the context manager will periodically keep trying until it has successfully performed the method, or until a new context change transaction is initiated. The intervals at which the context manager tries to retry this method is implementation-dependant.

11.3.4.3 ContextChangesCanceled

This method informs a participant in a common context system that a context change transaction has been canceled. The value of the input *contextCoupon* denotes the transaction that has been canceled.

If it is not possible to perform this method on an application because it is busy, the context manager will periodically keep trying until it has successfully performed the method, or until a new context change transaction is initiated. The intervals at which the context manager tries to retry this method is implementation-dependant.

1 **11.3.4.4 CommonContextTerminated**

2 This method informs a participant in a common context system that the system is being
3 terminated. The participant will not be subsequently informed about context changes, nor will
4 it be able to perform common context changes. If the system is re-established, the participant
5 must explicitly rejoin the system before performing the ContextManager::JoinCommon-
6 Context method.

7 **11.3.4.5 Ping**

8 This method provides a means for a context manager to determine whether or not a participant
9 in a common context system is still running. This method shall be implemented by all
10 participants to return immediately. The context manager can then perform this method to probe
11 a participant when its existence is in question.

12 In performing this method, the context manager will be able to indirectly exercise the
13 underlying communications infrastructure. The infrastructure will either indicate that the
14 method was successfully performed, that the method failed because the participant no longer
15 exists, or that the method failed but it cannot be determined whether or not the participant
16 exists. In this last case, the manager shall assume that the participant still exists.

11.3.5 ImplementationInformation (II)

```

interface ImplementationInformation {
    readonly string ComponentName
    readonly string RevMajorNum
    readonly string RevMinorNum
    readonly string PartNumber
    readonly string Manufacturer
    readonly string TargetOS
    readonly string TargetOSRev
    readonly string WhenInstalled
}

```

11.3.5.1 *ComponentName*

This read-only property is the name of the component, specifically, “Patient Link Mapping Agent”.

11.3.5.2 *RevMajorNum*

This read-only property is the major number for the software revision for the component, as assigned by its manufacturer. For example, in the full revision number Z.32, ‘Z’ is the major number and might indicate a particular functional release of the software.

11.3.5.3 *RevMinorNum*

This read-only property is the minor number of the software revision for the component, as assigned by its manufacturer. For example, in the full revision number Z.32, ‘32’ is the minor number and might indicate a particular build of the software.

11.3.5.4 *PartNumber*

This read-only property is the part number that the component’s manufacturer assigned to the component.

11.3.5.5 *Manufacturer*

This read-only property is the name of the organization that developed the component.

11.3.5.6 *TargetOS*

This read-only property is the name of the operating system on which the component is able to execute.

1 **11.3.5.7 *TargetOsRev***

2 This read-only property is the revision of the operating system named in target operating
3 system on which the component is able to execute.

4 **11.3.5.8 *WhenInstalled***

5 This read-only property is the date and time at which the component was installed on its host.

6

7

8

11.3.6 MappingAgent (MA)

```

interface MappingAgent {
    ContextChangesPending
    inputs(long mappingAgentCoupon, Principal contextMgr,
           long contextCoupon)
    outputs(string decision, string reason)
    raises()

    Ping
    inputs()
    outputs()
    raises()
}

```

11.3.6.1 ContextChangesPending

This method informs a mapping agent in a common context system that a change to the common context data is pending. The value of the input *contextCoupon* denotes the transaction within which the context changes occurred. The value of the input *mappingAgentCoupon* is a predefined coupon that denotes the specific type of mapping agent. (See Section 11.2.10, Pre-Defined Mapping Agent Coupons). The value of the input *contextMgr* is an interface reference to the context manager's principal interface. This is so that the mapping agent can easily obtain the context manager interface(s) it needs.

The agent shall respond with an indication of how it wants to deal with the context change:

- The changes are valid
- The changes are invalid

If the changes are valid, then the value of the output *decision* should be "valid". If the changes are invalid, then the value of the output *decision* should be "invalid". The changes should only be declared invalid if the set of identifiers in the proposed context data do not all represent the same patient. If the changes are invalid, then the value of the output *reason* will contain a succinct but detailed string describing why the changes were invalid. Otherwise the value of *reason* is null.

11.3.6.2 Ping

This method provides a means for a context manager to determine whether or not a mapping agent in a common context system is still running. This method shall be implemented by all agents to return immediately. The context manager can then perform this method to probe a mapping agent when the agent's existence is in doubt.

1 In performing this method, the context manager will be able to indirectly exercise the
2 underlying communications infrastructure. The infrastructure will either indicate that the
3 method was successfully performed, that the method failed because the agent no longer exists,
4 or that the method failed but it cannot be determined whether or not the agent exists. In this last
5 case, the manager shall assume that the agent still exists.

11.3.7 SecureBinding (SB)

```

interface SecureBinding {
    exception UnknownBindee {}
    exception UnknownPropertyName { string propertyName; }
    exception BadPropertyType { string propertyName; type actual;
        type expected; }
    exception BadPropertyValue { string propertyName;
        variant itemValue; string reason; }
    exception NameValueCountMismatch {long numNames; long numValues }
    exception ImproperKeyFormat { string reason; }
    exception ImproperMACFormat { string reason; }
    exception BindingRejected { string reason; }
    exception AuthenticationFailed { string reason; }

    InitiateBinding
    inputs(long bindeeCoupon, string[] propertyName,
        variant[] propertyValues)
    outputs(string binderPublicKey, string mac)
    raises(UnknownBindee, NameValueCountMismatch,
        UnknownPropertyName, BadPropertyType, BadPropertyValue,
        BindingRejected)

    FinalizeBinding
    inputs(long bindeeCoupon, string bindeePublicKey,
        string mac)
    outputs()
    raises(UnknownBindee, ImproperKeyFormat, ImproperMACFormat,
        AuthenticationFailed)
}

```

11.3.7.1 InitiateBinding

This method enables a context management component (“bindee”) to initiate the process of establishing a secure binding with another context management component (“binder”). The bindee shall complete the process of establishing a secure binding with the binder by performing the method SecureBinding::FinalizeBinding upon the binder.

A secure binding shall be established by the bindee before it attempts to interact with the binder via methods that entail the use of either the bindee’s or the binder’s digital signature. For example, an application or user mapping agent shall establish a secure binding with the context manager before it attempts to access the context manager in order to set (or, in the future, get) context item values that require the bindee’s digital signature. An application shall establish a secure binding with the authentication repository before attempting to set or get user authentication data from the authentication repository.

This method shall be performed only after the bindee has been provided by the binder with a coupon to denote itself. The value of the input *bindeeCoupon* is this coupon. The value of *bindeeCoupon* depends upon the role bindee and binder, as described below:

Bindee	Binder	Value of <i>bindeeCoupon</i>
Context Participant Application	Context Manager	Participant coupon, obtained by the participant from the context manager via ContextManager::JoinCommonContext.
Context Participant Application	Authentication Repository	Connection coupon, obtained by the participant from the authentication repository via AuthenticationRepository::Connect.
Mapping Agent	Context Manager	Mapping agent coupon, obtained from the context manager when it most recently performed MappingAgent::ContextChangesPending upon the mapping agent.

As part of the process of establishing a secure binding, it is necessary for the bindee and the binder to agree upon the properties of the underlying security algorithms that they will use in subsequent secure interactions. These properties may include the public key / private key scheme, the number of bits used to represent a key, and the type of one-way hash algorithm that is to be used to generate message digests and message authentication codes. The specific properties that must be agreed upon, and the allowed set of values for these properties, are defined in the each of the HL7 context management technology-specific component mapping specification documents.

The value of the input sequence *propertyNames* contains the names of the secure binding-related properties for which the bindee wishes to establish agreement. The values for each of these properties are contained in the input sequence *propertyValues*. The i^{th} element in *propertyValues* is the value for the property named by the i^{th} element in *propertyNames*. The data type for a property is the same as the data type of the element in *propertyValues* that contains the property's value.

The value of the output *binderPublicKey* is the binder's public key, and shall be used by the bindee in all subsequent secure interactions that involve the binder. The value of *binderPublicKey* is character-encoded binary data formed by the binder when it computes its public key / private key pair.

The value of the output *mac* is the message authentication code. This code shall be used by the bindee to prove the identity of the binder, and to ensure that the value of *binderPublicKey* has not been tampered with.

The value of *mac* is character-encoded binary data formed by the binder's computation of a one-way hash value. This hash value is computed using an input string formed by concatenating the bindee's passcode to the end of the character-encoded binary string containing the binder's public key. This passcode is a secret known only to the bindee and the binder. Upon receipt of the output *mac* and *binderPublicKey*, the bindee independently creates the same string as the binder and performs the same hash computation. If the resulting hash value matches the value of *mac*, then the binder shall be considered authentic and the value of *binderPublicKey* shall be considered valid.

The algorithms used to compute *mac* and *binderPublicKey* are technology-specific. The format of these outputs are also technology specific.

The exception *UnknownBindee* is raised if the input *bindeeCoupon* does not denote a context management component currently known to the binder.

The exception *NameValueCountMismatch* is raised if the number of items in the input *propertyNames* does not match the number of items in the input *propertyValues*.

The exception *BadPropertyType* is raised if the data type for one or more of the properties whose value is to be set is not the same as the expected data type.

The exception *BadPropertyValue* is raised if the data value for one or more of the properties whose value is to be set is determined to be unacceptable or incompatible.

The exception *BindingRejected* is raised if the bindee is not authorized to establish a binding with the binder. When this exception is raised by the context manager, it means that the context participant application has not been designated for authenticating users. When this exception is raised by the authentication repository, it means that the repository has not been configured to serve the application.

11.3.7.2 FinalizeBinding

This method enables bindee to finalize the process of establishing a secure binding with a context management component. This method shall be performed by a bindee only after it has successfully performed the method *InitiateBinding* upon a binder. The bindee denotes itself using the same value for the input *bindeeCoupon* that it used when it performed the method *InitiateBinding* upon the binder.

The input *bindeePublicKey* is the bindee's public key, and shall be used by the binder in all subsequent secure interactions that involve the bindee. The value of *binderPublicKey* is

1 character-encoded binary data formed by the bindee when it computes its public key / private
2 key pair.

3 The input *mac* is the message authentication code. This code shall be used by the binder to
4 prove the identity of the bindee, and to ensure that the value of *bindeePublicKey* has not been
5 tampered with.

6 The value of *mac* is character-encoded binary data formed by the bindee's computation of a
7 one-way hash value. This hash value is computed using an input string formed by
8 concatenating the bindee's passcode to the end of the character-encoded binary string
9 containing the bindee's public key. This passcode is a secret known only to the bindee and the
10 binder. Upon receipt of the inputs *mac* and *bindeePublicKey*, the binder independently creates
11 the same string as the bindee and performs the same hash computation. If the resulting hash
12 value matches the value of *mac*, then the bindee shall be considered authentic and the value of
13 *bindeePublicKey* shall be considered valid.

14 The algorithms used to compute *mac* and *bindeePublicKey* are technology-specific. The
15 format of these inputs are also technology specific.

16 The exception UnknownBinding is raised if the input *bindingCoupon* does not denote an
17 bindee currently known to the binder.

18 The exception ImproperKeyFormat is raised if the input *publicKey* is not properly formatted.

19 The exception ImproperMACFormat is raised if the input *mac* is not properly formatted.

20 The exception BindingDenied is raised if the input *mac* does not establish the identity of the
21 bindee and/or the integrity of the input *bindeePublicKey*.

11.3.8 SecureContextData (SD)

```

interface SecureContextData {
    exception UnknownItemName { string itemName; }
    exception BadItemNameFormat { string itemName; string reason }
    exception BadItemType { string itemName; type actual;
        type expected; }
    exception BadItemValue { string itemName; variant itemValue;
        string reason; }
    exception NameValueCountMismatch { long numNames; long numValues }
    exception ChangesNotPossible {}
    exception SignatureRequired {}
    exception AuthenticationFailed { string reason; }

    GetItemNames
    inputs(long contextCoupon)
    outputs(string[] itemNames)
    raises(InvalidContextCoupon)

    SetItemValues
    inputs(long participantCoupon, string[] itemNames,
        variant[] itemValues, long contextCoupon, string appSignature)
    outputs()
    raises(NotInTransaction, InvalidContextCoupon, UnknownParticipant,
        NameValueCountMismatch, BadItemNameFormat, BadItemType,
        BadItemValue, ChangesNotPossible, SignatureRequired,
        AuthenticationFailed)

    GetItemValues
    inputs(long participantCoupon, string[] itemNames,
        boolean onlyChanges, long contextCoupon, string appSignature)
    outputs(variant[] itemValues, string managerSignature)
    raises(InvalidContextCoupon, UnknownParticipant,
        BadItemNameFormat, UnknownItemName, SignatureRequired,
        AuthenticationFailed)
}

```

11.3.8.1 GetItemNames

This method is identical to ContextData::GetItemNames.

11.3.8.2 SetItemValues

This method is similar to ContextData::SetItemValues. The primary difference is that the context participant's digital signature shall be provided as the value of the input *appSignature* when user subject item values are among the items to be set. This signature enables the context manager to authenticate that they came from a designated application or from the real user mapping agent, and that the values were not tampered with between the time they were sent and were received.

A signature is not required when the values for the user subject items are null. This enables any application to set the user context to empty. When a signature is not provided, the value of the input *appSignature* shall be an empty string (“”).

Concatenating the string representations of the following inputs in the order listed shall form the data from which a message digest is computed by the participant:

- *participantCoupon*
- *itemNames* (i.e., All the elements in the order that they appear in the array.)
- *itemValues* (i.e., All the elements in the order that they appear in the array.)
- *contextCoupon*

A participant shall compute its digital signature by encrypting the message digest with its private key.

The exception *SignatureRequired* is raised if the value of *appSignature* is not a digital signature and a signature is required in order to perform this method.

The exception *AuthenticationFailed* is raised if a digital signature is required and provided, but the process of authentication determines that: the application that invoked this method did not previously provide its public key via the interface *SecureBinding*; that the input *appSignature* has been forged; that the input parameter values have been tampered with; that the participant has not been designated for performing user context changes.

11.3.8.3 *GetItemValues*

This method is similar to *ContextData::GetItemValues*. The primary difference is that the context manager’s digital signature shall be provided as the value of the output *managerSignature* when user subject identifier item values are among the items named for retrieval. This signature enables the recipient of the item values to authenticate that they came from the real context manager, and that the values were not tampered with between the time they were sent and were received.

Concatenating the string representations of the following inputs in the order listed shall form the data from which a message digest is computed by the context manager:

- *ItemValues* (i.e., All the elements in the order that they appear in the array.)
- *contextCoupon*

The context manager shall compute its digital signature by encrypting the message digest with its private key.

