

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

MICROSOFT CORPORATION

IPR2025-00270

U.S. Patent No. 8,073,911

**DECLARATION OF DR. JOHN BLACK, JR.
UNDER 37 C.F.R. § 1.68 IN SUPPORT OF
PETITION FOR *INTER PARTES* REVIEW**

TABLE OF CONTENTS

	Page
I. INTRODUCTION	14
II. QUALIFICATIONS AND PROFESSIONAL EXPERIENCE	16
III. LEVEL OF ORDINARY SKILL IN THE ART	19
IV. RELEVANT LEGAL STANDARDS	20
V. OVERVIEW OF THE 911 PATENT	25
A. The Disclosures of the 911 Patent.....	25
B. The Prosecution History of the 911 Patent.....	35
VI. CLAIM CONSTRUCTION	36
VII. IDENTIFICATION OF HOW THE CLAIMS ARE UNPATENTABLE.....	37
VIII. TECHNOLOGY BACKGROUND.....	37
A. Electronic mail systems.....	37
B. Message Databases.....	38
C. Pointers and References	39
D. De-duplication	44
IX. THE PRINCIPAL PRIOR ART	49
A. Rohall	49
B. Schiavone	57
C. Gomes.....	65
X. CLAIMS 1-13, 15, 18, 20-21, AND 27 ARE UNPATENTABLE	66
A. Ground 1: The Combination of Rohall and Schiavone renders obvious claims 1-8, 10-13, 15, 18, and 20-21	66
B. The Combination of Rohall and Schiavone	66
1. Independent Claims 1, 10, and 13:	75
a. Preamble: Electronic Messaging System	75
[1.0] A computerized messaging server in an electronic messaging system, comprising:	75
[10.0] A computer program product having a non-transitory computer-readable medium	

	having embodied thereon program code for use in an electronic messaging system, the program code comprising:	75
	[13.0] A computer-implemented method of managing messages in an electronic messaging system, comprising:	76
b.	Messaging Module Controlling a Message Database.....	80
	[1.1] a messaging module adapted to control a message database storing messages sent among users of the messaging system,.....	80
	[10.1] a messaging module adapted to control a message database storing messages sent among users of the messaging system to at least one other end user of the messaging system, ...	80
	[13.2] receiving a message sent by one end user of the messaging system to at least one other end user of the messaging system, ...	80
c.	Sent Message comprising a Message Container	86
	[1.2] at least one of the sent messages stored in the message database comprising a message container containing relational references pointing to a plurality of sent submessages stored externally to the message container in the message database,.....	86
	[10.2] at least one of the sent messages comprising a message container containing relational references pointing to a plurality of sent submessages stored in the message database externally to the message container,.....	87
	[13.3] the sent message comprising a message container containing relational references pointing to a plurality of sent submessages stored in a message database externally to the message container, ...	87

- d. Sent submessages Are a Reply-To or a Forward-Of.....105
 - [1.3] at least one of the sent submessages being one of a reply to and a forward of other sent submessages of the sent message; and.....105
 - [10.3] at least one of the sent submessages being one of a reply to and a forward of other sent submessages of the sent message; and.....106
 - [13.4] the plurality of sent submessages including a first sent submessage and a second sent submessage, at least one of the sent submessages being one of a reply to and a forward of another sent submessage of the plurality of sent submessages; and 106
- e. Governance Module controlling a Governance Policy 109
 - [1.4] a governance module adapted to control a governance policy database storing a compliance policy describing rules applicable to the sent submessages, ...109
 - [10.4] a governance module adapted to control a governance policy database storing a compliance policy describing rules applicable to the sent submessages, ...109
 - [13.1] defining a compliance policy describing rules applicable to sent submessages of sent messages in the electronic messaging system; 109
- f. Governance module determines rules described by Compliance Policy.....118
 - [1.5] the governance module further adapted to determine rules described by the compliance policy that are applicable to ones of the plurality of sent submessages,.....118
 - [10.5] the governance module further adapted to determine rules described by the compliance

	policy that are applicable to ones of the plurality of sent submessages;.....	118
g.	Rules of the compliance policy selectively apply to submessages	121
	[1.6] wherein different rules of the compliance policy are applicable to different ones of the plurality of sent submessages;.....	121
	[10.6] wherein different rules of the compliance policy are applicable to different ones of the plurality of sent submessages;.....	121
	[13.5] selectively applying rules of the compliance policy to the sent submessages	121
h.	Entities in a Population Structure.....	123
	[1.7] wherein the messaging system is utilized by a set of entities in a population structure and.....	123
	[10.7] wherein the messaging system is utilized by a set of entities in a population structure and.....	123
	[13.6] wherein the messaging system is utilized by a set of entities in a population structure and.....	124
i.	Rule Elements of Compliance policy.....	126
	[1.8] wherein a rule of the compliance policy includes one or more rule elements describing whether the rule is applicable to a sent submessage, ...	126
	[10.8] wherein a rule of the compliance policy includes one or more rule elements describing whether the rule is applicable to a sent submessage, ...	126
	[13.7] wherein a rule of the compliance policy includes one or more rule elements describing whether the rule is applicable to a sent submessage, ...	126
j.	Rule Elements selected from Source/Target/Rank..	129

	[1.9] at least one of the rule elements selected from the group consisting of: a source element describing a population structure entity that is a source of the sent submessage; a target element describing a population structure entity that is a recipient of the sent submessage; and a rank describing a ranking of the rule relative to other rules in the compliance policy.	129
	[10.9] at least one of the elements selected from the group consisting of: a source element describing a population structure entity that is a source of the sent submessage; a target element describing a population structure entity that is a recipient of the sent submessage; and a rank describing a ranking of the rule relative to other rules in the compliance policy.	129
	[13.8] at least one of the elements selected from the group consisting of: a source element describing a population structure entity that is a source of the sent submessage; a target element describing a population structure entity that is a recipient of the sent submessage; and a rank describing a ranking of the rule relative to other rules in the compliance policy.	129
2.	Dependent Claims	131
a.	Jobcode	131
	[2.0] The messaging server of claim 1, wherein the compliance policy describes a population structure of end-users of the messaging system, a set of jobcodes with which the sent submessages can be classified, and a set of rules applicable to the population structure, jobcodes, and sent submessages.	131
b.	Jobcode Rules applicable to all sent submessages..	135
	[3.0] The messaging server of claim 2, wherein a jobcode describes a set of jobcode rules	

	applicable to all sent submessages classified with the jobcode.....	135
c.	Associate a jobcode with a sent submessage	136
	[4.0] The messaging server of claim 2, wherein the messaging module is further adapted to receive an indication of a jobcode and of a sent submessage from an end-user, and is further adapted to associate the jobcode with the sent submessage responsive to receipt of the indication of the jobcode from the end-user.....	136
d.	Communication and Readership Rules	138
	[5.0] The messaging server of claim 1, wherein the rules described by the compliance policy include one or more types of rules from the set consisting of: communication rules describing which population structure entities can communicate with other population structure entities via the messaging system; and readership rules describing which population structure entities can read messages from other population structure entities via the messaging system.....	138
e.	Validity and Retention Rules.....	142
	[6.0] The messaging server of claim 1, wherein the rules described by the compliance policy include one or more types of rules from the set consisting of: validity rules describing a time period for which a sent submessage or entity in the population structure is useable; and retention rules describing a time period after which a sent submessage or entity in the population structure should be deleted from the messaging system	142
f.	Prior compliance policy active at a prior time	143
	[7.0] The messaging server of claim 1, wherein the governance policy database further stores at least one prior compliance policy active at a time prior to a time at which the compliance policy	

- was active, and wherein the sent submessages of the stored messages are associated with a compliance policy that was active at a time that content of the sent submessages was created. 144
- g. Rules are associated with submessages.....145
- [8.0] The messaging server of claim 7, wherein at least one of the messages stored in the message database has a first sent submessage associated with a first compliance policy and a second sent submessage associated with a second compliance policy different from the first compliance policy, and wherein the governance module applies the rules described by the first compliance policy to the first sent submessage and applies the rules described by the second compliance policy to the second sent submessage. 146
- h. Communication, Readership, Action, and Search Rules 148
- [11.0] The computer program product of claim 10, wherein the rules described by the compliance policy include one or more types of rules from the set consisting of: communication rules describing which population structure entities can communicate with other population structure entities via the messaging system; readership rules describing which population structure entities can read messages from other population structure entities via the messaging system; action rules describing actions a population structure entity can perform on a message component in the -messaging system; and search rules describing which population structure entities can view which message components when searching the messaging components stored by the messaging system. 148
- i. Validity and Retention Rules.....148