1 The value of the inputs *participantCoupon* and *appSignature* are not currently used and are
2 defined in anticipation of future uses of this method. In the future, the value of these inputs will
3 enable the context manager to enforce context data access rights as a function of the context
4 participant's identity and the properties of the requested context items, as listed in the input
5 *itemNames*. The value of *participantCoupon* will denote the participant. The value of
6 *appSignature* will be the digital signature of the participant.

7 Until stated otherwise in a future version of this specification, the value of the input
8 *participantCoupon* shall be zero (0). The value of the input *appSignature* input shall be an
9 empty string ("").

10 The exception *SignatureRequired* is raised if the value of *appSignature* is not a digital
11 signature and a signature is required to perform this method.

12 The exception *AuthenticationFailed* is raised if a digital signature is required and provided, but
13 the process of authentication determines that: the application that invoked this method did not
14 previously provide its public key via the interface *SecureBinding*; that the input *appSignature*
15 has been forged; that the input parameter values have been tampered with; that the participant
16 is not allowed to access the requested context items.

17

1 **12 Backwards Compatibility**

2 The HL7 Context Management Architecture specified in this document is fully compatible
3 with the Clinical Context Object Workgroup Patient Link 1.1 Architecture Specification. The
4 CMA is, however, a superset of the CCOW Architecture.

5

Appendix: Diagramming Conventions

There are four types of formal diagrams that are used throughout this document to describe the CMA architecture:


- A use case diagram depicts the actors (human and/or computer-based) and the roles that they play when participating in an interesting scenario.
- A use case interaction diagram illustrates the high-level interactions between the actors that participate in the use case.
- A component architecture diagram depicts components and their interfaces, and indicates the interfaces each component uses for communicating with other components.
- A component interaction diagram illustrates the series of method invocations that components perform on each other in order to implement a particular use case.

The conventions for each of these diagrams are explained below. Many of the conventions were leveraged from Ivar Jacobson's text *Object-Oriented Software Engineering*.[†] In the future, these conventions will be evolved to comply with the Unified Modeling Language specification, which is still being refined[★].

Use Case Diagram

The use case diagramming conventions are:


- A stick figure represents an actor, even if the actor is a computer-based entity, such as an application:




[†] Object-Oriented Software Engineering, Ivar Jacobson, Addison-Wesley, 1994.

[★] Unified Modeling Language Reference Manual, James Rumbaugh, Grady Booch, Ivar Jacobson, Addison-Wesley, 1997.

- 1 • An oval represents a use case. The name of the use case appears next to the oval:




- 4 • An arrow directed from an actor to the use case indicates that the actor participates in
5 the use case. A label near the arrow succinctly describes the actor's role in the use
6 case:



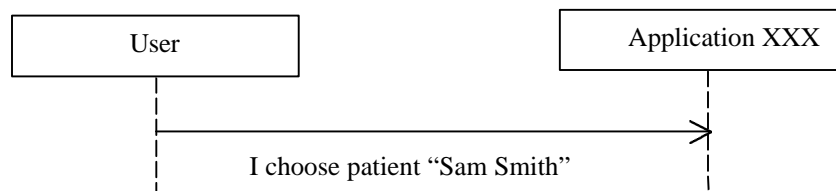
8 **Use Case Interaction Diagrams**

9 The use case interaction diagramming conventions are:

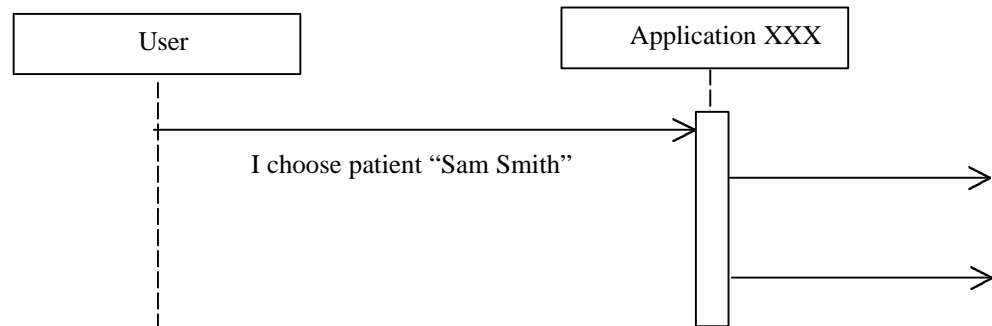
- 10 • The interacting actors are depicted by rectangles labeled with the actor's name,
11 arranged in a horizontal row. A vertical dashed bar descends from each of these
12 rectangles:




- 17 • An interaction that is initiated by an actor is represented as an arrow that emanates
18 from the actor. The arrow terminates on the actor to which the interaction is directed .
19 Each arrow is labeled with a short description of the interaction it represents:



- A vertical bar indicates the start and end of the actions that an actor performs in response to an interaction. These actions may include additional interactions:



- An actor can respond to an interaction. A response is shown as an arrow labeled with an indication of the response:




- The entire set of interaction arrows is temporally ordered, from left to right, top to bottom.


Component Architecture Diagrams

The component architecture diagramming conventions are:


- Each component is depicted as a rectangle. The name of the component appears within the rectangle:



- Each of the interfaces implemented by a component is illustrated as a circle tangent to the rectangle that depicts the component. Each circle is labeled with the name of the interface it represents. Two or three letter abbreviations are typically used:




- A directed arrow connects components that communicate with each other. Arrows emanate from a client component and point to the server components that it uses. Each arrow terminates on the circle representing the specific server component interface that is used. A distinct arrow is used for each interface for each server component that a client component uses:




Component Interaction Diagrams

The component interaction diagramming conventions are:

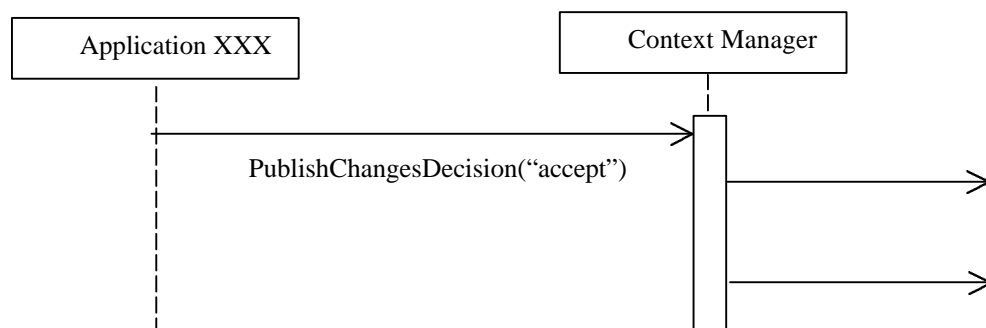
- The interacting components are depicted by rectangles labeled with the component's name, arranged in a horizontal row. A vertical dashed bar descends from each of these rectangles:




- A method that is invoked by a component is represented as an arrow that emanates from the bar and that terminates on the bar for the component that services the method. Each arrow is labeled with the name of the method it represents. Examples of actual parameter values *may* be included for clarity:



- A vertical bar indicates the start and end of the processing that a component performs in response to a method invocation. This processing may itself include method invocations:



- Method return values are indicated when this aids in understanding the use case. A return value is shown as an arrow labeled with an indication of the return value:



- The entire set of method invocation arrows is temporally ordered, from left to right, top to bottom.

Glossary

Accept	An application's response when it is willing to accept the context data changes and to change its internal state accordingly if the changes are published.
Accept-Conditional	An application's response when it is in the midst of a task that might cause work to be lost if the user does not complete the task; if the changes are published it is willing to terminate the task, accept the context data changes and change its internal state accordingly.
ACL	Access control lists, which determine the privileges and capabilities a particular user has, are presumed to be maintained by each application.
Apply	A user choice; the context data changes are applied to all of the applications, including those that indicated that they might loose work performed by the user; <i>this choice is allowed only when there are no busy applications.</i>
Authentication repository	Enables applications to securely maintain application-specific user authentication data. The repository is used by applications that do not have a built-in means to easily sign-on a user given only a logon name.
AuthenticationRepository (AR)	Interface used by applications to securely interact with the repository to store and retrieve user authentication data.
Automatic Log-off	Logs the current user off of the User Linked applications on a desktop when the user has not interacted with the applications for an appreciable period of time.

Break Link	A user choice; the context changes are applied just to the application with which the user initiated the context changes.
Busy	When an application is unable to apply the context change because it is blocked (e.g., it is a single threaded application that has a modal dialog open); these applications are referred to as <i>busy</i> .
Cancel	A user choice; when the context change is canceled; the context changes are not published.
CCOW	Clinical Context Object Workgroup.
Centralized model	In the <i>centralized model</i> of context management, the responsibility for managing the common context is centralized in a common facility that is responsible for coordinating the sharing of the context among the applications.
Chain of trust	With the <i>chain of trust</i> , User Link-enabled applications and User Link components work together to ensure that only authorized users are allowed access to a common context system.
Clinical context	State information that users establish and modify as they interact with healthcare applications. The context is common because it establishes parameters that should uniformly affect the behavior or operation of multiple healthcare applications.
Common context system	Applications that share the same common context, and have established and maintain a common context link.
Component architecture diagram	Depicts components and their interfaces, and indicates which interfaces each component use for communicating with other components.

Component interaction diagram	Illustrates the series of method invocations that components perform on each other in order to implement a particular use case.
Conditionally accept	When an application might lose work performed by the user if it applies the context changes (e.g., the user was in the process of entering data that would not be applicable in the new context); these applications are referred to as having <i>conditionally accepted</i> the context changes.
CMA	Context Management Architecture.
COM	Microsoft's Component Object Model.
Component model	The architecture of a system as described in terms of components and the interfaces they must implement in order to be participants in the system.
Context change coupon	Unique identifier that is assigned by the context manager to denote each context change transaction.
Context changes pending	During the context change survey, the context manager informs each of the applications in the common context system (except for the application that instigated the changes) that there are pending context data changes. Each application decides whether or not it wants to accept the changes. All applications must accept in order for the context to change.
Context change survey	In the first step of completing a context change transaction the context manager surveys the applications. Each application is informed that there are a candidate set of context data changes and is asked to indicate whether it can accept these changes.
Context change transaction	A multi-step process in that coordinates changes to the common context data. First, an application begins a transaction. The application sets a transaction-specific version of the common context data. Second,

the context manager conducts a context change survey. Third, the context manager reports the survey results to the application that began the transaction. Finally, the application indicates whether the changes are to be applied or cancelled. The decision as to how to proceed may involve the user. If changes are applied, then the transaction-specific version of the context data becomes the new context. Otherwise the transaction-specific context data is discarded.

Context manager	Coordinates applications each time there is a context change transaction. It is also the “owner” of the authentic context data for the system.
Context participants	Applications that set and/or get context data from the context manager. They must follow the policies established later in this document in order to behave as proper context management “citizens.”
Context subject	A subject represents a real-world entity or concept that is identified as part of the overall common clinical context.
ContextData (CD)	Interface implemented by the context manager; used by applications to set/get the data items that comprise the common context.
ContextManager (CM)	Interface implemented by the context manager; used by applications to join/leave a common context system and to indicate the start/end of a set of changes to the common context data.
ContextParticipant (CP)	Interface implemented by an application that wants to participate in a common context system; used by the context manager to inform an application that the context has changed.
CORBA	Common Object Request Broker Architecture.

Corroborating data	Corroborating data can be used by applications and/or users as a basis for checking further that the identified context subject is what was expected.
DCOM	Distributed version of Microsoft's Component Object Model.
Digital signature	Formed using public key / private key encryption techniques, a digital signatures enables
Dispose	A component performs an implicit or explicit action, which is technology-specific, when it no longer intends to use a particular reference. The latter action is referred to as <i>disposing</i> an interface reference.
Distributed model	In the <i>distributed model</i> of context management, the responsibility for managing the common context is uniformly distributed among the applications. There is no central point of common context management.
Empty context	A context is not defined for any subject, either because no context identifier items are present in the context data (as is the case when a context manager is first initialized) or because the values of all of the identifier items for the subject that are present in the context data are <i>null</i> (as is the case when an application explicitly indicates that the context is empty).
Empty context subject	A context subject is <i>empty</i> when a real-world entity or concept is not currently identified. For example, for the patient subject, this means that a patient is not currently identified.
Identifier data	Data that identifies a real-world entity or concept (such as a specific patient or a specific encounter). Identity information is required in order to establish a common context between applications that involves a real-world entity or concept. The string "id" indicates identifier data.

IDL	Interface Definition Language. IDL can specify: an interface's symbolic name, the set of component properties and methods that can be accessed via the interface, the name and data type of each property, the names and data types for each method's input and outputs, and the names and data content for each method's exceptions.
Instigator	The application that began the current context change transaction.
ImplementationInformation (II)	Interface implemented by the context manager and mapping agent; used by applications, context management components, and tools, to obtain details about a component's implementation, including its revision, when it was installed, etc.
Interface interrogation	The interfaces that a component implements can be determined by other components at run-time through direct interrogation.
Interface reference registry	A service that contains references to component interfaces. Components can use the registry to obtain interface references to each other.
Log-off	The termination of a user's session with an application; it assumed that logging-off does not require user authentication.
Mapping agent	A service component that from the perspective of an application is a transparent participant in a context change. A mapping agent's primary role is to add additional subject-specific context identifier items to the context data.
MappingAgent (MA)	Interface implemented by a mapping agent and used by a context manager to inform the mapping agent that the clinical context has changes pending and that the mapping agent should perform its context data mapping responsibilities.

Message authentication code	A secure hash value produced from a data stream that consists of data that is openly communicated between two parties, and “secret” data that they both know but do not openly communicate.
Message digest	A digital signature is formed by applying a secure hash function (alternatively known as a one-way hash function) to the data that is to be transmitted. The resulting hash value is referred to as the <i>message digest</i> , as it is a numeric surrogate for the plain-text message.
Null item value	The value of a context identifier item or corroborating data item can be set to the distinguished value of <i>null</i> to indicate that the item does not have a valid value.
OMA	Object Management Group’s Object Management Architecture.
Participant coupon	Unique identifier that is assigned by the context manager to denote each context participant within a system, including applications and mapping agents.
Passcode	Similar to passwords used by people. However, because passcodes are only used by computer programs, they can be much longer and complex than passwords typically are. This makes passcodes extremely hard to guess, even when brute force techniques are employed.
Patient Link	Enables the user to select the patient of interest once from any application as the means to automatically “tune” all of the applications to the selected patient.
Patient subject	The context subject of <i>Patient</i> is defined for Patient Link. The context data identifier item for this subject is the patient’s medical record number. The patient’s given name is not used as an identifier.

Principal interface	Every component implements at least one well-known interface, referred to as the component's <i>principal interface</i> . The principal interface enables components to perform initial interface interrogations because the name of the principal interface is known a priori, and because all components implement it.
Private key / Public key	An approach for encrypting data, and for creating digital signatures, wherein a matched set of security keys is used. The private key remains the secret of its owner. The matching public key can be disseminated. X can send a message that only Y can read by encrypting the message using Y's public key. Y decrypts the message using its private key. Alternatively, Y can digitally sign its messages using its private key. X can validate Y's signature using Y's public key.
Pull-model	A shared component is used to maintain the shared context data. Applications update this resource to change the data. Other applications periodically poll the component to determine if the data has changed.
Push-model	A shared component is used to maintain the shared context data. This component notifies applications whenever the data is changed. In order to receive a notification, an application must have first explicitly indicated its interest in being notified.
Reauthentication time-out	Requires the currently signed-on user to reauthenticate herself before being allowed to continue using the applications on a clinical desktop. The time-out occurs when the user has not interacted with the desktop for an appreciable period of time.
Repository	See <i>authentication repository</i> .
RMI	Java Remote Method Invocation mechanism.
RSA	A popular public key / private key algorithm.

Secure (or one-way) hash function	A function used for producing a unique numeric surrogate from an arbitrary data stream. It is improbable that two different data streams will yield the same hash value. A secure hash function is an essential part of the infrastructure needed to support the use of digital signatures.
SecureBinding (SB)	Interface used by applications to establish a secure communications binding with the context manager before using the SecureContextData interface. Also used by applications to establish a secure communications binding with the authentication repository before using the AuthenticationRepository interface.
SecureContextData (SD)	Interface similar to the ContextData interface defined for Patient Link; this interface is used by applications to securely set/get the values for the items (logically represented as name-value pairs) that comprise the clinical context.
S-HTTP	Secure Hyper-Text Transfer Protocol.
Sign on	The act of identifying oneself to an application, prior to initiating a user session, in a manner that can be authenticated by the application, typically involving a secret password or a biometric reading (such as a thumb-print scan).
SSL	Secure Socket Layer. SSL enables secure (i.e., encrypted) transmission of data between a client and a server. It also enables a client to authenticate a server (and a server to authenticate a client).
Stat admission	Occurs when an application needs to enable the user to record information about a patient even if an identifier for the patient is not known.
Technology neutral	Means that the common clinical context approach should work equally well with any one of a candidate set of relevant technologies.

Use case diagram	Depicts the actors (human and/or computer-based) and the roles that they play when participating in an interesting scenario.
Use case interaction diagram	Illustrates the high-level interactions between the actors that participate in the use case.
User Link	Enables the user to securely logon once to any application as the means to automatically “tune” all of the applications to the user.
User subject	The context subject of <i>User</i> is defined for User Link. The context data identifier item for this subject is the user’s logon name. The user’s given name is not used as an identifier.
User Link-enabled application	An application that implements the CMA User Link capability.

Health Level Seven Standard

Context Management Specification Component technology Mapping: ActiveX Version CM-1.0

DOCUMENT ID: HL7SIGVI_3_2_99
REVISION ID: March 17, 1999
FILE NAME: hl7_sigvi_activex_cm_1_0 .doc
SUPERCEDES: n/a

Copyright 1999 Health Level Seven

1 Contents

2	1 INTRODUCTION	7
3	1.1 ASSUMPTIONS	7
4	1.2 COMPATABILITY	7
5	1.3 TECHNOLOGY MAPPING	8
6	2 COMPONENT MODEL MAPPING	11
7	3 INTERFACE REFERENCE MANAGEMENT	15
8	4 DUAL INTERFACES.....	17
9	5 WINDOWS REGISTRY SETTINGS	19
10	6 ACTIVEX JAVA WRAPPERS	23
11	7 MICROSOFT'S CRYPTO32 APL.....	25
12	7.1 SECURE BINDING PROPERTIES.....	25
13	7.2 CRYPTOGRAPHIC SERVICE PROVIDER.....	26
14	7.3 CREATING DIGITAL SIGNATURES	26
15	7.4 SIGNATURE FORMAT	26
16	7.5 PUBLIC KEY FORMAT	26
17	7.6 HASH VALUE FORMAT.....	27
18	7.7 KEY CONTAINERS	27
19	7.7.1 <i>Required Containers</i>	27
20	7.7.2 <i>Key Container Naming Convention</i>	28
21	7.7.3 <i>Key Container Management</i>	28
22	7.7.4 <i>Key Container Security</i>	28
23	8 ERROR HANDLING	31
24	9 CHARACTER SET	35
25	10 MIDL LISTING.....	37
26	10.1 TYPE LIBRARIES.....	38
27	10.2 IAUTHENTICATIONREPOSITORY.....	39
28	10.3 ICONTEXTDATA	40
29	10.4 ICONTEXTMANAGER	41
30	10.5 ICONTEXTPARTICIPANT	42
31	10.6 IIMPLEMENTATIONINFORMATION	43
32	10.7 IMAPPINGAGENT	44
33	10.8 ISECUREBINDING	45
34	10.9 ISECURECONTEXTDATA.....	46
35		

1 **Figures**

2 Figure 1: Organization of HL7 Context Management Specification Documents.....9

3 Figure 2: Automation Interfaces in a Common Context System.....12

4

5 **Tables**

6 Table 1: How Interface References Are Obtained.....13

7 Table 2: Secure Binding Properties25

8 Table 3: Key Container Naming Scheme.....29

9 Table 4: Exception Codes.....33

10

11

1

Preface

2

3

4

5

6

This document was prepared by Robert Seliger, Sentillion, Inc., on behalf of Health Level Seven's Special Interest Group for Visual Integration (formerly the Clinical Context Object Workgroup --- CCOW). Comments about the organization or wording of the document should be directed to the author (robs@sentillion.com). Comments about technical content should be directed to ccow@list.mc.duke.edu.

1 Introduction

This document specifies the details needed to develop Microsoft ActiveX implementations of applications and components that conform to the HL7 Context Management Architecture (CMA). Using this specification, the resulting applications and service components will be able to communicate with each other per the CMA even if they were independently developed.

The scope of this document is limited to the details pertaining to implementing the CMA-specified application and component interfaces using ActiveX Automation (formerly known as OLE Automation). This sub-technology within the ActiveX portfolio of technologies is supported by a wide range of Microsoft and non-Microsoft development tools.

Visual Basic® 4.0 is used as the “lowest common denominator” baseline programming language for developing context participant applications. The collective capabilities of Visual Basic® 5.0 (as opposed to 4.0) , Visual C++® 5.0, and Visual J++® 1.1 (Microsoft’s implementation of Java) are used as the baseline programming language implementations for developing CMA components, including the context manager, patient and user mapping agents, and authentication repository. This specification is also forwards-compatible with more recent versions of these tools.

However, any development tool that supports the creation of Automation clients and servers, and in particular supports the IQueryInterface idiom, should enable the development of applications and components that conform to this specification.

1.1 Assumptions

It is assumed that the reader is familiar with Microsoft’s ActiveX technology and with the Microsoft’s underlying Component Object Model (COM).

1.2 Compatability

This specification is compatible with the following host operating systems:

- Windows NT Workstation 4.0 service pack 3, or later
- Windows 95 or later

This specification is compatible with at least the following programming language implementations:

- Visual C++ 5.0 or later

- 1 • Visual Basic 4.0 or later
- 2 • Visual J++ 1.1 or later with Microsoft's Java SDK 3.1 or later and Microsoft's Java
- 3 Virtual Machine 5.00.3161 or later

4 The specification is likely to be compatible with other implementations of these languages, as
5 well as with other programming languages.

6 **1.3 Technology Mapping**

7 The HL7 Context Management Architecture specification is technology-neutral. This means
8 that while an underlying component system is assumed, a specific system is not identified
9 within the architecture. It is the purpose of this document, and its companions for other
10 component technologies, to map the CMA to a specific target technology. For Automation, the
11 technology-specific details specified in this document include (but are not limited to):

- 12 • multiple interfaces
- 13 • interface reference management
- 14 • dual interface requirements
- 15 • registry settings
- 16 • ActiveX Java wrappers for ActiveX components
- 17 • error handling
- 18 • implementable interface definitions

19 It is beyond the scope of this document to provide all of the details that are needed in order to
20 fully implement conformant CMA applications and components. The necessary additional
21 details are covered in a series of companion specification documents, starting most notably
22 with the Health Level Seven Context Management Specification, Technology- And Subject-
23 Independent Component Architecture, CM-1.0.

24 As illustrated in Figure 1, these documents are organized to facilitate the process of defining
25 additional link subjects and to accelerate the process of realizing the CMA using any one of a
26 variety of technologies.

27

28

29

1

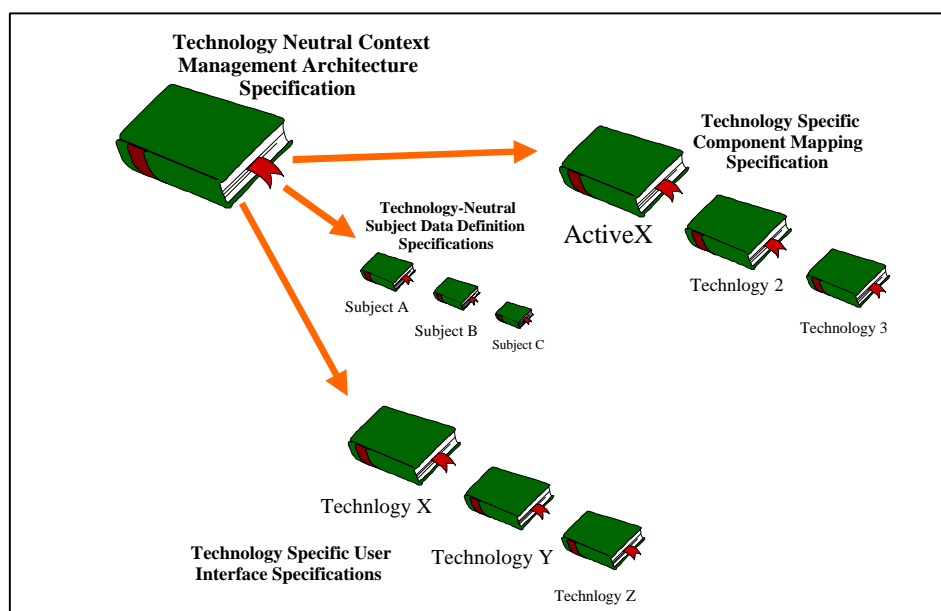


Figure 1: Organization of HL7 Context Management Specification Documents

The context management subjects and technologies that are of interest are determined by the HL7 constituency:

- There is an HL7 context management data definition specification document for each of the standard link subjects. Each document defines the data elements that comprise a link subject. Concurrent with the publication of this document, the following documents have been developed:

Health Level-Seven Standard Context Management Specification,
Data Definition: Patient Subject, Version CM-1.0

Health Level-Seven Standard Context Management Specification,
Data Definition: User Subject, Version CM-1.0

- There is an HL7 context management user interface specification document for each of the user interface technologies with which CMA-enabled applications can be implemented. Each document reflects the user interface requirements established in this document in terms of a technology-specific look-and-feel. Concurrent with the publication of this document, the following document has been developed:

Health Level-Seven Standard Context Management Specification,
User Interface: Microsoft Windows OS, Version CM-1.0

1 Finally, there is an HL7 context management component technology mapping specification
2 document for each of the component technologies. Each document provides the technology-
3 specific details needed to implement CMA-compliant applications and the associated CMA
4 components, as specified in this document. This document serves the role of specifying the
5 details for a CMA implementation using Microsoft's ActiveX technology.

6

2 Component Model Mapping

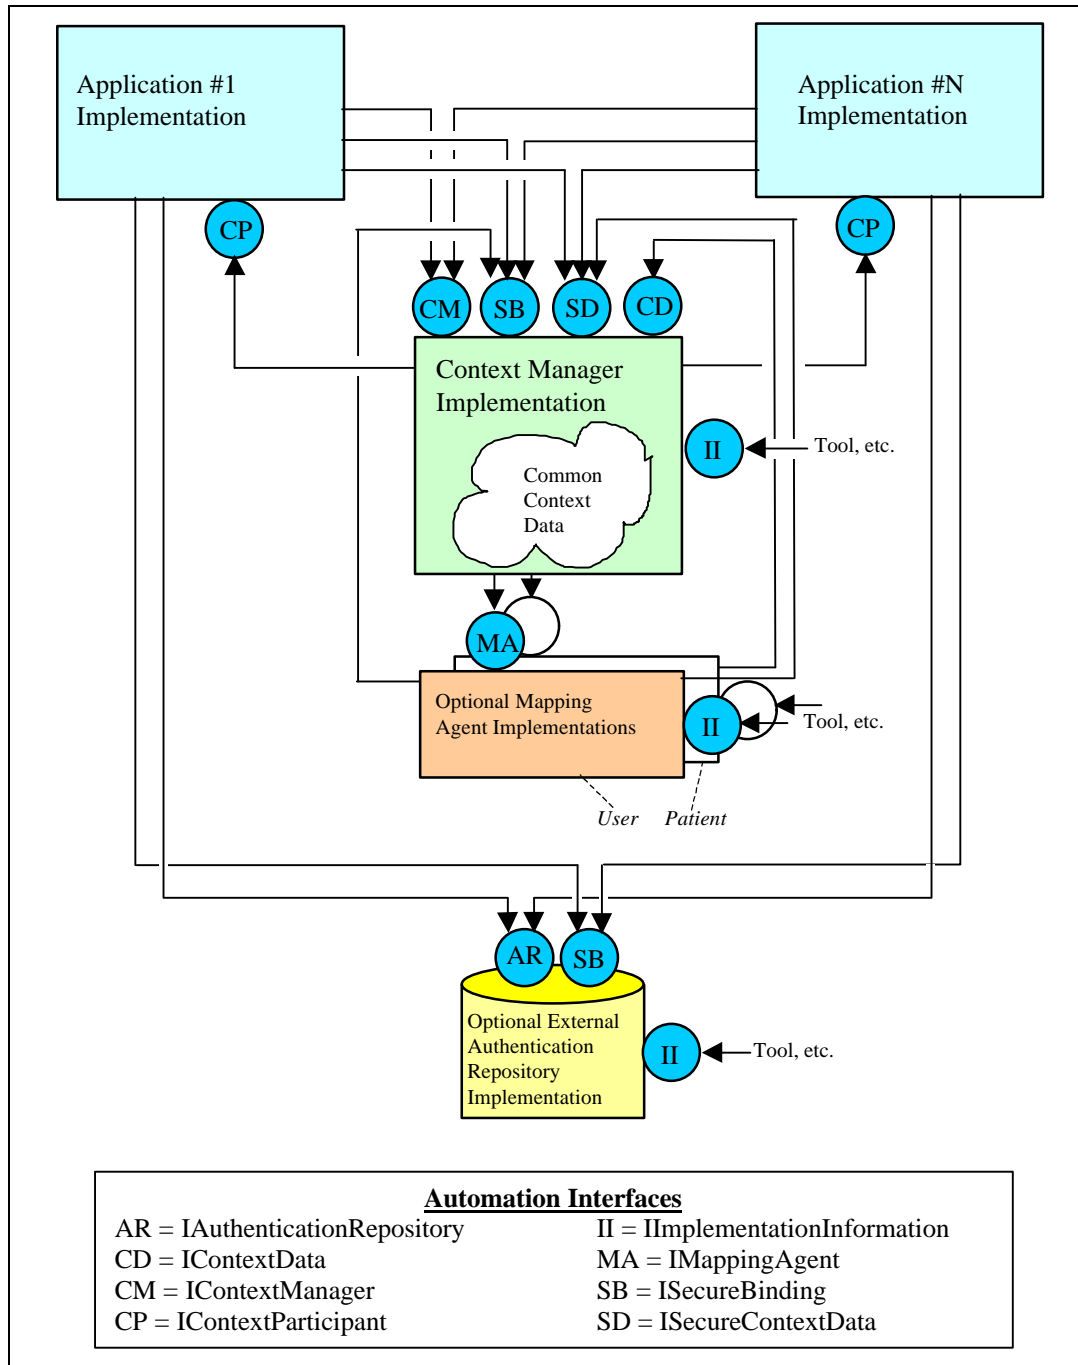
Each interface defined in the CMA specification is implemented as an ActiveX automation interface. All of the components defined in the CMA specification, including context participant applications, are clients as well as servers. In the parlance of ActiveX, they are all Automation clients and servers because they implement and use Automation interfaces.