	[12.0] The computer program product of claim 10, wherein the rules described by the compliance policy include one or more types of rules from the set consisting of: validity rules describing a time period for which a sent submessage or entity in the population structure is useable; and retention rules describing whether a sent submessage or entity in the population structure should be deleted from the messaging system.	148
j.	Group and Rules of Entities	149
	[15.0] The computer-implemented method of claim 13, wherein the compliance policy describes the set of entities in the population structure, the method further comprising: defining a group containing a set of entities, the entities in the group defined by their locations in a hierarchy defined by the hierarchical population structure; and defining a set of rules applicable to the entities in the group.....	149
k.	Communication, Readership, Action, and Search Rules	151
	[18.0] The computer-implemented method of claim 13, wherein the rules described by the compliance policy include one or more types of rules from the set consisting of: communication rules describing which population structure entities can communicate with other population structure entities via the messaging system;	151
	readership rules describing which population structure entities can read sent messages from other population structure entities via the messaging system;	152
	action rules describing actions a population structure entity can perform on a sent submessage in the messaging system; and.....	152

	search rules describing which population structure entities can view which sent submessages when searching the sent submessages stored by the messaging system.....	152
1.	Action Rules	152
	[20.0] The computer-implemented method of claim 13, wherein the rules described by the compliance policy include action rules describing actions an entity can perform on a sent submessage in the messaging system, and wherein the action rules describe a number of times within a given time period that the entity can perform a specified action.....	152
m.	Group and Rules of Entities	153
	[21.0] The computer-implemented method of claim 13, further comprising:	153
	defining a group of entities in the population structure;	153
	defining group rules applicable to sent submessages related to the group of entities; and ...	153
	applying the group rules to a sent submessage related to the group of entities.....	154
C.	Ground 2: The Combination of Rohall, Schiavone, and Gomes renders obvious claims 9 and 27	156
1.	The Combination of Rohall, Schiavone, and Gomes	156
2.	Dependent Claim 9	158
a.	Search Rules	159
	[9.0] The messaging server of claim 1, wherein the rules described by the compliance policy include search rules describing which population structure entities can view which sent submessages when searching the sent submessages stored by the messaging system.....	159
3.	Independent Claim 27	162

a.	Sent submessages stored exactly once within the email database	163
	[27.4] each of the sent submessages is stored exactly once within the email database, and upon receipt of a reply to a first email having a first email container containing a first sent submessage, the reply is added to the message container as an additional sent submessage but the first sent submessage is not duplicated within the email database; and.....	163
XI.	OBJECTIVE INDICIA OF NON-OBVIOUSNESS	167
XII.	CONCLUSION.....	167

EXHIBIT LIST

No.	Exhibit Description
1001	U.S. Patent No. 8,073,911 (“the 911 Patent”)
1002	File History of U.S. Patent No. 8,073,911
1003	Declaration of Dr. John R. Black, Jr.
1004	CV of Dr. John R. Black, Jr.
1005	U.S. Patent Pub. No. 2003/0163537 (“Rohall”)
1006	U.S. Patent Pub No. 2002/0120600 (“Schiavone”)
1007	File History of U.S. Patent Pub No. 2002/0120600 (“Schiavone-FH”)
1008	Document History of U.S. Patent Pub No. 2002/0120600 (“Schiavone-DH”)
1009	T. Denehy et al., “Duplicate Management for Reference Data,” IBM Research Report, RJ 10305 (A0310-017) October 7, 2003
1010	Sterling’s Preliminary Infringement Contentions (PIC)
1011	U.S. Patent Pub. No. 2003/0167310 (“Moody310”)
1012	U.S. Patent Pub. No. 2003/0177190 (“Moody190”)
1013	<i>Collaboration Support in Networked Distance Learning</i> , B. Krämer et al, INTERNATIONAL COMPUTER SCIENCE INSTITUTE, May 1997
1014	K. Sierra and B. Bates, “Head First Java: Your Brain on Java – A Learner’s Guide,” 1 st Edition, 2003
1015	U.S. Patent Pub. No. 2004/0177271 (“Arnold”)
1016	U.S. Patent No. 8,375,008 (“Gomes”)
1017	U.S. Patent Pub. No. 2004/260698 (“MacMillan”)
1018	U.S. Patent No. 6,978,274 (“Gallivan”)
1019	<i>Filtering Spam Using Search Engines</i> , O. Kolesnikov et al., published as Georgia Tech Technical Report, GIT-CC-03-58
1020	U.S. Patent No. 7,380,039 (“Miloushev”)
1021	“Main memory caching of Web documents,” E.Makatos, Computer Networks and ISDN Systems 28 (1996) 893-905

TABLE OF ABBREVIATIONS AND CONVENTIONS

Abbreviation	Meaning
911 Patent	Ex.1001: U.S. Patent No. 8,073,911
IPR	<i>inter partes</i> review
Petitioner	Microsoft Corporation (“Microsoft”)
Patent Owner or PO	Sterling Computers Corporation. (“Sterling”)
Preliminary Infringement Contentions	PIC: Ex.1010
<i>xx:yy–zz</i>	column <i>xx</i> , lines <i>yy</i> to <i>zz</i>

1. I, John R. Black, Jr., Ph.D., declare as follows:

I. INTRODUCTION

2. I am making this declaration in support of Petitioner’s request for an *Inter Partes* Review of U.S. Patent No. 8,073,911 (“the 911 Patent”) to Marston et al.

3. I am being compensated for my work in this matter at my standard hourly rate. I am also being reimbursed for reasonable and customary expenses associated with my work and testimony in this proceeding. My compensation is not contingent on the outcome of this matter or the specifics of my testimony.

4. I have been asked to provide my opinions regarding whether the subject matter of claims 1-8, 10-16, 18, 20-25, and 27-28 (“the Challenged Claims”) of the 911 Patent are obvious to a person having ordinary skill in the art (“POSITA”) at the time of the alleged invention, in light of the prior art. It is my opinion that the Challenged Claims would have been obvious to a POSITA.

5. In the preparation of this declaration, I have studied:

- the 911 Patent (Ex.1001),
- the prosecution history of the 911 Patent (“911 File History”) (Ex.1002),
- Complaint
- Preliminary Infringement Contentions (Ex.1010)

- the additional references and publications I refer to below and/or identified in the exhibit list.

6. In forming the opinions expressed below, I have considered: the documents listed above; the relevant legal standards, including the standard for obviousness; and my own knowledge and experience based upon my work in the field of network communications and security as described below.

7. Unless otherwise noted, all **emphasis** in any quoted material has been added.

8. Since I was retained, I have been in frequent contact with counsel while working on this matter to communicate my opinions. My opinions are presented both in this Declaration and in the Petition for Inter Partes Review, which it supports. In many, if not most, cases the language used to express my opinions in both documents is substantially the same. This is because I worked with counsel to capture my opinions, and generally speaking we captured them in the Petition itself as opposed to developing both the Petition and Declaration in tandem. My opinions, as reflected in the Petition we prepared, were then incorporated into my Declaration. I make this statement to explain the work process and underscore that the opinions in this Declaration are my own, were formed by me and were documented with the assistance of counsel in the manner described. As detailed herein, it is my opinion that each of the challenged claims is rendered obvious by prior art references that

predate the earliest priority date of the 911 patent. If requested, I am prepared to testify about my opinions expressed herein.

II. QUALIFICATIONS AND PROFESSIONAL EXPERIENCE

9. I am over the age of 18 and am competent to write this declaration. I have personal knowledge, or have developed knowledge, of these technologies based upon education, training, or experience, of the matters set forth herein.

10. My complete qualifications and professional experience are described in my Curriculum Vitae, a copy of which can be found in Ex.1004. The following is a brief summary of my relevant qualifications and professional experience.

11. I hold a Bachelor of Science from the California State University at Hayward (now “California State University, East Bay”) in Mathematics and Computer Science, conferred in 1988. I received a Doctor of Philosophy in Computer Science from the University of California at Davis in 2000.

12. I am an Associate Professor of Computer Science at the University of Colorado, Boulder. I have taught more than 30 classes in Computer Science, on subjects including Data Structures, Algorithms, Networking, Operating Systems, Software Engineering, Security, Cryptography, Discrete Mathematics, and Quantum Computing. In my networking classes I have lectured on topics including network routing, path selection, network efficiency and reliability and Voice over IP (VoIP) technologies. I have authored or coauthored more than 20 publications. I have been

involved with computers for over 35 years in both commercial and academic capacities.

13. My father owned a company that sold computing services and I was the sole employee engaged in the installation, configuration, customization and maintenance of computer systems and networks from 1976 to 1988. This experience exposed me to some of the earliest computer systems and networks such as MS DOS, MS Windows, Novell Networks, etc., as well as to storage interfaces such as IDE and SCSI and to network protocols such as Ethernet, IP, and TCP. After my B.S. degree in 1988, I worked for six years for Ingres Corp as a software developer writing and reviewing code written in C (a programming language). Then, in 1995, I started working toward a PhD in Computer Science at the University of California at Davis. My area of focus was cryptography and security. I have long been interested in network protocols and their security. I have several publications directed to research in these areas and I have taught more than 20 classes that include core concepts, techniques, and protocols involving network security and their underlying cryptographic algorithms.

14. After my B.S. degree in 1988, I worked for 6 years for Ingres Corp as a software developer, working on the Data Management Facility, which involved writing software for storage systems, clusters, RAIDs and logging systems. Then, in 1995, I started working toward a PhD in Computer Science at the University of

California at Davis, graduating in 2000 and working as an Assistant Professor of Computer Science at the University of Nevada at Reno (until 2002), and then the University of Colorado at Boulder (where I am still currently employed).

15. In 2011, I became Chair of my department's Computing Committee and as part of that work, I built a new data center for use by my department. This involved the installation and oversight of racks, servers, fiber optic cables, switches, power, and cooling as well as setting policy for resource allocation including the use of power and network bandwidth by data center tenants.

16. For 2016 and 2017, I was on Entrepreneurial Leave from CU Boulder to start a company called "SecureSet" in Denver Colorado. The objective of SecureSet is to take reasonably proficient technical people and turn them into computer and network security specialists in 5 months of intensive training. I returned to my position at the University of Colorado in Fall of 2018.

17. I have contributed to more than 20 peer-reviewed publications, including both journals and conference proceedings, and various patent applications. I authored particular chapters of various books including a chapter for UNESCO's "Encyclopedia of Life Support Systems" and Springer's "Encyclopedia of Cryptography and Security."

18. In all, I have written widely for a broad range of computing and networking publications and computer and networking-oriented articles. As

previously noted, my *curriculum vitae* provides as comprehensive a record of these publications as I have been able to assemble.

19. I received the National Science Foundation's CAREER award for my research in computer and network security, as well as numerous awards for teaching.

III. LEVEL OF ORDINARY SKILL IN THE ART

20. My opinions are provided based on what a person of ordinary skill in the art ("POSITA") in the technical field of the invention would have understood at the time of the purported invention of the 911 patent. I have been asked to assume the date of the purported invention for the 911 patent is May 12, 2004.

21. I understand there are multiple factors relevant to determining the level of ordinary skill in the pertinent art, including (1) the levels of education and experience of persons working in the field at the time of the invention; (2) the sophistication of the technology; (3) the types of problems encountered in the field; and (4) the prior art solutions to those problems.

22. Based on the materials and information I have reviewed, and on my experience in the technical areas relevant to the Asserted Patents, a POSITA with respect to the Asserted Patents as of the time of the invention, i.e., the earliest possible priority date of the 911 Patent, May 12, 2004, would have had a bachelor's degree in electrical engineering, computer science, computer engineering, or an equivalent, and two years of relevant experience involving computer science

fundamentals, including pointers, data structures, object-oriented programming, database principles, internet-based communication technologies and electronic mail systems that is pertinent to the 911 Patent. Lack of professional experience can be remedied by additional education, and vice versa.

23. For purposes of this Declaration, in general, and unless otherwise noted, my statements and opinions, such as those regarding my own experience and what a POSITA would have understood or known generally (and specifically related to the references I consulted herein), reflect the knowledge that existed in the relevant field as of the priority date of the 911 Patent.

IV. RELEVANT LEGAL STANDARDS

24. I am not an attorney. In preparing and expressing my opinions and considering the subject matter of the 911 Patent, I am relying on certain basic legal principles that counsel has explained to me.

25. I understand that prior art to the 911 Patent includes patents and printed publications in the relevant art that predate the priority date of the 911 Patent. For purposes of this Declaration, I am applying May 12, 2004, as the priority date of the 911 Patent, but I take no position on that issue.

26. I understand and have been informed that in order for an inventor to be entitled to a patent, the invention must be “new.” There are a number of ways that an invention can be “anticipated,” that is, not new.

27. A patent claim is anticipated if each limitation of the patent claim is found either expressly or inherently in a single item of prior art. While a prior art reference need not use the same words as a patent claim to anticipate the claim, the prior art reference must describe the requirements of the claim with sufficient clarity such that POSITA would have been able to make and use the claimed invention based on the reference and his or her knowledge in the applicable technical field.

28. The disclosure of a feature of a claimed invention can be either “express,” meaning that the text, figures, or other content of a reference actually teaches the aspects of the patent claim, or the disclosure can be “inherent.” An express disclosure does not need to use the same words as the patent claim being considered, and does not need to be depicted in exactly the same way, but the disclosure must actually convey the claimed invention to POSITA using some combination of words, figures, data, or other portions of the patent. I have been informed and understand that in order to establish that an element of a claim is “inherent” in the disclosure in a prior art reference, it must be clear to a person skilled in the art that the missing element is an inevitable part of what is explicitly described in the reference, and that it would have been recognized as necessarily present by a person skilled in the art. Inherency cannot be established by probabilities or possibilities, and the mere fact that something may result from a given set of circumstances is not enough to establish inherency.

29. I further understand that to anticipate a claim, the elements of the claim must be found in the prior art reference “arranged as in the claim.” I have been informed that this means that the elements of the prior art reference may not be recombined in a way that is not taught in the prior art to arrive at the invention; instead, the prior art must teach the claimed subject matter in the arrangement required by the claims.

30. I understand that if a reference incorporates other documents by reference, the incorporating reference and the incorporated reference(s) should be treated as a single prior art reference for purposes of analyzing anticipation.

31. I understand that it is acceptable to consider evidence other than the information in a particular prior art document to determine if a feature is necessarily present in or inherently described by that reference.

32. I have been informed and understand that subject matter claimed in a patent is obvious under 35 U.S.C. § 103 if a person of ordinary skill in the art at the time the alleged invention was made would have had reason to combine or modify the disclosures of one or more prior art references to arrive at the claimed subject matter. A claim may be unpatentable even if each claim limitation is not present or disclosed in a single prior-art item.

33. I have been informed and understand that, under the doctrine of obviousness, a claim is unpatentable if the differences between the invention and the

prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which the subject matter pertains. A person of ordinary skill in the art is presumed to have knowledge of the relevant prior art at the time of the claimed invention.

34. I have been informed and understand that obviousness is based on the scope and content of the prior art, the differences between the prior art and the claim, the level of ordinary skill in the art and secondary indicia of obviousness and nonobviousness to the extent such indicia exist. I have been advised that the following evidence, when present, must be considered before forming a conclusion that a claimed invention is obvious: (1) the invention's commercial success, (2) long felt but unresolved needs, (3) the failure of others, (4) skepticism by experts, (5) praise by others, (6) teaching away by others, (7) recognition of a problem, and (8) copying of the invention by competitors. In this instance, however, I am not aware of any evidence of any of these factors. I will of course consider any evidence on these issues which is presented to me.

35. I have been informed and understand the determination of whether the asserted claims would have been obvious to a person of ordinary skill in the art and, therefore, invalid, is not governed by any rigid test or formula. A determination that a claim is obvious is, instead, based on a common-sense determination that the claimed invention is merely a combination of known limitations to achieve

predictable results. Any of the following rationales are acceptable justifications to conclude that a claim would have been obvious: (1) the claimed invention is simply a combination of known prior art methods to yield predictable results; (2) the claimed invention is a simple substitution of one known element for another to obtain predictable results; (3) the claimed invention uses known techniques to improve similar devices (methods, or products) in the same way; (4) the claimed invention applies a known technique to a known device (method, or product) ready for improvement to yield predictable results; (5) the claimed invention was “obvious to try”—choosing from a finite number of identified, predictable solutions, with a reasonable expectation of success; (6) there is known work in one field of endeavor that may prompt variations of it for use in either the same field or a different one based on design incentives or other market forces if the variations would have been predictable to one of ordinary skill in the art; or, (7) there is some teaching, suggestion, or motivation in the prior art that would have led one of ordinary skill in the art to modify the prior art reference to combine prior art teachings to arrive at the claimed inventions.

36. I have been informed and understand that an analysis of whether a claimed invention is obvious must not rely on a hindsight combination of prior art. The analysis must proceed in the context of the time of the invention or claimed priority date and consider whether the invention as a whole would have been obvious

to a person of ordinary skill in the art, taking into consideration any interrelated teachings of the prior art, the effects of demands known to the design community or present in the marketplace, and the background knowledge possessed by a person having ordinary skill in the art, all in order to determine whether there was an apparent reason to combine any known elements in the fashion claimed by the patent at issue.

V. OVERVIEW OF THE 911 PATENT

37. The 911 Patent was based on U.S. Patent Appl. No. 11/129,212 filed on May 12, 2005, and claims priority to Prov. App. 60/570,848, filed May 12, 2004, and 60/570,861, filed May 12, 2004. Ex.1001, 1:7-10.

A. The Disclosures of the 911 Patent

38. The 911 Patent is generally related to a system and method for enforcing compliance on electronic messages delivered via a network such as the Internet. *Id.*, Abstract, 1:18-21. The 911 explains there is a need for an electronic messaging system to facilitate separating business records from general e-mail “chatter” and enforces compliance policies on the messages. *Id.*, 1:56–59. The claimed invention “allows an administrator to define and enforce one or more compliance policies for the electronic messages.” *Id.*, 1:63–65.

39. The patent describes the claimed invention as follows:


...the messaging system utilizes message components that can be independently stored, accessed, and controlled. The messaging system utilizes a compliance policy that is defined in terms of a population structure, rule set, and jobcodes, and describes the set of rules that apply to message components during their lifecycles in the messaging system. An enterprise can thus bring its electronic messaging system into compliance with internal or government-mandated regulations.

Id., 15:34-42.

40. With reference to Fig. 2, the patent describes that the above-described functionalities are implemented as follows:

The messaging system treats a set of related messages, such as an email string between two or more people, as a message container (200) having relational references to one or more submessages (210, 212, 214). A messaging server (112) stores the messages and submessages as discrete message components within a message database (416). In addition, the messaging server (112) includes a governance policy module (422) having a governance policy database (424) that stores one or more compliance policies. A compliance policy describes the set of rules that apply to message components during their lifecycles in the messaging system.

Id., 1:65–2:9.




Ex.1001, Figs. 2 and 4.

41. A submessage is defined as the message that the end user actually composes: “When an end-user composes and sends a message, she is actually composing a submessage, and then sending a message 200 containing a reference to the submessage 200 to other end-users.” *Id.*, 5:10–13.

42. With reference to Fig. 1, the 911 Patent describes that the messaging server 112 can “act[] as a central repository for messages received by the end-users,”

“communicate with the messaging clients 116 and proxy servers 114 via the network 110,” and “provide[] interfaces that allow other entities in the messaging system, such as the proxy servers 114 and/or messaging clients 116 to exchange messages.”

Id., 3:43–53. The messaging client may be “a computer system executing standard messaging software,” or “other types of electronic devices, such as personal digital assistants (PDAs); cellular telephones with text messaging functionality, portable email devices.” *Id.*, 4:32-45. Messages include emails, Short Message Service (SMS) messages, Instant Messages (IMs), Multi-Media Message (MMS) and/or other types of messages, as well as media files, such as audio, video, or still images. *Id.*, 3:8–14.



Ex.1001, Fig. 1.

43. A crucial component within the messaging server 112 is the governance module 422 that is responsible for applying the governance policies for the system, which the patent describes as follows:

The governance module 422 (also called a “compliance module”) applies governance policies (also called “compliance policies”) to the messaging system and controls the governance policy database 424. In one embodiment, the governance module 422 also maintains a compliance log logging compliance-related events. The database 424 stores compliance policies that are established on the messaging system. A compliance policy describes the set of rules that apply to message components during their lifecycles in the messaging system. In one embodiment, each message component in the messaging system is associated with the compliance policy that was active at the time the content was created. Accordingly, when an action occurs that involves a piece of content, the messaging system identifies the relevant compliance policy in the governance policy database 424 and applies it. A compliance administrator establishes and administrates the compliance policies for the messaging system.

Id., 7:62–8:11.