Context participant applications are only currently required to implement a single Automation interface. However, context managers and mapping agents are required to implement multiple distinct Automation interfaces.

The COM IUnknown::QueryInterface idiom is used to enable context components to acquire each others' interface references through interface interrogation. (Note that Visual Basic implements IUnknown::QueryInterface "under the covers" via the Visual Basic assignment operator.) The COM interface IUnknown serves as a context component's principal interface. See the chapter *Component Model* in the document HL7 Context Management Specification, Technology- And Subject- Independent Component Architecture, CM-1.0 for a discussion about interface interrogation and principal interfaces.

In some cases a component obtains a reference to IQueryInterface for another component from the Windows registry. This registry serves as the interface reference registry described in the chapter *Component Model* in the document HL7 Context Management Specification, Technology- And Subject- Independent Component Architecture, CM-1.0. In other cases, components pass interface references to each other as method parameters.

The various Automation interfaces employed in a common context system are shown in Figure 1. The means by which the various CMA compliant applications and components obtain interface references to each other are described in Table 1.



1

2 **Figure 2: Automation Interfaces in a Common Context System**

3

4

Automation Server	Client's means for obtaining server's interface reference(s) ...	
	Automation Client	Means for obtaining reference
Context Manager's IContextManager and IContextData interfaces.	Context Participant	A context participant obtains a reference to the context manager's IUnknown interface from the Windows registry. The context participant then performs IUnknown::QueryInterface on the context manager to get the desired interface references.
Context Manager's IContextData interface.	Mapping Agent	The context manager provides a reference to its IUnknown interface to the mapping agent when the context manager calls IMappingAgent::ContextChangesPending. The mapping agent then performs IUnknown::QueryInterface on the context manager to get the desired interface reference.
Mapping Agent's IMappingAgent and IImplementationInformation interfaces.	Context Manager	The context manager obtains a reference to the mapping agent's IUnknown interface from the Windows registry. The context manager then performs IUnknown::QueryInterface on the mapping agent to get the desired interface references.
Context Participant's IContextParticipant interface.	Context Manager	A context participant provides a reference to its IContextParticipant interface to the context manager when the context participant calls IContextManager::JoinCommonContext.
Authentication Repository's IAuthenticationRepository	Context Participant	A context participant obtains a reference to the authentication repository's IUnknown interface from the Windows registry. The context participant then performs IUnknown::QueryInterface on the authentication repository to get the desired interface references.

1 **Table 1: How Interface References Are Obtained**

2

1 **3 Interface Reference Management**

2 In order to “possess” an interface reference, as described in the chapter *Component Model* in
3 the HL7 Context Management Specification, Technology- And Subject- Independent
4 Component Architecture, CM-1.0 document, COM interface reference counts should be
5 incremented and decremented in accordance with COM conventions. In general, a component
6 performs IUnknown::AddRef to “possess” an interface reference. Conversely, a component
7 performs IUnknown::Release to “dispose” an interface reference.

1 **4 Dual Interfaces**

2 Dual Interfaces are a COM optimization that enables an Automation interface to be called
3 using a run-time dispatching mechanism (i.e., so called *dispatch interfaces*), or directly via a
4 compile-time binding mechanism (i.e., so called v-table interfaces). The latter approach
5 generally results in better performance. Dual interfaces accommodate the widest possible range
6 of application development tools, from interpreted late binding languages like Smalltalk and
7 VisualBasic to compiled early binding languages like C and C++.

8 Context manager, mapping agent, and authentication repository implementations shall expose
9 their CMA-defined Automation interfaces as dual interfaces. This may limit the choice of
10 programming language for these components to just those that support the development of dual
11 interfaces. However, the advantage is better overall run-time performance.

12 Context participant applications can choose to implement their CMA-defined
13 IContextParticipant interface as a dispatch interface or as a dual interface. This enables
14 application developers to use a wide range of programming languages, as not all languages
15 support dual interfaces.

5 Windows Registry Settings

ActiveX components can have a wide variety of Windows registry entries. It is not unusual for these entries to become quite complex. An objective of this document is to specify the simplest registry entries that will enable applications and components that conform to the CMA specifications to be implemented using any of the common ActiveX-capable programming languages and still seamlessly interoperate.

The context manager shall be registered in the Windows registry. This enables context participant applications to locate and bind to the context manager. If present, a mapping agent shall also be registered in the Windows registry. This enables the context manager to locate and bind to the mapping agent. Finally, if present, the authentication repository shall be registered in the Windows registry. This enables context participant applications to locate and bind to the authentication repository.

ActiveX component registry entries often include implementation-specific information, such as the file name and path to the component's executable code, and may vary depending upon how the component has been implemented (e.g., executable vs. dynamic link library). However, the registry entry for an ActiveX component can use a program identifier (ProgID), which is a symbolic name for the type of component, as a registry key. A registry key is used to locate a registry entry (known as a *value*).

The value associated with a ProgID is the component's class identifier (CLSID), which denotes an implementation of the component. By fixing the ProgID, it is possible to write client's for a type of component such that the client does not need to know anything about the component's implementation. Instead, the client uses the ProgID to locate the component's CLSID at run-time. The CLSID is then used to create an instance of the component, or to connect to an existing instance of a running component.

In summary, ProgID's are invariant across implementation. Therefore, no matter how they are implemented, all of the CMA compliant applications and components shall use the ProgID's defined below¹:

- The context manager shall be registered using the ProgID sub-key string, "CCOW.ContextManager". The CLSID under which a context manager is registered shall be different for different context manager implementations.

¹ These ProgID's are the same as defined by the Clinical Context Object Workgroup, upon whose original specification this specification is based.

- 1 • The patient mapping agent shall be registered using the ProgID sub-key string,
2 “CCOW.MappingAgent_Patient”. The CLSID under which the patient mapping agent
3 is registered shall be different for different patient mapping agent implementations.
- 4 • The user mapping agent shall be registered using the ProgID sub-key string,
5 “CCOW.MappingAgent_Patient”. The CLSID under which the user mapping agent is
6 registered shall be different for different user mapping agent implementations.
- 7 • The authentication repository shall be registered using the ProgID sub-key string,
8 “CCOW.AuthenticationRepository”. The CLSID under which the authentication
9 repository is registered shall be different for different authentication repository
10 implementations.

11 The ProgID prefix “CCOW” is reserved for use by HL7 for creating future CMA-related
12 ProgIDs. A CMA-compliant application or component shall not use this prefix other than as
13 specified in this document.

14 The use of a common ProgID but implementation-specific CLSID requires additional effort on
15 the part of context manager and mapping agent developers. It may also require additional
16 effort on the part of context participant developers:

- 17 • Context manager, mapping agent, and authentication repository implementations shall
18 provide ActiveX Java wrapper classes for their CMA coclasses and interfaces as part
19 of their installation package. The details of how these wrapper classes should be
20 prepared and packaged are described below. These wrapper classes are needed in
21 order to hide the ActiveX implementation details of these components, including their
22 CLSIDs, from J++ Automation clients for these components.
- 23 • Context manager, mapping agent, and authentication repository implementations shall
24 each provide ActiveX-compliant registry entries in
25 HKEY_CLASSES_ROOT\Interface\ for each of their CMA-specified
26 Automation interfaces. This information is needed so that the Automation clients for
27 these components can create instances of these interfaces.
- 28 • Context manager, mapping agent, and authentication repository implementations shall
29 each provide an ActiveX-compliant registry entry
30 HKEY_CLASSES_ROOT\TypeLib\ for their respective type libraries. This
31 information is needed so that the Automation clients for these components can create
32 calls to these interfaces using the dispatch mechanism.
- 33 • Developers of CMA-compliant context participant applications and components shall
34 use the ProgId, not the CLSID, to bind to any of the CMA-defined components that
35 are registered in the registry. This enables implementations to be changed without
36 affecting interoperability.

- 1 • Developers of J++ CMA-compliant context participant applications and components
2 shall use the ActiveX Java wrapper classes provided with the CMA-defined
3 components of which they are clients. This is as opposed to client-generated wrappers,
4 which require that the client have development time (versus run-time) access to the
5 implementation of the wrapped component's type library. This is not only impractical,
6 but introduces the probability that a J++ client would only work with a specific
7 Automation server implementation.

8 When these rules are followed, context participant applications and CMA components will
9 interoperate independently of each other's implementations.

6 ActiveX Java Wrappers

Context manager and mapping agent implementations must provide ActiveX Java wrapper classes:

- The Java package name "ccow.contextmanager" shall be used for the context manager package.
- The Java package name "ccow.mappingagent_patient" shall be used for the patient mapping agent package.
- The Java package name "ccow.mappingagent_user" shall be used for the user mapping agent package
- The Java package name "ccow.authenticationrepository" shall be used for the authentication repository package.
- The context manager package shall minimally contain the Java wrapper classes ContextManager.class, IContextManager.class, IContextData.class, ISecureContextData.class, ISecureBinding.class, ImplementationInformation.class and IContextParticipant.class.
- Both of the mapping agent packages shall minimally contain the Java wrapper classes MappingAgent.class, IMappingAgent.class, and ImplementationInformation.class.
- The authentication repository package shall minimally contain the Java wrapper classes AuthenticationRepository.class, IAuthenticationRepository.class, ISecureBinding.class, and ImplementationInformation.class.

The wrapper classes hide component implementation details. One specific detail hidden is the CLSID to be used by J++ Automation clients for these objects. In order to hide these details, the wrapper classes must be created with knowledge of the details that they hide, hence the need for them to be provided with each component implementation.

From the perspective of a J++ Automation client, the wrapper classes will look and behave the same across component implementations. The wrapper classes are dynamically loaded by a J++ client whenever it first accesses the corresponding Automation client.

The installation of a new component will simply cause J++ clients to automatically access a different version of a seemingly identical component.

1 The wrapper classes for the context manager should be packaged as "package
2 ccow.contextmanager" and located in:

3 %windir%\java\trustlib\ccow\contextmanager
4

5 The wrapper classes for the patient mapping agent should be packaged as "package
6 ccow.mappingagent_patient" and located in:

7 %windir%\java\trustlib\ccow\mappingagent_patient
8

9 The wrapper classes for the user mapping agent should be packaged as "package
10 ccow.mappingagent_user" and located in:

11 %windir%\java\trustlib\ccow\mappingagent_user
12

13 The wrapper classes for authentication repository should be packaged as "package
14 ccow.authenticationrepository" and located in:

15 %windir%\java\trustlib\ccow\authenticationrepository
16

17 Note that ccow, contextmanager, mappingagent_patient,
18 mappingagent_user, and authenticationrepository are all lower case.

19

7 Microsoft's CRYPTO32 API

All ActiveX implementations of CMA-compliant applications and components that use the CMA-defined secure interfaces shall use the RSA public key / private key scheme and shall use the MD5 one-way hash algorithm. It is recommended that Microsoft's Cryptography Application Programming Interface (CRYPTO32) be used, and that the Microsoft RSA Base Provider be selected as the cryptographic service provider.

However, a different API and/or cryptographic service provider implementation can be used as long as it employs algorithms and binary data formats that are functionally identical to those employed by the Microsoft RSA Base Provider as accessed via the CRYPTO32 API.

7.1 Secure Binding Properties

The CMA-defined interface ISecureBinding requires that the bindee indicate to the binder various security properties that depend upon the bindee's implementation. The properties that must be indicated, and the allowed value or values for each property, depend upon the underlying implementation technology.

For an ActiveX implementation, the following secure binding property names and values defined in Table 2: Secure Binding Properties shall be used.

Property Name	Allowed Value	Meaning
Technology	CRYPTO32	Microsoft CRYPTO32 or equivalent.
PubKeyScheme	RSA_EXPORTABLE ²	Exportable version of RSA public key / private key scheme (employs 40 bit keys).
HashAlgo	MD5	MD5 secure hash algorithm (creates 128 bit hash).

Table 2: Secure Binding Properties

² Public key / private key schemes are subject to United States export restrictions. Specifically, The U.S. Government limits the size (in bits) of the security keys that can be used as part of applications exported by U.S. vendors. The Microsoft Base Service Provider has been approved for export by the U.S. Government. Applications that use this CSP via the CRYPTO32 API should not require additional export approvals.

1 Property names are not case sensitive. Property values shall be character-encoded per the
2 convention stated in the CMA specification.

3 **7.2 Cryptographic Service Provider**

4 The CRYPTO32 API enables applications to select from a set of cryptographic service
5 providers (CSP). Each CSP provides cryptographic services that can be accessed via the
6 CRYPTO32 API. For CMA-compliant applications and components that are implemented
7 using the CRYPTO32 API, the Microsoft RSA Base Provider shall be used as the
8 cryptographic service provider. This means that the value of the *dwProvType* to the
9 CRYPTO32 function *CryptAcquireContext* shall be *PROV_RSA_FULL*.