44. The 911 Patent explains that “the governance module 422 specifies compliance policies in terms of populations, rules, and jobcodes,” and that:

“the populations represent the universe of end-users and/or other entities that use the messaging system. The rules represent the terms of the compliance policies. The jobcodes describe which rules are applicable in which circumstances, and can represent exceptions to the standard rules.”

Id., 8:12-19.

45. In the compliance policy population structure, an end-user is associated with three types of attributes. “The first type of end-user attribute specifies the domain and groups to which the end-user belongs,” where “each end-user is assigned

to exactly one domain.” *Id.*, 8:22-25. “For example, all of the end-users of a given domain can have the same minimum password length, password expiry period, and number of incorrect login attempts that are allowed before being locked out.” *Id.*, 8:28-32. An end-user can belong to multiple groups. *Id.*, 8:36-57. Examples of groups include geography, seniority, and project based. *Id.*

46. The second type of end-user attributes “is a set of allowable actions,” which “define the activities that the end-user can participate in using the messaging system, and can include ‘can use mail,’ ‘can print,’ ‘can set passwords,’ etc.” *Id.*, 8:52-57. The third type of end-user attributes “relate to the specific end-user... such as the end-user’s line of reporting, human resources advisor, time spent with the company, etc.” *Id.*, 8:64-9:2.

47. The 911 Patent further explains that “the compliance policy specifies a collection of rules that govern how content (i.e., message components and related data) is handled by the messaging system.” *Id.*, 9:26-28. Rules may be divided “into the following rule types: communication rules, readership rules, action rules, retention rules, and search rules.” *Id.*, 9:32-34.

48. “Communication rules determine which end-users, domains, and groups are allowed to communicate with other end-users, domains, and groups.” *Id.*, 9:37-39. For example: “If an end-user attempts to send a submessage to a recipient with whom he is not allowed to communicate, along with other recipients with whom

he is allowed to communicate, an embodiment of the messaging system indicates to the sender that some of the recipients will not receive that submessage and asks for confirmation.” *Id.*, 9:65-10:5.

49. “Readership rules determine which end-users, domains, and groups are allowed to read content written by other end-users, domains, and groups when such content is forwarded to them in some manner.” *Id.*, 10:10-13. “Action rules determine which end-users, domains, and groups are allowed to perform various actions on content written by other end-users, domains, and groups,” and “the action rules act as a form of digital rights management (DRM) on the message components in the messaging system.” *Id.*, 10:38-43. “Retention rules specify how long content authored by end-users, domains, and groups is retained by the messaging system before being deleted.” *Id.*, 11:7-10. If the messaging system includes functionalities that allows for message searching, “[s]earch rules determine which end-users, domains, and groups are allowed to view content written by other end-users, domains, and groups when searching.” *Id.*, 11:52-58. For example, “[i]f the search produces submessages or other content that the searcher is not allowed to view, one embodiment of the messaging system replaces the content with an error message.” *Id.*, 12:18-20.

50. To apply a rule, the 911 Patent explains: “when the governance policy module 422 determines what communication rule to apply, it first considers the

highest ranked rule and, if the communication in question is not governed by this rule, it proceeds to the next highest ranked rule until it discovers the first rule that applies. The lowest ranked rule represents the default and applies to all communications.” *Id.*, 9:59-65.

51. Jobcodes provide exceptions for the standard rule set. *Id.*, 12:35–36. Jobcodes also “allow submessages and submessages to be tied to external business processes.” For example, “jobcodes allow messages to be associated with a business deal such as a merger, with a patient at a hospital, with particular clients of a law firm, etc.” *Id.*, 12:36–40.

52. To use the system, “the governance policy module 422 generates a graphical user interface (GUI) on a messaging client 116 and/or other terminal that an administrator can use to establish, edit, and view compliance policies and related information,” allowing the administrator to “select population structure entities and rules, specify available jobcodes, and set various options using GUI elements.” *Id.*, 14:2-9. For example, the messaging client 116 includes the components shown below:




FIG. 5

Id., Fig. 5.

53. The client module 510 is “an application dedicated to sending and receiving messages” that “includes standard functionality for composing messages, viewing messages, replying to and forwarding messages, etc.” *Id.*, 14:15-19. The client module may also include “an audit and security cache 514 for caching audit and/or security information received by the client module,” to, for example, use the “digitally-signed tamper-detection data” or “to access secured submessages stored in the message cache.” *Id.*, 14:36-43.

54. Fig. 6, shown below, provides a flow chart of the steps of the invention.




FIG. 6

Id., Fig. 6. The 911 Patent explains that: “Initially, an administrator establishes 610 a compliance policy for the messaging system,” using “the GUI to specify the population structure, rule set, and jobcodes that define the compliance policy,” and that “[t]he compliance policy is stored in the governance policy database.” *Id.*, 14:60-67.

55. Then, an end-user performs an action, such as creating a new submessage, or “attempting to view, print, copy, reply to, forward, and/or delete a submessage.” *Id.*, 15:1–5. “The submessage is received by the messaging server 112 and stored in the message database 416.” *Id.*, 15:9-10.

56. The 911 Patent describes the application of the compliance policy to the message component as follows:

In one embodiment, the messaging server 112 applies 614 the compliance policy to the message component each time an entity performs (or attempts to perform) an action on it. As part of this process, the messaging server 112

determines the jobcode assigned to the message component and then determines the rules applicable to it. The messaging server 112 applies the rules, and thereby allows or denies the action as may be appropriate.

Id., 15:11-18.

57. Further, “the messaging server 112 periodically scans 616 stored message components and/or other stored data for compliance with the validity and retention rules specified by the compliance policy.” *Id.*, 15:23-25. Doing so, “the messaging server 112 may apply the validity rules to identify message components, end-users, groups, rules, compliance policies etc. as invalid, passive, and/or orphaned,” or “may apply the retention rules to identify message components or other stored data having expired retention periods.” *Id.*, 15:26-33.

B. The Prosecution History of the 911 Patent

58. I have reviewed the file history of the 911 Patent, including that of the aforementioned U.S. Patent applications, to which the 911 Patent claims priority.

59. On May 12, 2005, the Applicant filed the application no. 11/129,212, which matured into the 911 patent, with 30 claims. On June 30, 2009, the U.S. Patent and Trademark Office (PTO) issued a non-final office action that rejected all 30 pending claims as anticipated by U.S. Pub. no. 2004/0177271 to Arnold et al. (“Arnold”). The Applicant amended the claims to add, among other things, the limitations for applying rules to submessages and argued that Arnold does not teach this limitation. Ex.1002, 68, 78 (9/30/09 Amendment). On January 19, 2010, the

PTO issued a final office action that rejected all pending claims based on Arnold in view of U.S. Pub no. 2004/0039753 to Albaugh et al. (“Albaugh”). The Applicant further amended the independent claims to recite “message container” that references submessages, as well as argued that the claims as amended are distinguishable from Arnold, Albaugh, and other references the Examiner identified. Ex.1002, 176, 183–184 (3/18/2010 Amendment).

60. On July 13, 2010, the PTO rejected the pending claims based on an additional reference, U.S. Patent No. 6,138,168 to Kelly (Kelly). Ex.1002, 196 (7/13/10 Rejection at 3). In response, the Applicant amended the independent claims to recite that the of plurality of submessage are stored externally to the message container. Ex.1002, 232, 235 (11/15/2010 Amendment). On February 16, 2011, the PTO issued a final rejection based on Arnold in view of a new prior art reference, U.S. Patent No. 6,725,228 to Clark and in view of Nassiri (U.S. Pub. No. 2006/0265464). Ex.1002, 264. The Applicant again narrowed the claims such that the “messaging module” and “governance module” control only sent messages. On September 28, 2011, the PTO issued a notice of allowance of all pending claims. Ex.1002, 323.

VI. CLAIM CONSTRUCTION

61. It is my understanding that in order to properly evaluate Challenged Claims, the terms of the claims must first be interpreted. It is my understanding that

for the purposes of this *inter partes* review, the claims are to be construed under the so-called *Phillips* standard, under which claim terms are given their ordinary and customary meaning as would have been understood by a POSITA in light of the specification and prosecution history, unless the inventor has set forth a special meaning for a term. I have also been informed that claim terms only need to be construed to the extent necessary to resolve the unpatentability inquiry.

VII. IDENTIFICATION OF HOW THE CLAIMS ARE UNPATENTABLE

62. The discussion in this Declaration provides a detailed analysis of how the asserted prior art teaches each limitation of the Challenged Claims.

63. As part of my analysis, I have considered, and discuss in detail, the scope and content of the prior art and any differences between the alleged invention and the prior art.

64. It is my opinion that the alleged invention recited in the Challenged Claims would have been obvious in view of the teachings of the asserted prior art and the knowledge of a POSITA, as I discuss in detail below.

VIII. TECHNOLOGY BACKGROUND

65. As I explain next, the technology disclosed and claimed in the 911 Patent was well-known to a POSITA by the earliest possible priority date of the 911 Patent, May 12, 2004.

A. Electronic mail systems

66. Electronic mail systems, or email systems, were well-established and widely-used by the time of the 911 Patent's priority date. The Background of the Invention section of the 911 Patent makes this clear when referring to email system as containing both informal and formal content, before describing the addition of filtering systems existing email systems. Ex. 1001, 1:40-55. However, the 911 Patent claims there were "few, if any, ways to automate the process of filtering and storing business-related e-mails." Ex. 1001, 1:36-38. The patent goes on to describe a further problem when email systems store the same message multiple times, thereby resulting in multiple copies of the same message being stored which is "inefficient and undesirable." Ex. 1001, 1:39-45. The 911 Patent proceeds to then offer solutions to these problems via its purported inventions, as described above when I summarized the patent's disclosure. However, as explained in detail below, nothing in the 911 Patent was new by May 12, 2004.

B. Message Databases

67. The concept of storing email messages in a database was well-established and would have been known to a POSITA by the priority date. General-purpose databases for holding information of any kind, including messages, existed by the 1980s, and I worked for a major database company from 1988-1994 where many of our customers stored messages in our database product. Further, the prior art application cited by the Examiner during prosecution of the patent, i.e., Arnold,

specifically references “copies of the communication are preserved in a database.” Arnold, [0023].

68. Email systems at the time typically stored messages using the computer’s underlying file system. For example, when I was an undergraduate student in the mid-1980s, the most prevalent MTA (Mail Transfer Agent) was “sendmail.”¹ Sendmail would store messages in a container for each recipient labeled by username; for example, my incoming mail in 1984 would be found in /var/spool/mail/jrblack of my school’s mail server. I used a program simply called “mail” to read emails from the incoming mail spool, and I would either delete them after reading, or save them if I needed to refer to them later. In order to do the latter, I would specify a container name where each message should be stored; for example I might have containers named “games,” “homework,” “hiking,” and so forth. Again, these containers were simply implemented as files within the client computer’s file system.

C. Pointers and References

69. The concept of a pointer (or a reference) is an idea that long predates computer science. For example, when I was in high school I learned how to cite to

¹ See <https://en.wikipedia.org/wiki/Sendmail>

references in my essays or scientific papers using the proper format; this might include the name of the reference, author, date, publisher, page numbers, etc.

70. In Computer Science, we often use the concept of a pointer or reference to some other object or resource that resides on the computer (or on a remote computer). The usual requirements for a pointer are that it (1) it is unambiguous and (2) can be readily used to locate the referenced object. For example, a memory location is a clear and easily-used type of pointer: memory locations point to a single specific place in the computer's memory, and computers are optimized to quickly read or write to or from a memory location. A vague reference such as "picture of a dog" is not a "pointer" because it is ambiguous (for instance, there could be many pictures of dog on my computer, and in some pictures it may be difficult to determine whether there is a dog shown) and searching for a "picture of a dog" through the entirety of the computer's memory would potentially take a long time.

71. By the priority date of the 911 Patent, pointers had been used for decades. In the C programming language, one of the first concepts we teach undergraduate students is how to create, read, and write pointers, which are simply variables containing memory locations of other variables. For example, in C, I might write

```
int counter;  
  
int *pointer = &counter;
```

72. This would create an integer variable called “counter” and a pointer variable called “pointer” that would contain the memory address of (i.e., a pointer to) the “counter” variable.

73. Virtually every computer language has pointers or references to objects as a basic concept. In Java, a language invented by James Gosling while at Sun Microsystems in 1995, the programmer creates “objects” and then can make “references” to those objects when passing these objects around to different functions (functions are called “methods” in Java).


74. When I taught Java in the early 2000s, I used a book called “Head First Java: Your Brain on Java – A Learner’s Guide” (Ex.1014). Chapter 3 of that book is entitled “Primitives and References.”

Variables come in two flavors: primitive and reference. So far you’ve used variables in two places—as object **state** (instance variables), and as **local** variables (variables declared within a *method*). Later, we’ll use variables as **arguments** (values sent to a method by the calling code), and as **return types** (values sent back to the caller of the method).

Ex.1014, 47.

75. A “primitive” is a simple variable like an integer or real number. But an object is a more versatile and complex type of storage container created using a “class.” Before an object can be created, the programmer has to first declare the class that describes what the object will contain and how it is structured. For

example, the program snippet below shows an example class called “Dog” that contains an integer named “size,” strings named “breed” and “name,” and a method called “bark.”



Ex.1014, 34.

76. Once a class is declared, one or more objects can be created using this class. Creating an object causes computer memory to be allocated to hold/store the contents of that object. To keep track of the object in a program, e.g., to read from and/or write into the contents of an object, a “reference” is used. For example, the below snippet shows how the Dog class could be instantiated, i.e., how an object can be created from the Dog class:

The Dot Operator (.)

The dot operator (.) gives you access to an object's state and behavior (instance variables and methods).

// make a new object

Dog d = new Dog();

// tell it to bark by using the
// dot operator on the
// variable *d* to call bark()

d.bark();

// set its size using the
// dot operator

d.size = 40;

Ex.1014, 34.

77. The code above does the following: a new object is created from the Dog class, and then a reference called “d” is created and made to point to the newly-created object. This means that somewhere in the computer’s memory enough space is allocated to hold the three variables defined in the class (size, breed, and name) in a single contiguous block, and the address of that memory location is stored in the reference called “d.” After this, the reference can be used to execute a method


associated with the Dog object (`d.bark()`) and a new value of size 40 can be assigned to the corresponding object's variable (`d.size = 40`). Using the value or a pointer or a reference to find the referenced object is called "dereferencing."

78. The above examples show how memory locations are commonly used as pointers in computer languages, but references can be applied for other attributes as well. For example, a file on my computer is referenced by its path and filename. This declaration might be stored in `/home/users/john/consulting/sterling/reports/declaration.docx` for example. That path and filename uniquely describes a file on my computer in a manner that is both unambiguous and allows efficient access to the file. Another familiar type of reference is a URL used with the World Wide Web. A URL gives a reference to an object that is found over the Internet, such as [https://en.wikipedia.org/wiki/Reference_\(computer_science\)](https://en.wikipedia.org/wiki/Reference_(computer_science)).

D. De-duplication

79. The general idea that it is more space-efficient to store an object only once, and then use multiple references to that single object, is well-established in computer science. For example, if someone has a digital file of phone numbers and he/she wants that file to exist in two different folders on his/her computer, there are two options: (1) store a copy of the file in each folder, or (2) store the file once and use pointers or references within each folder to the phone number file. The ability

to perform both option (1) and option (2) have been around since the advent of Unix around the 1970s. In Unix or Unix-like systems (such as Solaris, AIX, Linux, etc.) the user can create a file and then create “links” to the file, making it appear that there are multiple copies of the file when in fact all links are just pointers to the sole copy. Links can be created using the “ln” command. So if I create a file called “Link A.TXT” using a text editor and then type “ln LINK A.TXT LINK B.TXT,” it would result in two links to the same file:




https://en.wikipedia.org/wiki/Hard_link

80. The image above shows two hard links, named “LINK A.TXT” and “LINK B.TXT”, pointing to the same physical data. Thus, for example, when the editor (or any other app) opens “LINK B.TXT”, it can see those changes made to “LINK A.TXT.” Any number of hard links to the physical data may be created. To access the data, a user only needs to specify the name of any existing link; the operating system will resolve the location of the actual data.

81. Another well-known method that uses this same idea of using references to a common set of information to save space is the Lempel-Ziv compression algorithm invented in 1978 (called the “LZ78” algorithm). LZ78 compresses a file by finding common repeated portions and storing them in a dictionary with the repetitions in the file being replaced by pointers. *See, e.g.*, https://en.wikipedia.org/wiki/LZ77_and_LZ78 discussing the Lempel-Ziv compression algorithm.

82. Removing duplicates in a message store was also known in the art by the priority date of the 911 Patent. For example, U.S. Patent No. 6,978,274 to Gallivan et al. (Ex.1018), filed Aug. 31 2001, directed at a document processing algorithm describes removing duplicates of documents (such as unstructured data, including electronic message stores, such as electronic mail (email) folders, word processing documents or Hypertext documents). *See* Ex.1018, 4:54-58.



Ex.1018, Fig. 14.

83. In addition, detecting and removing duplicate emails to save space was also well-known. For example, a research report from IBM, entitled “Duplicate Management for Reference Data” published in 2003 , describes methods of detecting and eliminating duplicate data, *e.g.*, emails stored in the Lotus Notes format. Ex.1009, 8 (“The emails are stored.”); Table 1 (“Industry Research Email - Lotus Notes format”); 5 (“[W]e present three methods for detecting duplicate data in a set of objects: simple blocking, content chunking, and sliding blocking.”); 8 (“We used several real-world data sets to evaluate the different methods for duplicate detection.”).

When a block of data is written to the object system, the system first determines how many copies of that data already exist logically in the file system. Based on this occurrence count, and the performance and reliability settings, the system then decides how many copies of that data it should maintain. If it currently has fewer than the desired number of copies of the data, the system proceeds to create a new copy of the data. If the system already has the desired number of copies of the data, it selects one of the existing copies of the data.

Ex.1009, 4. Thus, the system disclosed in the IBM paper provides greater in allowing IT system administrators to decide on “the desired number of copies of the data.” Thus, for example, if a system admin chooses to have just one copy of the data (*e.g.*, emails), then only one copy of the data is stored and additional copies are eliminated.

84. As another example, in a 2003 paper entitled “Filtering Spam Using Search Engines,” Kolesnikov described compiling a list of URLs extracted from emails into categories that were used to determine whether the email was spam or not. As part of this compilation of URLs, they removed duplicates: “For uncategorized URLs in the user’s mailbox, we remove duplicates.” Ex.1019, 3.

IX. THE PRINCIPAL PRIOR ART

A. Rohall

85. U.S. Patent Publication No. 2003/0163537 (“Rohall,” Ex.1005) was based on application number 10/334,087, filed on December 30, 2002, and published on August 28, 2003. Rohall, titled “Method and apparatus for handling conversation threads and message groupings as a single entity,” is generally directed to “data processing systems and, more specifically, to a technique for effectively reviewing and processing electronic mail and electronic mail threads.” Ex.1005, [0005].

86. As background, Rohall explains that given the increase in the volume and complexity of emails, there is a need to better organize and access complete mail threads, as well as allowing the generated electronic mail conversation threads to be stored in memory and manipulated with conventional operations. *Id.*, [0006]–[0011].

87. Rohall describes “an improved inbox or viewer for electronic mail that allows for greater integration of functions to enhance usability and productivity.” *Id.*, [0012]. Rohall discloses an electronic mail inbox that is based on the principles of: “1) bringing all communications together into one place; 2) helping to focus on what’s important; 3) finding the information and people needed for a task or activity; and 4) keeping things moving forward over time.” *Id.* The conversation thread, “may include a parent/child hierarchy of a plurality of original electronic mail documents,

summaries of electronic mail documents, shadow documents of original electronic mail documents, or any combination thereof.” *Id.*, [0013].