10 **7.3 Creating Digital Signatures**

11 The CRYPTO32 function *CryptSignHash* is used to create a digital signature. The function
12 *CryptVerifySignature* is used to verify a signature. Both of these functions accept an optional
13 pointer to a character string for the parameter *sDescription*. The value of this parameter shall
14 be NULL for all calls to these functions as it pertains to creating or comparing signatures used
15 to implement User Link.

16 **7.4 Signature Format**

17 Digital signatures passed via any of the CMA-defined Automation interfaces shall be
18 represented as a string. This string contains binary data that has been character-encoded per
19 the convention defined in CMA specification. The binary data from which a signature string is
20 created is the byte array produced by *CryptSignHash*. This string must be converted back to
21 binary data in order to be used as an input to *CryptVerifySignature*.

22 **7.5 Public Key Format**

23 Public keys passed via any of the CMA-defined Automation interfaces shall be represented as
24 a string. This string contains binary data that has been character-encoded per the convention
25 defined in CMA specification. The binary data from which a public key is created is the byte
26 array produced by *CryptExportKey* with the parameter *dwBlobType* set to
27 *PUBLICKEYBLOB*. This string must be converted back to binary data in order to be used as
28 an input to *CryptImportKey*.

7.6 Hash Value Format

Hash values passed via any of the CMA-defined Automation interfaces shall be represented as a string. This string contains binary data that has been character-encoded per the convention defined in CMA specification. The binary data from which a hash value is created is the byte array produced by CryptGetHashParam. Hash values shall be compared for equality by comparing their character-encode string representations. Character case shall not be considered when comparing these strings.

7.7 Key Containers

With CRYPTO32, public keys and public key / private key pairs are maintained in key containers. These containers can be created and deleted using the CRYPTO32 API function CryptAcquireContext. Keys can be imported into a container, or keys can be directly generated within an empty container.

7.7.1 Required Containers

An application shall maintain the following key containers:

- A key container for holding its own public key / private key pair.
- A key container for holding the context manager's public key.
- Optionally, a key container for holding the authentication repository's public key.

The context manager shall maintain the following key containers:

- A key container for holding its own public key / private key.
- A key container for holding each designated application's public key.
- A key container for holding the user mapping agent's public key.

The user mapping agent shall maintain the following key containers:

- A key container for holding its own public key / private key.
- A key container for holding the context manager's public key.

The authentication repository shall maintain the following key containers:

- A key container for holding its own public key / private key.

- A key container for holding the public keys for each of applications that use the repository.

The convention for naming these containers and for managing their creation and deletion are described next.

7.7.2 Key Container Naming Convention

All of the key containers shall have unique names when they are co-resident on the same Windows host. The naming convention is defined in Table 3: Key Container Naming Scheme.

Note that all of the letters in a container's name shall be capitalized. Also note that the portion of a container name shown as *APPLICATION-NAME* is the same string that an application provides to the context manager when it joins the common context system.

7.7.3 Key Container Management

An application, context manager, user mapping agent, and authentication repository shall delete any containers that its has created prior to terminating.

However, an application, context manager, user mapping agent, or authentication repository that terminates prematurely might fail to delete some or all of the containers that it has created. When the failed component is next launched it will not be able to create a new container if a previously created container with the same name still exists. This situation shall be handled as follows: The existing container shall be deleted and a new container created in its stead. The necessary keys shall be created and/or imported into the new container.

7.7.4 Key Container Security

CMA-compliant applications and components that maintain key containers shall protect their containers from unauthorized access. This means that only the application or component that created the container should be able to access the container.

If key containers are not protected then they are vulnerable to unintended uses. For example, a rogue application might access the keys within a container created by valid CMA-compliant application as a means to impersonate the application within a context management system.

There are a variety of ways to protect key containers. In order to maximize design flexibility for CMA-compliant applications and components, a particular approach is not defined in this specification.

1

Container created by	Container purpose ...	Container name ...
Application	Holding own key pair. Holding context manager's public key. Holding authentication repository's public key.	<i>CCOW.APPLICATION-NAME.SELF</i> <i>CCOW.APPLICATION-NAME.CM</i> <i>CCOW.APPLICATION-NAME.AR</i>
Context Manager	Holding own pair. Holding an application's public key. Holding user mapping agent's public key.	<i>CCOW.CM.SELF</i> <i>CCOW.CM.APPLICATION-NAME</i> <i>CCOW.CM.MA_USER</i>
User Mapping Agent	Holding own key pair. Holding context manager's public key.	<i>CCOW.MA_USER.SELF</i> <i>CCOW.MA_USER.CM</i>
Authentication Repository	Holding own key pair. Holding an application's public key.	<i>CCOW.AR.SELF</i> <i>CCOW.AR.APPLICATION-NAME</i>

2

Table 3: Key Container Naming Scheme

3

8 Error handling

The CMA specifies a set of exceptions that can be raised by CMA components. (Context participant applications do not currently throw exceptions).

ActiveX Automation exceptions are implemented in a two-stage process. First, all Automation and dual interface methods return a 32-bit encoded error value, called an HRESULT, to their caller. Secondly, ActiveX components that support the Microsoft-defined IErrorInfo and ISupportErrorInfo interfaces can provide additional error information to clients when requested. This information includes a textual description of the error and the guid³ of interface that threw the error.

Each of the CMA-specified exceptions is identified by a distinguished HRESULT. Additionally, the context manager, both mapping agents, and the authentication repository shall support the IErrorInfo and ISupportErrorInfo interfaces. Automation clients for these objects should check the HRESULT after each method invocation to determine if an exception has occurred. Clients may then optionally access additional error information via the server component's IErrorInfo interface.

In the Win32 COM implementation there is at most one error object associated with each logical thread of execution (i.e. a thread can logically span multiple processes on the same or different hosts), and that the error object may be overwritten by a subsequent error. Clients should access IErrorInfo immediately after detecting an exception to insure that the error information they obtain is pertinent.

Visual Basic developers should note that the Visual Basic Err object handles all the IErrorInfo manipulations automatically. In the event that a Visual Basic client encounters an exception, the Visual Basic Err object will contain the exception information.

The list of CMA-defined HRESULTS values is shown in Table 4: Exception Codes.

³ A guid is a globally unique identifier. Every COM interface definition is denoted by a different guid.

Exception	Hexadecimal value	Explanation
NotImplemented	0x80004001L	Method not implemented. This is the same value as defined for the Win32 E_NOT_IMPL HRESULT.
GeneralFailure	0x80004005L	An error was detected or a failure occurred. This is the same value as defined for the Win32 E_FAIL HRESULT.
ChangesNotEnded	0x80000201L	Attempt to publish context changes before ending the context change transaction.
InvalidContextCoupon	0x80000203L	A context coupon does not match most recently committed coupon or current transaction coupon.
NameValueCountMismatch	0x80000206L	A name array and its corresponding value array do not have the same number of elements.
NotInTransaction	0x80000207L	Attempt to perform a context management transaction method when a transaction is not in progress.
TransactionInProgress	0x80000209L	Attempt to perform a context management method when a transaction is in progress.
UnknownItemName	0x8000020AL	An item name not known.
UnknownParticipant	0x8000020BL	Participant coupon does not denote a known participant.
TooManyParticipants	0x8000020CL	Attempt to join a context that can't accommodate another participant.
AcceptNotPossible	0x8000020DL	Attempt to publish an "accept" decision but there were participants for which it was not possible to obtain a survey response (e.g., these participants were blocked)
BadItemNameFormat	0x8000020EL	An item name does not conform to format rules.
BadItemType	0x8000020FL	An item data type does not conform to data definition for the item.
BadItemValue	0x80000210L	An item value does not conform to the allowed set of values as defined by the data definition for the item.

InvalidTransaction	0x80000211L	A transaction has been invalidated and aborted because it violates one or more semantic integrity constraints.
UndoNotPossible	0x80000212L	Attempt to undo context changes after the transaction has ended.
ChangesNotPossible	0x80000213L	Attempt to set or delete context data after the transaction has ended.
ChangesNotAllowed	0x80000214L	Mapping agent attempts set or delete a context data item that has been set by the participant that instigated the transaction.
AuthenticationFailed	0x80000215L	A signature could not be authenticated.
SignatureRequired	0x80000216L	A signature is required to perform the method.
UnknownApplication	0x80000217L	An application name is not known.
UnknownConnection	0x80000218L	A connection is not known to the authentication repository.
LogonNotFound	0x80000219L	The desired user logon is not found in the authentication repository.
UnknownDataFormat	0x8000021AL	The format of user authentication data requested could not be found in the authentication repository.
UnknownBindee	0x8000021BL	A security binding coupon does not denote a known bindee.
ImproperKeyFormat	0x8000021CL	A public key is not properly formatted.
BindingRejected	0x8000021DL	The identity of a bindee could not be verified.
ImproperMACFormat	0x8000021EL	A message authentication code is not properly formatted.
UnknownPropertyName	0x8000021FL	A property name is not known.
BadPropertyType	0x80000220L	A property data type does not conform to specification.
BadPropertyValue	0x80000221L	A property data value does not conform to specification.
AlreadyJoinedContext	0x80000222L	The application has already joined the context.

1

Table 4: Exception Codes

1 **9 Character Set**

2 The Unicode character set shall be used to represent all character strings that are transmitted
3 amongst and between CMA-compliant applications and components. The Unicode character
4 set enables representation of virtually any local character set.

5 The use of ActiveX Automation, in which character strings are represented by the Automation
6 data type BSTR, provides built-in support for Unicode. This means that an ActiveX
7 implementation of a CMA-compliant applications and components will inherently support
8 Unicode for the character strings that are communicated via the CMA-defined ActiveX
9 Automation interfaces.

10 MIDL Listing

The interfaces defined below are an implementable translation of the abstract interfaces definitions documented in the CMA specification. The following rules were applied to produce the translation:

- The prefix “I” is prepended to the names of each interface, per COM conventions.
- The closest available data types supported by Automation were employed (see table below).
- Outputs are mapped as return values (retval) and in/out parameters. Plain out parameters are not used because they are not easily implemented using Visual Basic 5.0. (Note: the use of in/out parameters requires special attention to proper memory management techniques when implementing context managers or context participants with the C++ programming language.)
- Exceptions names are mapped as HRESULTs. Support for exception data values is optional. If supported, the data values should be mapped to formatted strings and made available through the IErrorInfo interface.
- An interface reference to a component’s principal interface is mapped as an IUnknown pointer. A reference to any other component interface is mapped as an IDispatch pointer.
- Sequences are mapped as safe arrays.
- Abstract data types are mapped to Automation data types as follows:

Abstract Data Type	Automation Data Type
byte	unsigned char
short	short
long	long
float	float
double	double
boolean	VARIANT_BOOL
string	BSTR
date	DATE
type	VARTYPE

variant	VARIANT
---------	---------

1

2 The MIDL below must be used by all ActiveX implementations of context managers and
3 context participants. This includes interface and class names, and method signatures.

4 **10.1 Type Libraries**

5 All CMA-compliant Automation server component implementations shall provide a type
6 library that is consistent with the interface definitions specified below. A default interface
7 should not be specified for any of these components. Clients should not assume that an
8 Automation server has a default interface. An explicit call to IUnknown::QueryInterface is
9 necessary to obtain a reference to a specific interface from an Automation server.

10.2 *IAuthenticationRepository*

```

1  10.2 IAuthenticationRepository
2
3
4  import "oaidl.idl";
5  import "ocidl.idl";
6
7  [
8      object,
9      uuid(12B28736-2895-11d2-BD6E-0060B0573ADC),
10     dual,
11     helpstring("IAuthenticationRepository Interface"),
12     pointer_default(unique)
13 ]
14 interface IAuthenticationRepository : IDispatch
15 {
16     [helpstring("Establish connection with authentication repository")]
17     HRESULT Connect([in] BSTR applicationName,
18                    [out, retval] long *bindingCoupon);
19
20     [helpstring("Terminate connection with authentication repository")]
21     HRESULT Disconnect([in] long bindingCoupon);
22
23     [helpstring("Set user authentication data for specified logon name")]
24     HRESULT SetAuthenticationData([in] coupon,
25                                  [in] BSTR logonName,
26                                  [in] BSTR dataFormat,
27                                  [in] BSTR appSignature);
28
29     [helpstring("Delete user authentication data for specified logon name")]
30     HRESULT DeleteAuthenticationData([in] coupon,
31                                      [in] BSTR logonName,
32                                      [in] BSTR dataFormat,
33                                      [in] BSTR appSignature);
34
35     [helpstring("Retrieve user authentication data for specified logon name")]
36     HRESULT GetAuthenticationData([in] coupon,
37                                   [in] BSTR logonName,
38                                   [in] BSTR dataType,
39                                   [in] BSTR appSignature,
40                                   [in, out] BSTR *userData,
41                                   [out, retval] BSTR *repositorySignature);
42 };

```

10.3 IContextData

```

1  10.3 IContextData
2
3
4  import "oaidl.idl";
5  import "ocidl.idl";
6
7  [
8      object,
9      uuid(2AAE4991-A1FC-11D0-808F-00A0240943E4),
10     dual,
11     helpstring("IContextData Interface"),
12     pointer_default(unique)
13 ]
14 interface IContextData : IDispatch
15 {
16     [helpstring("get the names of all of the context items")]
17     HRESULT GetItemNames([in] long contextCoupon, [out, retval] VARIANT *itemNames);
18
19     [helpstring("delete an item(s) from the set of context items")]
20     HRESULT DeleteItems([in] long participantCouppn,
21                         [in] VARIANT names,
22                         [in] long contextCoupon);
23
24     [helpstring("set the value of one or more context items")]
25     HRESULT SetItemValues([in] long participantCoupon,
26                          [in] VARIANT itemNames,
27                          [in] VARIANT itemValues,
28                          [in] long contextCoupon);
29
30     [helpstring("get the value of one or more context items")]
31     HRESULT GetItemValues([in] VARIANT names,
32                          [in] VARIANT_BOOL onlyChanges,
33                          [in] long contextCoupon,
34                          [out, retval] VARIANT *itemValues);
35 };

```

10.4 *IContextManager*

```

1  import "oaidl.idl";
2
3  import "ocidl.idl";
4
5  [
6
7      object,
8      uuid(41126C5E-A069-11D0-808F-00A0240943E4),
9      dual,
10     helpstring("IContextManager Interface"),
11     pointer_default(unique)
12 ]
13
14 interface IContextManager : IDispatch
15 {
16     [propget, helpstring("property MostRecentContextCoupon")]
17     HRESULT MostRecentContextCoupon([out, retval] long *pVal);
18
19     [helpstring("enables an application to join a common context system")]
20     HRESULT JoinCommonContext([in] IDispatch *contextParticipant,
21                               [in] BSTR sApplicationTitle,
22                               [in] VARIANT_BOOL survey,
23                               [in] VARIANT_BOOL wait,
24                               [out, retval] long *participantCoupon);
25
26     [helpstring("enables an application to leave a common context system")]
27     HRESULT LeaveCommonContext([in] long participantCoupon);
28
29     [helpstring("enables an application to start a context change transaction")]
30     HRESULT StartContextChanges([in] long participantCoupon,
31                                 [out, retval] long *pCoupon);
32
33     [helpstring("enables the application that instigated a context change transaction to
34 indicate that it has completed its changes")]
35     HRESULT EndContextChanges([in] long contextCoupon,
36                               [in, out] VARIANT_BOOL *someBusy,
37                               [out, retval] VARIANT *vote);
38
39     [helpstring("enables an application to discard any context data changes that it has
40 already made")]
41     HRESULT UndoContextChanges([in] long contextCoupon);
42
43     [helpstring("enables the application that instigated a context change transaction to
44 inform the other applications in a context system about whether the changes are to be
45 applied or have been canceled")]
46     HRESULT PublishChangesDecision([in] long contextCoupon,
47                                    [in] BSTR decision);
48
49     [helpstring("enables an application to indicate that it wants to suspend its active
50 participation in a common context system while remaining registered as a
51 participant")]
52     HRESULT SuspendParticipation([in] long participantCoupon);
53
54     [helpstring("enables an application to indicate that it wants to resume active
55 participation in a common context system")]
56     HRESULT ResumeParticipation([in] long participantCoupon,
57                                 [in] VARIANT_BOOL wait );
58 };
59

```


10.5 IContextParticipant

```

1  10.5 IContextParticipant
2
3
4  import "oaidl.idl";
5  import "ocidl.idl";
6
7  [
8      object,
9      uuid(3E3DD272-998E-11D0-808D-00A0240943E4),
10     dual,
11     helpstring("IContextParticipant Interface"),
12     pointer_default(unique)
13 ]
14 interface IContextParticipant : IDispatch
15 {
16     [helpstring("informs a participant that a change to the common context data is
17     pending")]
18     HRESULT ContextChangesPending([in] long contextCoupon,
19     [in, out] BSTR* reason,
20     [out, retval] BSTR *returnValue);
21
22     [helpstring("informs a participant that the common context data has changed")]
23     HRESULT ContextChangesAccepted([in] long contextCoupon);
24
25     [helpstring("informs a participant that a context change transaction has been rejected
26     by one or more of the other participating applications")]
27     HRESULT ContextChangesCanceled([in] long contextCoupon);
28
29     [helpstring("informs a participant that the system is being terminated")]
30     HRESULT CommonContextTerminated(void);
31
32     [helpstring("used to test if the participant is alive")]
33     HRESULT Ping(void);
34 };
35

```

10.6 Implementation Information

```

1
2
3
4 import "oaidl.idl";
5 import "ocidl.idl";
6
7 [
8     object,
9     uuid(41123600-6CE1-11d1-AB3F-E892F5000000),
10    dual,
11    helpstring("ImplementationInformation Interface"),
12    pointer_default(unique)
13 ]
14 interface ImplementationInformation : Idispatch
15 {
16     [propget, helpstring("property ComponentName")]
17     HRESULT ComponentName([out, retval] BSTR *pVal);
18
19     [propget, helpstring("property RevMajorNum")]
20     HRESULT RevMajorNum([out, retval] BSTR *pVal);
21
22     [propget, helpstring("property RevMinorNum")]
23     HRESULT RevMinorNum([out, retval] BSTR *pVal);
24
25     [propget, helpstring("property PartNumber")]
26     HRESULT PartNumber([out, retval] BSTR *pVal);
27
28     [propget, helpstring("property Manufacturer")]
29     HRESULT Manufacturer([out, retval] BSTR *pVal);
30
31     [propget, helpstring("property TargetOS")]
32     HRESULT TargetOS([out, retval] BSTR *pVal);
33
34     [propget, helpstring("property TargetOSRev")]
35     HRESULT TargetOSRev([out, retval] BSTR *pVal);
36
37     [propget, helpstring("property WhenInstalled")]
38     HRESULT WhenInstalled([out, retval] BSTR *pVal);
39 };
40

```

10.7 IMappingAgent

```

1  10.7 IMappingAgent
2
3
4  import "oaidl.idl";
5  import "ocidl.idl";
6
7  [
8      object,
9      uuid(753D98C0-6CE1-11d1-AB3F-E892F5000000),
10     dual,
11     helpstring("IMappingAgent Interface"),
12     pointer_default(unique)
13 ]
14 interface IMappingAgent : Idispatch
15 {
16     [helpstring("informs a mapping that a change to the common context data ready for
17 mapping")]
18     HRESULT ContextChangesPending([in] long mappingAgentCoupon,
19                                   [in] IUnknown *contextMgr,
20                                   [in] long contextCoupon,
21                                   [in, out] BSTR* reason,
22                                   [out, retval] BSTR *returnValue);
23
24     [helpstring("used to let Context Manager mapping agent is alive")]
25     HRESULT Ping(void);
26 };
27

```

10.8 ISecureBinding

```

1  10.8 ISecureBinding
2
3
4  import "oaidl.idl";
5  import "ocidl.idl";
6
7  [
8      object,
9      uuid(F933331D-91C6-11D2-AB9F-4471FBC00000),
10     dual,
11     helpstring("ISecureBinding Interface"),
12     pointer_default(unique)
13 ]
14 interface ISecureBinding : IDispatch
15 {
16     [helpstring("Initiate secure binding")]
17     HRESULT InitiatlizeBinding([in] long bindeeCoupon,
18                               [in] VARIANT propertyNames,
19                               [in] VARIANT propertyValues,
20                               [in, out] BSTR *binderPublicKey,
21                               [out, retval] BSTR *mac);
22
23     [helpstring("Finalize secure binding")]
24     HRESULT FinalizeBinding([in] long bindeeCoupon,
25                             [in] BSTR bindeePublicKey,
26                             [in] BSTR mac);
27 };
28

```

10.9 ISecureContextData

```

1  10.9 ISecureContextData
2
3  import "oaidl.idl";
4  import "ocidl.idl";
5
6  [
7      object,
8      uuid(6F530680-BC14-11D1-90B1-76C60D000000),
9      dual,
10     helpstring("ISecureContextData Interface"),
11     pointer_default(unique)
12 ]
13 interface ISecureContextData : IDispatch
14 {
15     [helpstring("return collection of the names in the context")]
16     HRESULT GetItemNames([in] long contextCoupon,
17                          [out, retval] VARIANT *itemNames);
18
19     [helpstring("set the value of one or more context items")]
20     HRESULT SetItemValues([in] long participantCoupon,
21                          [in] VARIANT itemNames,
22                          [in] VARIANT itemValues,
23                          [in] long contextCoupon,
24                          [in] BSTR appSignature);
25
26     [helpstring("obtain the value of one or more context items")]
27     HRESULT GetItemValues([in] long participantCoupon,
28                          [in] VARIANT names,
29                          [in] VARIANT_BOOL onlyChanges,
30                          [in] long contextCoupon,
31                          [in] BSTR appSignature,
32                          [in, out] BSTR *managerSignature,
33                          [out, retval] VARIANT *itemValues);
34 };
35

```

Health Level Seven Standard

Context Management Specification Data Definition: Patient Subject Version CM-1.0

DOCUMENT ID: HL7SIGVI_3_4_99

REVISION ID: March 17, 1999

FILE NAME: hl7_sigvi_patient_cm_1_0 .doc

SUPERCEDES: n/a

Copyright 1999 Health Level Seven

1	Contents	
2	1 INTRODUCTION.....	5
3	1.1 CONTEXT MANAGEMENT DOCUMENT OVERVIEW	5
4	1.2 CONTEXT DATA SUBJECT.....	7
5	1.3 CONTEXT DATA ITEM FORMAT.....	7
6	1.4 CASE SENSITIVITY	8
7	1.5 ITEM VALUES AND DATE TYPES.....	8
8	1.6 LOCALIZATION	8
9	2 PATIENT SUBJECT	9
10	2.1 STANDARD PATIENT CONTEXT DATA ITEMS.....	9
11	2.2 EXAMPLES OF PATIENT SUBJECT ITEMS.....	10
12	3 HL7 DATA TYPE REFERENCE	13
13	<i>IS - coded value for user-defined tables (HL7 Spec 2.8.20).....</i>	<i>13</i>
14	<i>ST - string data (HL7 Spec 2.8.38).....</i>	<i>13</i>
15	<i>PN - person name (HL7 Spec 2.8.28).....</i>	<i>13</i>
16	<i>DLN - driver's license number (HL7 Spec 2.8.11)</i>	<i>13</i>
17	<i>DT - date (HL7 Spec 2.8.13).....</i>	<i>13</i>
18	<i>TS - time stamp (HL7 Spec 2.8.42).....</i>	<i>14</i>

1

2

Preface

3

4

5

6

7

8

This document was prepared by Kyle Marchant, 3M Health Information Systems, on behalf of Health Level Seven's Special Interest Group on Visual Integration (formerly the Clinical Context Object Workgroup --- CCOW). Comments about the organization or wording of the document should be directed to the author (krmarchant@mmm.com). Comments about technical content should be directed to ccow@list.mc.duke.edu.

1 Introduction

The goal of this document is to provide a specification of the standard context data items that shall be supported for patient subject for the HL7 Context Management Architecture (CMA). For the patient subject this document specifies the standard context data items that are available for applications to use in setting and accessing the common clinical context.

1.1 Context Management Document Overview

It is beyond the scope of this document to provide all of the details that are needed in order to fully implement conformant CMA applications and components. The necessary additional details are covered in a series of companion specification documents, starting most notably with the Health Level Seven Context Management Specification, Technology- And Subject- Independent Component Architecture, CM-1.0.

These documents are organized to facilitate the process of defining additional link subjects and to accelerate the process of realizing the CMA using any one of a variety of technologies:

- There is an HL7 context management user interface specification document for each of the user interface technologies with which CMA-enabled applications can be implemented. Each document reflects the user interface requirements established in this document in terms of a technology-specific look-and-feel. Concurrent with the publication of this document, the following document has been developed:

Health Level-Seven Standard Context Management Specification,
User Interface: Microsoft Windows OS, Version CM-1.0

- There is an HL7 context management component technology mapping specification document for each of the component technologies. Each document provides the technology-specific details needed to implement CMA-compliant applications and the associated CMA components, as specified in this document. Concurrent with the publication of this document, the following document has been developed:

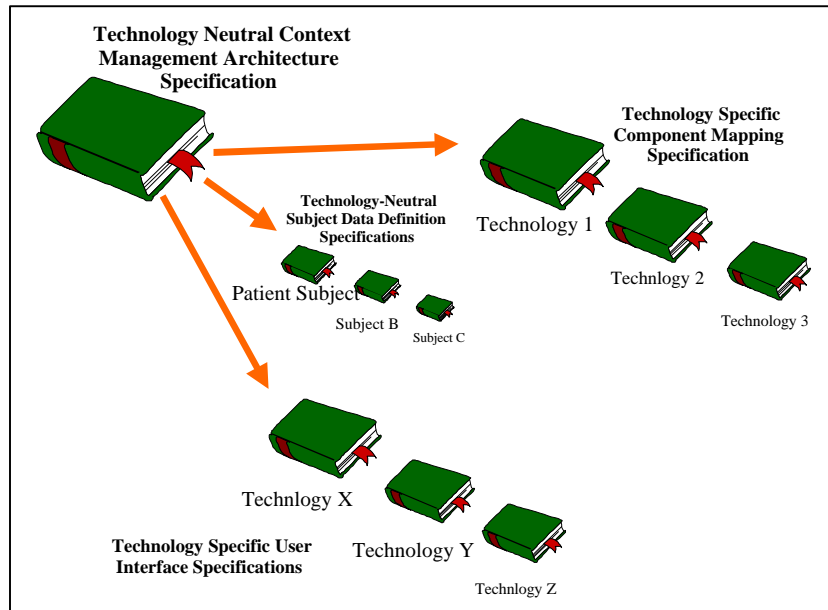
Health Level-Seven Standard Context Management Specification,
Component Technology Mapping: ActiveX, Version CM-1.0

Finally, the context management subjects and technologies that are of interest are determined by the HL7 constituency. There is an HL7 context management data definition specification document for each of the standard link subjects. Each document defines the data elements that comprise a link subject. Concurrent with the publication of this document for the patient subject, the following document has been developed:

Health Level-Seven Standard Context Management Specification,
Data Definition: User Subject, Version CM-1.0

1 The organization of this set of documents is illustrated in Figure 1.

2



3 **Figure 1: Organization of HL7 Context Management Specification Documents**

1.2 Context Data Subject

Context data is grouped by subject. Each subject represents a real-world entity or concept. Each subject is described by a set of context data items. Each context data item is structured as a name/value pair. This document specifies the items for the patient subject. The specific names and data types for each of the patient subject context data items are specified later in this document.

1.3 Context Data Item Format

The general format of a context data item name is:

`Item_subject_label.role.item_name_prefix.optional_item_name_suffix`

Item_subject_label is the name of the subject to which the item belongs.

Role indicates the role of the item, as follows:

- “Id” = standard identifier data, which is used to identify a real-world entity or concept.
- “Co” = standard corroborating data, which is used by applications and/or users to corroborate the identity of a real-world entity or concept.
- “Zz” = non-standard organizationally defined data, the meaning of which is specified by the organization that defined the item.

Item_name_prefix is the name of the item within the context of its subject.

Optional_item_name_suffix is optional for identifier and corroborating data items. It’s purpose is to two-fold:

- For identifier items, the suffix enables multiple items to represent the same logical concept. For example, at a particular site, patients may be identified by multiple medical record numbers. Each item that represents a patient medical record number would have the same item subject label, role, and item name prefix. However, each item name would have a different site-defined item name suffix.
- For non-standard items, the suffix shall always identify the name of organization that defined the item.

The HL7 Standard Context Management Specification, Technology-and-Subject-Independent Component Architecture specification document should be consulted for additional details on the definition and structure of context item names.

1.4 Case Sensitivity

Item names, and item values whose data type is a character string, shall be treated as “case insensitive” unless specifically noted otherwise. This means that unless specifically stated in one of the HL7 subject data definition documents, context participants, context managers, mapping agents etc. shall not rely on the case of a context item name or value when applying decision or comparison logic.

1.5 Item Values and Date Types

Where applicable, the HL7 Version 2.3 Specification for healthcare messaging data elements is used as the basis for context data item names and values.

1.6 Localization

Context data item names shall be in English, regardless of the country and/or location that the context manager and context participants are being used in. This enables those developing both context managers and context participants to code to a known language standard for each context subject area, while still allowing the user interface guidelines to take into account localization issues where appropriate.