88. In one exemplary embodiment, Rohall discloses a networked system that performs: “(A) detecting the occurrence of an triggering event related to an electronic mail document; (B) processing the electronic mail document and storing the results therefrom in a data structure; (C) identifying an ancestor document of the electronic mail document, if any; (D) processing the ancestor document and storing the results therefrom in a data structure; (E) repeating (B) through (D) until all ancestor documents of the electronic mail document have been processed and the data structure represents all ancestor documents in a conversation thread.” *Id.*, [0014].

89. Rohall teaches a computer system, such as a Dell Dimension 8200 that were commercially available. *Id.*, [0040]. As shown below, Rohall’s computer system 100 include a central processing unit (CPU), memory, memory controller, and input and output devices, such as a keyboard, mouse, and video display. *Id.*, [0041] – [0042]. Rohall also contemplates using a tablet. *Id.* Rohall’s computer is generally controlled by software that were commercially available, such as Windows, AIX, UNIX, Linux, and Lotus Notes. *Id.*, [0043]

90. Rohall teaches using the databases and replication thereof based on the Lotus Notes structure. *Id.*, [0048]. For example, “[t]he Notes Storage Facility (NSF)

architecture defines the manner in which documents and databases are created, modified and replicated among Notes servers across a computer network.” *Id.*, [0050].

91. Fig. 3 of Rohall shows an example of its networked environment.



Id., Fig. 3. “The Notes client processes 312, 314, 316 and 318, which include mail functionality, may likewise be implemented as part of an all software application that run on a computer system.” *Id.*, [0052]. The “servers 302-310 and Notes client process 314 may include in memory a copy of database 350 which contains document 360.” *Id.* Rohall states that “document” includes “a document, database, electronic mail message code, a “Note” or any file which is accessible and storable by a computer system.” *Id.*, [0050]

92. Rohall teaches “a Mail Agent 230 [that] creates a shadow document from an original document by generating a file containing data related to the

document,” where the shadow documents are stored in a database and are accessible. *Id.*, [0056]. A “new electronic mail message is considered a parent document and serves as the root from which a new shadow tree may be derived,” and “[a]ny replies to the original electronic mail message is/are considered a child/children document(s).” *Id.* Each document comprises an ID and optional metadata fields. *Id.*, [0057]. Rohall further explains that “the meta data may include such logistical information as sender, receiver, original size, subject, date, any carbon copy recipients, etc. associated with the document,” and that “key words or summaries of the content of the document or any attachments may likewise be included.” *Id.*, [0058].

93. The Parser 232, a component of the Mail Agent 230, “sets up a shadow document and parses the original document and any header of the following data fields: sender, receiver, original size, subject, date, any carbon copy receivers, attachment names, etc. and makes call to filtering software modules, as necessary.” *Id.*, [0059].

94. Fig. 5A of Rohall shows the process flow.




Figure 5A

Id., Fig. 5A.

95. Rohall describes the process flow shown in Fig. 5A as follows:

As illustrated in FIG. 5A, Mail Agent 230 first detects the occurrence of a triggering event as illustrated by decisional step 500. Such event may include the sending or receipt of an electronic message, or, alternatively a request to delete an electronic message. Next, Mail Agent 230 determines if the electronic message is a new message, as illustrated by decisional step 502. Within a Lotus Notes electronic mail domain, it is possible to determine with existing object methods to which message an incoming message is a reply. If the incoming message is from the Internet or from a non-Notes environment, a conventional algorithm can be used to determine if the message is a new message or a reply to an existing message by determining if the message has an “In-Reply-To” header, or whether the subject lines of the message match an existing message or shadow document. If so, Root Document ID 406 and Parent Document ID 404 are both set to null, as illustrated by procedural step 504.

Otherwise, Mail Agent 230 sets the Parent Document ID 404 to a pointer value referencing the parent document and simultaneously modifies one of the Child Document IDs 408 a-n of the parent document to reference the subject shadow document, as illustrated by procedural step 506. Additionally, Mail Agent 230 sets Root Document ID 406 to reference the root of the conversation thread tree, as illustrated by procedural step 508. Mail Agent 230 then sets the Original Document ID 402 to reference the original document from which the shadow document was created, as illustrated by procedural step 510. If the original document has been deleted, the value of Original Document ID 402 is set to null. Finally, Parser 232 parses the header information of the original electronic message for meta data and populates Meta Data fields 410 a-n accordingly, as illustrated by procedural step 512. Parser 232 may optionally make procedure calls for scanning of the document content or any of its attachment for key words or phrases to be saved as meta data. Thereafter, the shadow document is stored in memory, which, in the illustrative embodiment, is a mail database, as illustrated by procedural step 514.

Id., [0060]

96. Rohall also teaches how to construct a conversation thread tree based on examining the parent and child documents. *Id.*, [0064]. The system can also display graphical representation of the conversation thread trees, including highlighting the relationship among the messages. *Id.*, [0066]. The message trees can also be displayed on a timeline, showing when a message was received, and can also include information from a calendar to reflect certain milestone, such as when the

user was on vacation. *Id.*, [0069]. Rohall's system also provides a user interface. *Id.*, [0071].

97. The content of the messages may also be summarized using the summarization algorithm, whereby the resulting data may be stored. *Id.*, [0072] – [0080]. Rohall also “recognizes that there are specific domains in which identifying features, such as names, dates, and company names, product names, becomes useful,” and that “commercially-available feature extraction software extracts relevant features in documents, including names, numbers, and names of organizations and products.” *Id.*, [0084].

98. Rohall also teaches an inbox that can “receive not only electronic messages, but data and documents from other sources such as databases, templates and other information sources.” *Id.*, [0096]. “Using the shadow document generation techniques described herein, a surrogate document including meta data such as size, date, heading information and a pointer to the pointer actual data, is generated by Notes Mail Agent 230 and placed in inbox 900.” *Id.*

99. As an additional feature, Rohall explains that its system can monitor data specific to users and databases:

Notes mail agent **230** may be provided with the names of selected individuals and the addresses of the databases or other inboxes necessary to monitor such external information. Mail agent **230** upon receiving data associated with a particular user generates a shadow or

circuit [sic:surrogate] document, in a manner as described herein and transmits the surrogate document to inbox **902**. In this manner, inbox **900** becomes the central location for receiving not only electronic mail but other sources of information useful to a user. Selection of the surrogate document **916** from inbox **900** causes the pointer data to be resolved and the actual data retrieved and displayed as item **918** within the inbox **900**, as illustrated in FIG. 10.

Id., [0097].

100. Rohall teaches that a user can create rules to discriminate the different types of communications. For example, “the user may designate specific topics by name and have a folder or Notes mail database created into which the document may be stored until accessed.” *Id.*, [0100]. Then, the “Mail agent 230 processes each received item and, if related to one of the designated topics, stores either the complete document or a surrogate document, complete with a pointer to the document location, within the folder or mail database,” and “[t]hereafter, the UI object for inbox 900 is notified to highlight the relevant button.” *Id.*

101. As another example, Rohall also “provides a special functionality that allows the user to view and take action on classes of similar documents or correspondence which are of a different priority and which the user does not wish to have mingled, and possibly, ignored in the main list of electronic mail.” *Id.*, [0103]. An example of this is an approval request. *Id.*

B. Schiavone

102. U.S. Patent Publication No. 2002/0120600 to Schiavone et al. (“Schiavone”), was based on application number 09/793,263, filed on February 26, 2001, and published on August 29, 2002. Ex.1006, Cover. Schiavone, titled “System and method for rule-based processing of electronic mail messages,” relates to “a method and apparatus for rule-based processing of electronic mail messages, e.g. to negotiate transactions, enable automated verification of compliance with privacy, legal, regulatory and/or best practices standards, to perform spam filtering, anonymous profiling, target marketing through dynamic offers, and prioritization and standardization of e-mail to transactions.” *Id.*, [0001].

103. While Schiavone was published containing five drawings numbered Figs. 1-5 (see Schiavone, pp. 2-6), Schiavone includes description in support of six drawing figures—including specifically disclosures pertaining to Fig. 6 that is missing from the Schiavone publication. See Schiavone, [0014] and [0061]-[0066]. Indeed, the file history of Schiavone (“Schiavone-FH” or Ex.1007) and its document history (“Schiavone-DH” or Ex.1008) both indicate that the as-filed, original application which published as Schiavone contained six drawing figures. Specifically, the records on Patent Center indicate that Applicant filed the original application filed on February 26, 2001 containing six drawing figures—including Fig. 6—and subsequently a record entered on June 6, 2001 indicates five drawing

figures. See Schiavone-FH, 40-45; Schiavone-DH, 1. Therefore, for purposes of my declaration and the petition it supports, I have considered Fig. 6 to inform my analysis.

104. Schiavone generally describes its functionalities as follows:

An electronic mail message's mail type specifier is identified. A data store stores data specific to the intended recipient. The data store is referenced and the electronic mail message is processed as a function of the mail type specifier and data retrieved from the data store. The mail may be processed at the recipient's mail client device and the information maintained confidential. The processing may be based upon reference to a rule applicable to that particular message. The rule includes a conditional instruction conditioned upon a result of a logic or other function of recipient profile data....

Id., Abstract.

105. As background, Schiavone explains that unsolicited emails often cause annoyance, increase costs, and dilute the email's value as a communication medium.

Id., [0002]. And although certain filter measures were available, they were of limited protection. *Id.*, [0005]. For example, "mail reader software programs such as Microsoft® Outlook® allow the user to define filters, for refusing receipt of mail based on keywords and/or a sender's identity." *Id.* Describing the problem and the need for a solution, Schiavone notes that "[p]resently, e-mail is used as a communications tool but few standards, rules, or laws apply to govern how the e-mail communications medium is used," and that "even with respect to laws,


regulations and best practices adopted and enacted, there are no satisfactory techniques for ensuring compliance of e-mail messages with applicable standards and/or laws and/or otherwise determining appropriateness of a message for an intended recipient.” *Id.*, [0006].