2 Patient Subject

The item subject label for the patient subject is “Patient”.

A single patient may be identified using multiple patient subject identifier (id) items. Each item is differentiated by a different site-specific suffix. An application shall be configurable such that it can be instructed on-site as to which suffix (or suffices) it is to use when it interacts with the context manager to set or get patient context data. Use of this suffix, and the values that may be assigned to this suffix, is at the discretion of each healthcare institution at which a context management system is deployed.

2.1 Standard Patient Context Data Items

The standard context data items for the patient subject are described below.

<u>Patient Subject Identifier Item Name</u>	<u>Meaning</u>	<u>HL7 Data Type</u>	<u>Semantic constraints on values</u>	<u>Case Sensitive</u>
<code>Patient.Id.MRN.site_name</code> where <i>site_name</i> is a site-specified name of locale or site, or a set of locales or sites, for which this particular identifier item is valid.	Patient’s medical record number.	ST	none	No. For example, “01JSB0034” and “01jsb0034” are the same medical record numbers.
<code>Patient.Id.NationalIdNumber</code>	Patient’s national identifier number.	ST	none	no
<code>Patient.Id.Alternate.site_name</code> where <i>site_name</i> is a site-specified name of locale or site, or a set of locales or sites, for which this particular identifier item is valid.	Alternate patient identifier.	ST	none	no

An application shall set a value for at least one of items defined above whenever it sets the patient context.

<u>Patient Subject Corroborating Item Name</u>	<u>Meaning</u>	<u>HL7 Data Type</u>	<u>Semantic constraints on values</u>	<u>Case Sensitive</u>
Patient.Co.PatientName	Patient's name.	PN	none	no
Patient.Co.PatientName	Patient's name.	PN	none	no
Patient.Co.AliasName	Alias name for the patient.	PN	none	no
Patient.Co.DateTimeOfBirth	Patient's Date and time of birth.	TS	none	no
Patient.Co.Sex	Patient's gender.	IS	F for female M for male O for other U for unknown	no
Patient.Co.DLN	Patient's drivers license number.	DLN	none	no
Patient.Co.SSN	Patient's Social security number	ST	none	no

1

2 An application may optionally set a value for items defined above when it sets the patient context.

3

4

5 **2.2 Examples of Patient Subject Items**

6 Below are examples of patient subject items:

7

Example Item Names**Example Item Values**

Context Management Specification Data Definition: Patient Subject

Patient.Id.MPI	001KM002130-JJXXX-98
Patient.Id.MRN.St_Elsewhere_Clinic	SEC-KMAR-00hjd7792
Patient.Id.MRN.St_Somewhere_Clinic	SSC-KMAR-00WSB887455
Patient.Co.DateTimeOfBirth	19580317
Patient.Co.PatientName	Marchant^Kyle^^^^

1

3 HL7 Data Type Reference

The item data types referenced in Section 2, Patient Subject, are the same as those specified in the HL7 Version 2.3 Specification, Section 2.8, as described below:

<u>DATA TYPE</u>	<u>DATA TYPE NAME</u>	<u>HL7 Section Reference</u>
IS	Coded Value For User Defined Table	2.8.20
ST	String	2.8.38
PN	Person Name	2.8.28
DLN	Drivers License Number	2.8.11
DT	Date	2.8.13
TS	Time Stamp	2.8.42

The formatting information for each of these fields is specified below, with its corresponding description and HL7 specification section identifier. Only the encoding characters and escape sequences indicated below shall be used:

IS - coded value for user-defined tables (HL7 Spec 2.8.20)

ST - string data (HL7 Spec 2.8.38)

PN - person name (HL7 Spec 2.8.28)

Components: <family name (ST)> ^ <given name (ST)> ^ <middle initial or name (ST)> ^ <suffix (e.g., JR or III) (ST)> ^ <prefix (e.g., DR) (ST)> ^ <degree (e.g., MD) (ST)>

DLN - driver's license number (HL7 Spec 2.8.11)

Components: <license number (ST)> ^ <issuing state, province, country (IS)> ^ <expiration date (DT)>

DT - date (HL7 Spec 2.8.13)

Format: YYYY[MM[DD]]

1 **TS - time stamp (HL7 Spec 2.8.42)**

2 Format: YYYY[MM[DD[HHMM[SS[.S[S[S[S]]]]]]]]][+/-ZZZZ]^<degree of precision>

3

4

1
2
3
4
5
6
7
8
9
10
11

12
13
14
15
16
17

Health Level Seven Standard

**Context Management Specification
Data Definition: User Subject
Version CM-1.0**

DOCUMENT ID: HL7SIGVI_3_3_99
REVISION ID: March 17, 1999
FILE NAME: hl7_sigvi_user_cm_1_0 .doc
SUPERCEDES: n/a

Copyright 1999 Health Level Seven

1	Contents	
2	1 INTRODUCTION.....	5
3	1.1 CONTEXT MANAGEMENT DOCUMENT OVERVIEW.....	5
4	1.2 CONTEXT DATA SUBJECT.....	7
5	1.3 CONTEXT DATA ITEM FORMAT.....	7
6	1.4 CASE SENSITIVITY	8
7	1.5 ITEM VALUES AND DATE TYPES	8
8	1.6 LOCALIZATION	8
9	2 USER SUBJECT.....	9
10	2.1 STANDARD USER CONTEXT DATA ITEMS	9
11	2.2 EXAMPLES OF USER SUBJECT ITEMS	10
12	3 HL7 DATA TYPE REFERENCE	11
13	<i>IS - coded value for user-defined tables (HL7 Spec 2.8.20).....</i>	<i>11</i>
14	<i>ST - string data (HL7 Spec 2.8.38).....</i>	<i>11</i>
15	<i>PN - person name (HL7 Spec 2.8.28).....</i>	<i>11</i>
16	<i>DLN - driver's license number (HL7 Spec 2.8.11).....</i>	<i>11</i>
17	<i>DT - date (HL7 Spec 2.8.13).....</i>	<i>11</i>
18	<i>TS - time stamp (HL7 Spec 2.8.42).....</i>	<i>12</i>

1

2

Preface

3

4

5

6

7

8

This document was prepared by Kyle Marchant, 3M Health Information Systems, on behalf of Health Level Seven's Special Interest Group on Visual Integration (formerly the Clinical Context Object Workgroup --- CCOW). Comments about the organization or wording of the document should be directed to the author (krmarchant@mmm.com). Comments about technical content should be directed to ccow@list.mc.duke.edu.

1 Introduction

The goal of this document is to provide a specification of the standard context data items that shall be supported for user subject for the HL7 Context Management Architecture (CMA). For the user subject this document specifies the standard context data items that are available for applications to use in setting and accessing the common clinical context.

1.1 Context Management Document Overview

It is beyond the scope of this document to provide all of the details that are needed in order to fully implement conformant CMA applications and components. The necessary additional details are covered in a series of companion specification documents, starting most notably with the Health Level Seven Context Management Specification, Technology- And Subject- Independent Component Architecture, CM-1.0.

These documents are organized to facilitate the process of defining additional link subjects and to accelerate the process of realizing the CMA using any one of a variety of technologies:

- There is an HL7 context management user interface specification document for each of the user interface technologies with which CMA-enabled applications can be implemented. Each document reflects the user interface requirements established in this document in terms of a technology-specific look-and-feel. Concurrent with the publication of this document, the following document has been developed:

Health Level-Seven Standard Context Management Specification,
User Interface: Microsoft Windows OS, Version CM-1.0

- There is an HL7 context management component technology mapping specification document for each of the component technologies. Each document provides the technology-specific details needed to implement CMA-compliant applications and the associated CMA components, as specified in this document. Concurrent with the publication of this document, the following document has been developed:

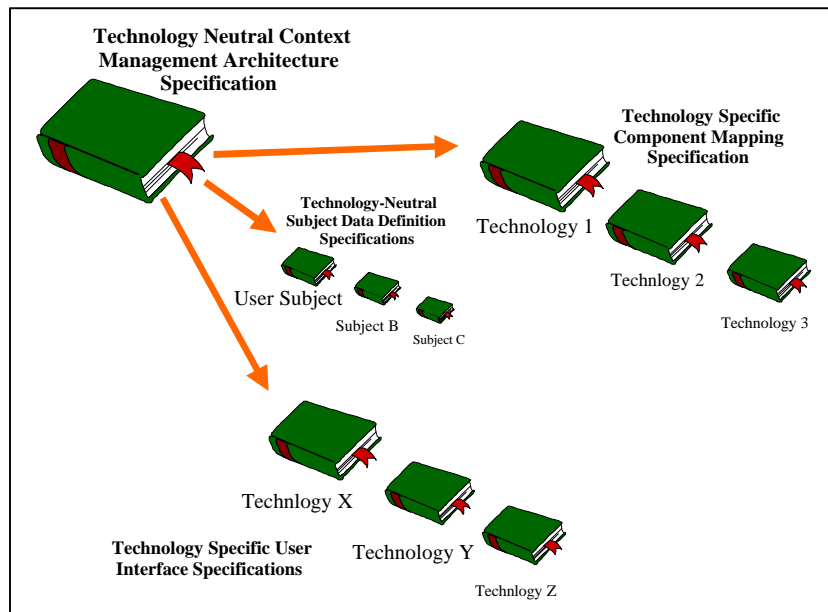
Health Level-Seven Standard Context Management Specification,
Component Technology Mapping: ActiveX, Version CM-1.0

Finally, the context management subjects and technologies that are of interest are determined by the HL7 constituency. There is an HL7 context management data definition specification document for each of the standard link subjects. Each document defines the data elements that comprise a link subject. Concurrent with the publication of this document for the user subject, the following document has been developed:

Health Level-Seven Standard Context Management Specification,
Data Definition: Patient Subject, Version CM-1.0

1 The organization of this set of documents is illustrated in Figure 1.

2



3 **Figure 1: Organization of HL7 Context Management Specification Documents**

1.2 Context Data Subject

Context data is grouped by subject. Each subject represents a real-world entity or concept. Each subject is described by a set of context data items. Each context data item is structured as a name/value pair. This document specifies the items for the user subject. The specific names and data types for each of the user subject context data items are specified later in this document.

1.3 Context Data Item Format

The general format of a context data item name is:

`Item_subject_label.role.item_name_prefix.optional_item_name_suffix`

Item_subject_label is the name of the subject to which the item belongs.

Role indicates the role of the item, as follows:

- “Id” = standard identifier data, which is used to identify a real-world entity or concept.
- “Co” = standard corroborating data, which is used by applications and/or users to corroborate the identity of a real-world entity or concept.
- “Zz” = non-standard organizationally defined data, the meaning of which is specified by the organization that defined the item.

Item_name_prefix is the name of the item within the context of its subject.

Optional_item_name_suffix is optional for identifier and corroborating data items. It’s purpose is to two-fold:

- For identifier items, the suffix enables multiple items to represent the same logical concept. For example, at a particular site, patients may be identified by multiple medical record numbers. Each item that represents a patient medical record number would have the same item subject label, role, and item name prefix. However, each item name would have a different site-defined item name suffix.
- For non-standard items, the suffix shall always identify the name of organization that defined the item.

The HL7 Standard Context Management Specification, Technology-and-Subject-Independent Component Architecture specification document should be consulted for additional details on the definition and structure of context item names.

1.4 Case Sensitivity

Item names, and item values whose data type is a character string, shall be treated as “case insensitive” unless specifically noted otherwise. This means that unless specifically stated in one of the HL7 subject data definition documents, context participants, context managers, mapping agents etc. shall not rely on the case of a context item name or value when applying decision or comparison logic.

1.5 Item Values and Date Types

Where applicable, the HL7 Version 2.3 Specification for healthcare messaging data elements is used as the basis for context data item names and values.

1.6 Localization

Context data item names shall be in English, regardless of the country and/or location that the context manager and context participants are being used in. This enables those developing both context managers and context participants to code to a known language standard for each context subject area, while still allowing the user interface guidelines to take into account localization issues where appropriate.

2 User Subject

The item subject label for the user subject is “User”.

A single user may be identified using multiple user subject identifier (id) items. Each item is differentiated by a different application-specific suffix. An application shall be configurable such that it can be instructed on-site as to which suffix (or suffices) it is to use when it interacts with the context manager to set or get user context data. Use of this suffix, and the values that may be assigned to this suffix, is at the discretion of each healthcare institution at which a context management system is deployed.

2.1 Standard User Context Data Items

The standard context data items for the user subject are described below.

<u>User Subject Identifier Item Name</u>	<u>Meaning</u>	<u>HL7 Data Type</u>	<u>Semantic constraints on values</u>	<u>Case Sensitive</u>
User.Id.Logon.application_name where <i>application_name</i> is a site-specified name of an application, or a set of applications, for which this particular identifier item is valid.	User's logon name.	ST	none	Value is case sensitive. For example, “ksmith” and “Ksmith” are two different logon id values.

An application shall set a value for the item defined above whenever it sets the user context.

<u>User Subject Corroborating Item Name</u>	<u>Meaning</u>	<u>HL7 Data Type</u>	<u>Semantic constraints on values</u>	<u>Case Sensitive</u>
User.Co.Name	User's name	PN	none	no

An application may optionally set a value for items defined above when it sets the user context.