106. Schiavone “provides a system and method for rule-based processing of electronic mail messages.” *Id.*, [0007]. “The rule may be associated with messages by default, in an automated fashion, or as the result of a sender’s specification, e.g., by including a mail-type specifier in the message.” *Id.* The system also provides, among other things, automated “compliance checking and/or other actions with respect to electronic mail messages based on a shared rule set.” *Id.*, [0015].

107. Schiavone states that “rules may be implemented for a wide variety of applications, including compliance with applicable laws.” *Id.*, [0016]. For example, “a message may be rejected, and not delivered to the intended recipient, if a recipient’s profile data is incompatible with applicable laws for receiving the message, as specified in a rule associated with the message.” *Id.* In another example, “the rule may require modification of the content of the message as a function of such recipient profile data,” where the message may be modified before it is sent, or before it is received, e.g., by adding or deleting certain content, by modifying the message’s color, length, etc., or by assigning values to variables within the mail message.” *Id.*, [0017].

108. Schiavone teaches that the “recipient’s communications device may be any communications device capable of receiving electronic mail or instant messages, such as a general purpose computer running mail client software, a Web-enabled wireless telephone, a wireless personal digital assistant such as a Palm VII® device, a pager, etc.” *Id.*, [0020]. The user profile data is stored in a database on the communication device or in a recipient data store. *Id.*, [0021], Fig. 1.

109. Fig. 2, shown below, is a flow diagram of an exemplary rule-based processing.



Id., Fig. 2.

110. A “rule may be constructed to determine compliance with any applicable law, rule, regulation, code, best practice, etc. of any region, e.g., local, state, federal, foreign, etc. (‘regional law’).” *Id.*, [0029]. A sender could “select several different rule sets to automate their compliance with selected US federal and local laws and regulations.” *Id.* The rule controls how the e-mail message will be processed. *Id.*, [0030]. The rule may be contained within the message or may be obtained over a network from a central source, where it is stored in memory. *Id.*

111. In one example, “the message type specifier identified in step 26 is ‘COPPA’” indicating “that the sender intends the message to be delivered in accordance with The Children’s Online Privacy Protection Act, a law applicable in the United States (“US”) which forbids delivery of commercial electronic mail communications to children aged thirteen (13) and under unless a parent has given prior consent.” *Id.*, [0029]. When the recipient’s profile data is referenced, the profile data “indicates that the recipient lives in the US and is 45 years of age, and therefore, “it is determined that the recipient’s profile data is in compliance, i.e., compatible, with the applicable rule in step 32 and the electronic mail message is delivered to the recipient, e.g. placed in the recipient’s electronic ‘inbox’ and/or displayed to the recipient.” *Id.*, [0033]. If the recipient profile data is not in compliance with the applicable rule, the electronic message is rejected and may be deleted. *Id.*, [0034].

112. Schiavone explains that its system can also be used for workflow processing transactions, “e.g., by providing a rule requiring approval and/or handling of a given electronic mail message by more than one recipient and/or according to a predetermined deadline schedule.” *Id.*, [0035]. Rules may also be directed to individuals or groups of individuals. *Id.*, [0036]–[0039]. Rules may be used for security, “e.g., by implementing a rule requiring notary-like certification and authentication by a third party or process prior to sending of an e-mail message,” or to “allow[] the third party intermediary to ensure compliance with the system, e.g.,

by implementing rules requiring a digital signature to communicate within the system, and requiring such digital signature to be based on a digital certificate revocable by the trusted third-party intermediary upon abuse of the system, etc.” *Id.*, [0040], [0044].

113. Schiavone teaches using its system for filtering purposes. For example, “a recipient may include in the recipient profile data that for a period of two months, e.g. after purchasing a house, the recipient wishes to receive sale offers relating to purchases of new homes, or home furnishings, etc.,” and the rule “can be used to reject or discard the electronic mail message, e.g. before viewing by a recipient, if the appropriate information is not contained in the recipient’s profile data.” *Id.*, [0067].

114. With reference to Fig. 6 (reproduced below), Schiavone discloses

[A] recipient compliance engine 154 [that] is configured to identify a rule associated with the message, including referencing a rule base 155 as necessary, and to process the rule, e.g., by referring to public preference data 120 a, private profile data 120 b and/or private commerce data 120 c stored in the recipient client’s data store 120. [T]he recipient compliance engine may reference private profile data and modify the message in accordance with the rule as a function of the recipient’s profile data. The recipient compliance engine 154 may store transaction data in a transaction data data [sic] store 158, which may be reference by a third party such as trusted authority 200’s accounting engine 220 for monitoring such transactions, e.g., for advertising revenue purposes.

Schiavone, Fig. 6.



C. Gomes

115. U.S. Patent No. 8,375,008 to Gomes (“Gomes”) was filed on January 16, 2004 and claims earliest priority to Prov. App. Nos. 60/440,855 entitled “System and Method for Data Extraction in a Non-Native Environment, Data De-Duplication, Database Creation and Manipulation, Image Back-up and PST File Monitoring” (filed Jan. 17, 2003) and 60/440,728 entitled “Method and System for Enterprise-Wide Retention of Digital or Electronic Data” (filed Jan. 17, 2003). Gomes is directed at managing data to retain (or store or process) electronic data within an enterprise. Ex.1016, 1:65-67. Gomes teaches processing email files:

by an email extraction engine with the extracted data fed to a collective database and simultaneously de-duplicated, including **de-duplication of attachments, message bodies and metadata**. The other electronic data files (from the file extraction engine) and data from backup tapes or a network (or other sources) can be processed by a de-duplication engine to remove duplicates of the content while keeping metadata from each unique location of the copy(s) of the content. Also, **file content can be de-duplicated such that only one copy of the content is in the database regardless of the location of the file (e.g., an email attachment, a loose file, a re-named file with no change to the content)**.

Gomes, 2:4–13.

X. CLAIMS 1-13, 15, 18, 20-21, AND 27 ARE UNPATENTABLE

Ground 1	Claims 1-8, 10-13, 15, 18, and 20-21	Rohall and Schiavone
Ground 2	Claims 9 and 27	Rohall, Schiavone, and Gomes

A. Ground 1: The Combination of Rohall and Schiavone renders obvious claims 1-8, 10-13, 15, 18, and 20-21

B. The Combination of Rohall and Schiavone

116. A POSITA would have been motivated to combine Rohall with Schiavone. In particular, it would have been obvious to include Schiavone’s teachings related to “rule-based processing of electronic mail messages,” e.g., including identifying rules, referencing a rule base storing a compliance policy, processing rules (including identifying rule elements such as senders and recipients), determining that certain rules are applicable selectively and applying its teachings (*see e.g.*, Schiavone, Abstract, [0001], [0007], [0010], [0015]-[0017], [0020]-[0021], [0028]-[0030], [0031]-[0035], [0064], [0066], and Figures 1 and 2) to messages (and submessages) sent using Rohall’s “electronic messaging system” for the reasons explained below.

117. Both Rohall and Schiavone are analogous art as the 911 Patent for at least two reasons. *First*, both references are directed toward the same field of endeavor of the 911 Patent—electronic mail messaging systems. Ex.1001, Abstract;

Rohall, Abstract; Schiavone, Abstract. *Second*, both Rohall and Schiavone are reasonably pertinent to a problem described by the 911 Patent: automating the process of filtering e-mails, which is a problem of conventional electronic mail systems identified in the 911 Patent. Ex.1001, 1:36-40. Rohall describes its invention as “an improved inbox or viewer for electronic mail that allows for greater integration of functions to enhance usability and productivity.” Rohall, [0012]. Rohall explains that by making calls to filtering software modules, its inbox helps to focus on which messages are important. Rohall, [0059], [0012]. Schiavone, which is directed at “rule-based processing of electronic mail messages” expressly describes that its “invention can be used for filtering.” Schiavone, [0001], [0007].

118. Rohall describes the problem associated with email usages, stating, for example, that “people increasingly feel frustrated by their electronic mail,” and they “are overwhelmed by the volume, lose important items, and feel pressure to respond quickly.” Rohall, [0006]. Rohall recognizes that email systems have not materially changed since they were first invented. *Id.* Schiavone solves these problems because it discloses a “rule-based processing of electronic mail messages” that can “perform spam filtering,” “prioritization and standardization of e-mail to transactions,” as well as “enable[ing] automated verification of compliance with privacy, legal, regulatory and/or best practices standards.” Schiavone, [0001]. As such, a POSITA would have

understood Schiavone's teachings to expand or otherwise improve the capabilities of Rohall.

119. Schiavone also identifies the shortcomings of conventional electronic mail systems stating that “few standards, rules, or laws apply to govern how the e-mail communications medium is used. Yet increasingly, a patchwork of local, federal and international laws, regulations and best practices are being enacted and adopted, many attempting to legislatively eliminate spam and/or marketing and mailing techniques utilized by spam senders. Applicable laws, however, have limited regional application, and knowledge and enforcement of such laws is difficult due to the global nature of the e-mail infrastructure.” Schiavone, [0006]. Accordingly, to overcome these shortcomings (and others), Schiavone proposes as its title describes, “rule-based processing of electronic mail messages.” Schiavone, Cover. As discussed above, Schiavone explains that such processing would allow electronic mail messages “to negotiate transactions, enable automated verification of compliance with privacy, legal, regulatory and/or best practices standards, to perform spam filtering, anonymous profiling, target marketing through dynamic offers, and prioritization and standardization of e-mail to transactions.” Schiavone, [0001]. Additional benefits of Schiavone include “maintaining confidentiality of a recipient's profile data by storing ... data privately.” Schiavone, [0019].