1

2 **2.2 Examples of User Subject Items**

3 Below are examples of user subject items:

4

Example Item Names	Example Item Values
User.Id.Logon.3M_Clinical_Workstation	k_marchant
User.Id.Logon.Logician	kylem
User.Id.Logon.Carevue	KM01230
User.Co.Name	Kyle Marchant

5

3 HL7 Data Type Reference

The item data types referenced in Section 2, User Subject, are the same as those specified in the HL7 Version 2.3 Specification, Section 2.8, as described below:

<u>DATA TYPE</u>	<u>DATA TYPE NAME</u>	<u>HL7 Section Reference</u>
IS	Coded Value For User Defined Tables	2.8.20
ST	String	2.8.38
PN	Person Name	2.8.28
DLN	Drivers License Number	2.8.11
DT	Date	2.8.13
TS	Time Stamp	2.8.42

The formatting information for each of these fields is specified below, with its corresponding description and HL7 specification section identifier. Only the encoding characters and escape sequences indicated below shall be used:

IS - coded value for user-defined tables (HL7 Spec 2.8.20)

ST - string data (HL7 Spec 2.8.38)

PN - person name (HL7 Spec 2.8.28)

Components: <family name (ST)> ^ <given name (ST)> ^ <middle initial or name (ST)> ^ <suffix (e.g., JR or III) (ST)> ^ <prefix (e.g., DR) (ST)> ^ <degree (e.g., MD) (ST)>

DLN - driver's license number (HL7 Spec 2.8.11)

Components: <license number (ST)> ^ <issuing state, province, country (IS)> ^ <expiration date (DT)>

DT - date (HL7 Spec 2.8.13)

Format: YYYY[MM[DD]]

1 **TS - time stamp (HL7 Spec 2.8.42)**

2 Format: YYYY[MM[DD[HHMM[SS[.S[S[S[S]]]]]]]]][+/-ZZZZ]^<degree of precision>

3

4

Health Level Seven Standard

Context Management Specification User Interface: Microsoft Windows OS Version CM-1.0

DOCUMENT ID: HL7SIGVI_3_5_99

REVISION ID: March 17, 1999

FILE NAME: hl7_sigvi_windows_cm_1_0.doc

SUPERCEDES: n/a

Copyright 1999 Health Level Seven

1	Contents	
2	1. INTRODUCTION.....	4
3	2. TECHNOLOGY NEUTRALITY.....	5
4	3. INTERFACE APPROACHES.....	6
5	4. ARCHITECTURE.....	7
6	4.1 JOINING A CONTEXT.....	8
7	4.2 CHANGING THE CONTEXT.....	8
8	4.3 REJOINING THE CONTEXT.....	10

Preface

This document was prepared by Jeff Amfahr, Component Software International., on behalf of Health Level Seven's Special Interest Group on Visual Integration (formerly the Clinical Context Object Workgroup --- CCOW). Comments about the organization or wording of the document should be directed to the author (amfahr@csi-corporate.com). Comments about technical content should be directed to the Clinical Context Object Workgroup ccow@list.mc.duke.edu.

1. Introduction

This document specifies the user interface when using the common clinical context component as specified by the Clinical Context Object Workgroup (CCOW). It is assumed that the reader is familiar with the architectural specification for that software. The user interface explained in this document is intended to describe only those user interface features that are directly concerned with the clinical context component itself. Since a goal of this software is to be as seamless and natural for the user as possible, only the minimum set of user interface features is specified. No attempt to standardize the overall look and feel of clinical applications is made.

It is intended that this user interface be applicable to the full breadth of user interfaces in use today. The clinical context object itself has no specified user interface. All the information specified here is implemented by the clients of the clinical context object. Below is a table that indicates which of the following user interface concepts are required in order to be considered CCOW compliant and which are recommendations that should be used when possible and appropriate.

	Required	Recommended
Context status	Continuous, consistent, representation of state of context link using either text or icons or both.	Use icons and text
Status icons	If using icon, use standard icons.	
Status text	If using text, use standard text.	
Status location		At least one of the context link indicators should be prominently displayed.
Individual context status		Some visual indicator near the most related data to context data itself.
Joining context		Show text to indicate that the application is joining the context manager. If implementing this recommendation it is a requirement that you use the specified text.
Changing context	The dialog for indicating that there are conditional, busy or mapping agent problem responses to the survey is a required element.	
Rejoining Context	There must be a mechanism for rejoining the context if there is any way to break the context link from the application.	On rejoining, the user should be given the option of using the application's current context values or the values from the global context.
Modifying link status		Clicking on the link indicator instigates reversing the state of the link.

Table 1 Required and recommended features

2. Technology Neutrality

Since user interface tends to be a technology specific detail, this document attempts to address the user interface from a generic standpoint. The primary assumption is that the application uses a graphical user interface (as opposed to a command line) that provides support for overlapping windows. This would include all flavors of Windows™ and web-based applications, for example. Because there is no desire to specify the overall look and feel of clinical applications, the guidelines use only standard constructs available in all these systems.

The ability to internationalize this specification has been considered, but this initial version contains only English strings. Some minor modifications may be necessary for other languages, especially multibyte language systems. For example, maximum string lengths may be increased.

3. Interface Approaches

The overall architecture of the clinical context component is described in the Architectural Specification, but several features described in that document are especially applicable in the user interface scope.

- An application may choose to defer applying a context change until some time in the future. For example, an application that retrieves large medical image files (that require substantial processing) might choose to not retrieve images each time a different patient is selected as part of the clinical context. Instead, the application might wait for an explicit directive or gesture from the user before actually retrieving the image. An application that behaves in this manner must be sure that it does not show data for an earlier context. Blanking-out its data displays or minimizing itself are possible ways that this can be accomplished.
- An application for which a change in the context might result in the loss of work performed by the user can request that the user explicitly decide whether to proceed with the context change anyway, or to cancel the change. The solicitation of user input is performed by the application that is being used to change the context. The solicitation includes an identification of the application for which work might be lost and a description of the work that might be lost. An application that behaves in this manner is expected to be able to discard its work in progress and apply the context changes if instructed to do so. For example, a medication ordering application might indicate that the inputs for a medication order, which has not yet been completed by the user, will be lost if the context is changed to a different patient.
- When an application is unable to respond to a context change, perhaps because the user left it waiting for user input, the user is asked to explicitly decide how to proceed. The solicitation of user input is performed by the application that is being used to change the context. The solicitation includes the identification of the non-responsive application and indicates that the application cannot respond to a context change. For patient safety reasons, when there are applications that cannot respond to the changes, context changes will not be automatically applied to the applications that share a common context.
- When it is not desirable or possible for context changes to be automatically applied, either because there are applications for which work might be lost or there are busy applications that cannot be notified about context changes, the user can explicitly interact with these applications to correct the situation, and then apply the context changes. For example, the user might complete or terminate a dialog that was left open in order to enable an application to apply the context changes.
- When it is not desirable or possible for context changes to be automatically applied, the user can also decide to apply the context change only to the application that is being used to change the context. The decision to do this is typically in response to an interruption during which the user needs to momentarily divert his attention to a different context for a specific application. The application is, in effect, disconnected from the common context, and must clearly indicate this fact to the user in a visual manner. The application can be subsequently instructed by the user to reconnect and apply the common context. The common context may have changed between the time the application was disconnected and the time it is reconnected to the common context.

4. Architecture

Although the desire is for the common context to be seamless, the user does need to be aware when the component is and is not a part of the context. A common vocabulary for the context is required so users can be familiar with the meaning of common dialogs. In all text that the user sees, the context is referred to as the “clinical link.” When the component is currently using the common context, the clinical link is “on”. When a component leaves the context, the clinical link is “broken”.

It is vital that the current state of the link be visible and apparent to the user. An application can use two methods to indicate the status of the link. The first is textual, in which case the text should be “Clinical link on”, “Clinical link changing”, or “Clinical link broken”. The second method is using icons which are shown below. At least one of these indicators must be used, although both are recommended. If both indicators are used, they must always show the same state for the link.

Please note: the picture objects shown below were created from the 40 pixel wide by 21 pixel high bitmaps that are recommended for CCOW-compliant applications. Programmers can copy them from this document, past them into Paint Stop Pro, or any other bitmap editor and save them as a bitmap file (.BMP) extension. For this reason the bitmaps are not published separately.



Figure 2 Clinical link on icon



Figure 3 Clinical link broken icon



Figure 4 Clinical link changing icon

If the application provides a user interface, this overall status indicator should be located prominently and persistently on the screen. It is also recommended that the application indicate the subjects (i.e. areas of clinical context) to which it is linked. If the application indicates these linked subjects, there are two recommended formats:

- The first format is to provide the indicator via a menu item, preferably in the same menu that allows the user to re-join the context. This menu should be hierarchical, with the main menu item labeled as “Supported clinical subjects” and the submenus having labels for each subject area (for example, “Patient”, “User”, etc.).
- The alternative format (which may be used in conjunction with the first approach) is to provide a visual indicator near a representative piece of clinical data. For example, an icon near the patient name to indicate that the patient subject is supported, an icon near the user name to indicate that the user subject is supported, etc. This icon should be similar to the normal clinical link status icons shown above, however they may be smaller in order to accommodate the actual area of the screen used. These icons should change when the link is broken to indicate that the subject is no longer linked.

There are three stages to using the common clinical context:

- Joining a context
- Changing a context
- Rejoining the context.

4.1 *Joining a context*

It is assumed that most applications will attempt to join the common context when launching, with no user intervention. When joining the context, a transaction may be in progress. No components are allowed to join the context while a transaction is in progress. If the component chooses to block on the call to join the transaction (by setting the wait parameter to true), then a dialog must be displayed informing the user that the thread is blocked. This dialog must only be displayed if the transaction takes longer than one second. This dialog should be a standard system dialog with the text “Establishing clinical link”. Once the context has been joined, this dialog should close itself with no further user interaction. At this point, the visual indicator for link status should indicate that the link is on, or if the context cannot be joined for whatever reason (e.g., the TooManyParticipants exception is raised), the link status should indicate a broken link.

4.2 *Changing the Context*

Once the context has been joined, the component will normally be in synch with the current context data. When another application desires to change the context, the context manager informs other applications by calling their ContextChangesPending method. At this point, the other applications need to change their status to indicate the patient link is changing. This tells the user that no context data change is presently possible. At some later point, the context manager will call either the ContextChangesAccepted or ContextChangesCanceled method. The link status should then change to indicate it is on.

If the current component desires to change the current context itself, it needs to inform the other CCOW clients. This document does not specify how new context data is selected by the user. Once the data has been entered, however, the context manager calls a mapping agent (if available) to confirm and add to the data. If the mapping agent finds no inconsistencies, the context manager surveys the registered applications. If all applications can change the current context, then the context data is changed as described above. If, however, any of the applications are unable to change because they are busy or because they require user interaction, or the mapping agent finds an inconsistency in the identifying information, the user must be queried for the appropriate action.

An application that requires user interaction (i.e., conditionally accepts) sends the context manager a string stating the reason and the potential consequences if the data change occurs. The format of that string should be declarative (i.e., state what the consequences will be, rather than query the user for information) and should be less than 64 characters in length. Some examples are “The open patient record will not be saved”, “Annotations for record 123 will be lost” or “Current information will be saved”. If an application is busy, meaning it is unable to respond to the survey, the context manager will return the string “This application appears to be busy” for that application.

In either of these cases, the component that requested the context change should then present the user with a dialog (Figure 5) containing the following text: “There may be a problem changing the current clinical data. The following application(s) reported a problem. Would you like to continue with the change?” followed by a list showing the application name(s) and reason(s) returned. This dialog contains three action buttons that must be labeled “OK”, “Cancel”, and “Break link”. All of these choices dismiss the dialog box. The default choice is “Cancel”. If at least one of the responses indicates an application is busy, then the “OK” choice must be disabled. If an application does not support the choice of “Cancel” (for example, because choosing a patient is a difficult process), then that choice should be disabled and “Break link” should be the default choice (Figure 6). Even if no other choice is possible other than “Break link”, the dialog should be presented so that the user is aware they are breaking the patient link.

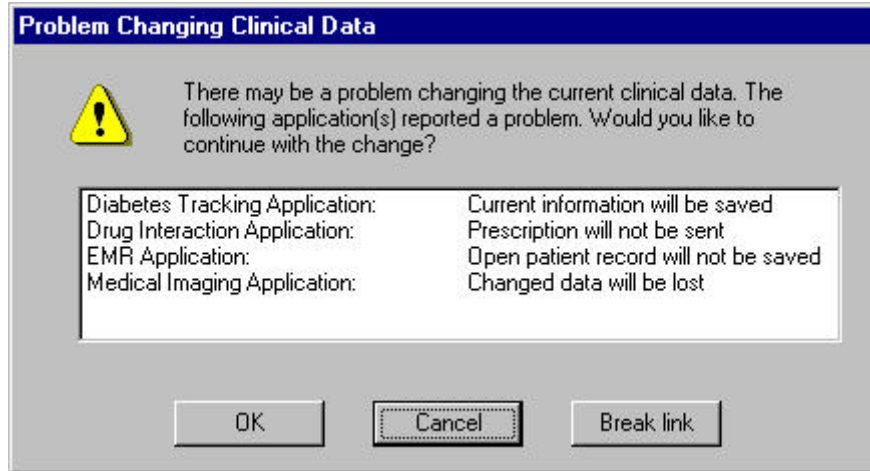


Figure 5 Context Dialog When No Applications Are Busy

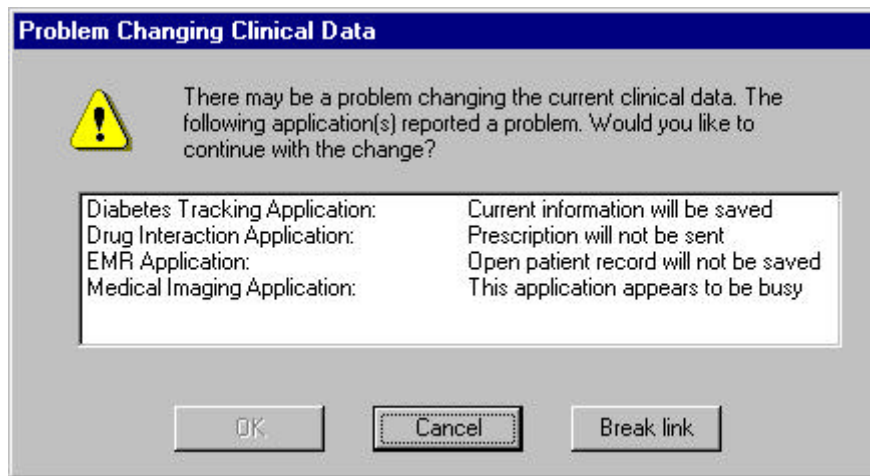


Figure 6 Context Dialog When At Least One Application is Busy

If the mapping agent finds an inconsistency in the identifying data set by the instigating application, it should send the context manager a string stating the reason for the inconsistency. The format of that string should be declarative and should be less than 64 characters in length. Some examples are “The IDs map to two different patients”, or “Patient identifiers are not consistent with data recorded”. The context manager should log these for future reporting as appropriate. The requesting application will receive a programmatic response identical to that received in the case of a busy application. The application name returned will be “Data mapping System” and the string returned by the context manager will be “Some of the IDs used were conflicting”. This result should be reported as described above for a busy application.

An application may be unable to interpret some or all of the context data that is established when another application changes the common context. For example, the selected patient might not be a patient known to all of the applications. In this case, the application should clearly indicate that it is unable to apply the context data. It should either minimize itself or blank-out its data displays. It is also acceptable for an application to do this if it is not the current application and accessing its data is a lengthy process. It would begin this process when it became the front most application. Note in either case that the application is still part of the common context and its status should indicate this.

4.3 Rejoining the context

When an application is currently not a part of the common context (because the user has explicitly broken the link), there must be a mechanism for the user to rejoin the context (e.g. a menu item). If a menu item or equivalent is used it must be labeled “Rejoin Clinical Link”. This menu will include a sub-menu with two choices: “Use this applications’ data” and “Use global data”. When an application is part of the context this menu item or equivalent should be disabled.