120. Because of the aforementioned benefits Schiavone teaches, a POSITA would have been motivated to additionally include Schiavone’s teachings of “rule-based processing of electronic mail messages” to improve Rohall’s electronic messaging system. The teachings of Rohall and Schiavone thus provide a POSITA with a strong motivation to combine them. *First*, both references are in the same field and address the same issue in the same context—managing electronic mail messaging systems. Second, in Rohall’s system architecture, shadow documents (e.g., created from electronic mail messages sent among users) are represented by data structures including pointers pointing to other messages in a conversation thread such as children documents or shadows thereof. Rohall, [0017]. Indeed, Rohall explains the benefits of implementing shadows documents using data structures. **“Such implementation may be more efficient in an environment where the thread trees can be recomputed quickly relative to storing and retrieving the thread trees from a memory.”** Rohall, [0064]. *See also* Rohall, [0063], [0068], [0129] and Figs. 4, 6C.

121. A POSITA would have recognized that Schiavone’s teachings of “rule-based processing of electronic mail messages,” would be complimentary to the system of Rohall, which seeks to develop an improved inbox for messages by enhancing usability of messages and user productivity, e.g., by implementing shadows documents using data structures. Rohall, [0012]. A POSITA would have

thus recognized that combining Rohall with Schiavone would yield predictable benefits and results because the combined electronic mail message processing system would enforce various rules to ensure compliance, as well as providing filtering features, which is especially useful in today's increased use of emails. *See e.g.*, Rohall, [0001], Schiavone, [0002].

122. It would have been obvious to a POSITA that Schiavone's engines and data stores can be implemented on the messaging server in Rohall's system and further that such server would include a governance module as in Schiavone. Ex.1001, 14:44-46. *See also* Ex.1006, [0064] (disclosing "the recipient's communications device ... store[ing] and run[ning] ... a recipient compliance engine"). Numerous passages in Schiavone teach or suggest that functionalities of these engines and data stores can be implemented on a server. For example, Schiavone teaches that (1) the "recipient compliance engine" may reside at "the commercial mailer, trusted authority and/or another third party," instead of at a client device, and (2) "**mail server software ... is well known in the art.**" Ex.1006, [0066], [0064], [0030]. Schiavone explains that in "an alternate embodiment **the rule base 155 may also or alternatively be stored at the commercial mailer, trusted authority and/or another third party** and referenced as necessary via a communications network," and "[i]n another alternate embodiment, **a recipient compliance engine may reside at commercial mailer, trusted authority and/or**

another third party.” Ex.1006, [0066] (emphases added). Schiavone defines “**a third party trusted intermediary**” as a “**central source**” that is “responsible for defining rules, maintaining and updating rules, approving proposed rules, and distributing such rules to senders and recipients known to the trusted intermediary.” Schiavone, [0066], [0030]. From these teachings, a POSITA would have recognized that there are only two possible ways of implementing engines and data stores: at the server-side or at the client-side, and moreover, both implementations were known in the prior art. Schiavone, [0066].

123. Accordingly, a POSITA would have understood that server-side and client-side implementations of Schiavone’s engines and data stores represent a finite number of identified, predictable solutions to the design need of implementing engines and data stores. Thus, it would have been obvious to a POSITA that servers in Rohall’s systems (e.g., “*messaging system server*”) as modified by Schiavone’s teachings of rule-based processing of electronic mail messages would include Schiavone’s teachings of “recipient compliance engine,” “transaction data,” “rule base,” “public preference data,” “private profile data,” and “private commerce data,” collectively serving the claimed functionality of the “*governance module*.” Therefore, the combination of Schiavone and Rohall teaches or suggests that the “*messaging server*” includes the “*governance module*” as taught by Schiavone.

124. Furthermore, the references make clear that their systems require no specialized hardware. For example, Rohall's system includes a CPU, memory (temporary and permanent), and input (keyboard, mouse, microphone) and output devices (video display, speaker), and network components for local area network and wide area network (WAN). Rohall, [0041]–[0042]. Schiavone states that its techniques can be implemented on a “general-purpose computer running mail client software.” Schiavone, [0020]. The systems of both references also run on commercially available operating systems, such as Windows or Unix. Rohall, [0043]; Schiavone, [0065].

125. Rohall also teaches using metadata associated with email messages, which can be used to implement Schiavone's structure—i.e., message type specifier or a rule. See Rohall, [0057]–[0058], [0096]; Schiavone, [0007], [0024]–[0030]. Indeed, both Rohall and Schiavone are concerned about reducing memory/storage requirements. For example, Rohall describes “[a]n object is actually created by the program at runtime by means of a special function called a constructor which uses the corresponding class definition and additional information, such as arguments provided during object creation, to construct the object. Likewise, objects are destroyed by a special function called a destructor.” Rohall, [0046]. It is well-known that a constructor is used to allocate memory to an instance or object and a destructor is used to free up memory space occupied by objects created by the constructor.

Ex.1020, 25:10-26 (“For each type of resource, regardless of whether it is of the transparent or opaque kind, there is a particular way to create or allocate a unit of the resource, and to destroy or deallocate it. For example, memory is being allocated by calling the malloc() function and freed by calling the free() function. File objects are being created by calling the open() function and destroyed by calling the close() function. The creation function is generally called a “constructor” for the resource, and the destruction function is generally called a “destructor for the resource.”). Therefore, using well-known computer programming techniques, such as constructors and destructors for memory management demonstrates Rohall’s concern for reducing memory/storage requirements. Similarly, Schiavone identifies that “there are significant and increasing costs associated with the cost of ... storage ... absorbed by e-mail service providers and corporate domains.” Schiavone, [0002]. A POSITA would have thus recognized that Schiavone’s email processing techniques are compatible with Rohall’s system in many ways, and therefore would have had a reasonable expectation of success for combining them.

126. Additionally, a POSITA would have had a reasonable expectation of success in implementing Schiavone’s teachings (which include its enforcement and filtering methods) with Rohall’s system, because doing so would involve minimal changes to Rohall’s system architecture and the combined system uses the hardware and techniques already described in each reference, and each system would perform

the same in the combined system. For instance, a POSITA would recognize that Schiavone's teachings related to rule-based processing of electronic mail messages, e.g., identifying, storing, and processing rules would merely be a few additional process steps for messaging servers (e.g., servers 302-310) used in Rohall's architecture of an "electronic messaging system." Rohall, [0051]-[0052], and Figures 1 and 3. *See also* Rohall, [0060] (provides an exemplary implementation using the Lotus Notes/Domino electronic email system.). Rohall's invention, the "Mail Agent module" already implements² several process steps. Rohall, [0012], [0054]-[0056], [0073]. And, therefore, in one implementation, a POSITA would have been motivated to modify Rohall's system architecture (e.g., the Mail Agent Module) to additionally include, for instance, one or more extra sub-modules incorporating Schiavone's teachings related to "rule-based processing of electronic mail messages." And in another implementation, for instance, a POSITA would have understood to incorporate Schiavone's teachings related to "rule-based processing of electronic mail messages" directly into Rohall's Mail Agent Module without

² For example, parser 232, shadow document generator 234, both of which are included in the Mail Agent Module 230 implement the process steps outlined in Figures 5A and 5B. *See* Rohall, Figures 2, 5A, 5B, and [0060].

developing a sub-module. A POSITA would have the ability to make the combination cited above, and would have further understood these options as obvious-to-try and therefore had good reason to pursue them.

127. A POSITA would be motivated to combine Rohall with Schiavone because to do so would have been the arrangement of old elements (an electronic mail system comprising messaging servers, use of data structures for representing conversation threads and messages within the electronic mail system, modules applying rules of compliance policies to electronic messages, database servers) with each performing the same function it has been known to perform (e.g., determining and applying different rules to different messages, storing and applying rules of compliance policies to messages) and yielding no more than what one would expect from such an arrangement (an improved electronic mail system/inbox that includes automated filtering capabilities), as Rohall demonstrates. *See*, Rohall, [0012].

1. Independent Claims 1, 10, and 13:

a. Preamble: Electronic Messaging System

[1.0] A computerized messaging server in an electronic messaging system, comprising:

[10.0] A computer program product having a non-transitory computer-readable medium having embodied thereon program code for use in an electronic messaging system, the program code comprising:

[13.0] A computer-implemented method of managing messages in an electronic messaging system, comprising:


128. Elements [1.0], [10.0], [13.0] relate generally to “[a] computerized messaging server in an electronic messaging system.” Rohall discloses these elements. Given the similarity in the claim language recited in these elements, they are addressed below simultaneously.

129. Rohall discloses “an electronic messaging system.” Rohall describes that its invention “may be implemented using any number of **electronic mail systems**” (“*electronic messaging system[s]*”). Rohall provides an exemplary implementation in which the “electronic mail system” (“*an electronic messaging system*”) is the Lotus Notes/Domino electronic mail domain. Rohall, [0060]. *See also* Rohall, [0048] (disclosing “The illustrative embodiment of the invention may be implemented as part of Lotus Notes® and a Lotus Domino server, both commercially available from Lotus Development Corporation, Cambridge, Mass.”) While the exemplary implementation in Rohall uses “Lotus Notes and Domino products from IBM Corporation,” Rohall explains that its “**inventive algorithm ... is not specific to Lotus Notes, Domino** or Intelligent Miner for Text.” Rohall, [0073].


130. Furthermore, with reference to Figs. 1 and 3, Rohall explains that the “*electronic messaging system*” executes on computer architecture (e.g., as depicted

in Fig. 1³) and operates in a network environment (e.g., as depicted in Fig. 3). “[A] **computer system operatively connectable to a network and capable of executing a communication process for sending and receiving electronic mail documents [and includes] servers 302-310.**” Rohall, [0014]. Specifically, the network environment of Fig. 3 “in which the invention may be practiced” includes “a packet-switched data network 300” comprising “servers 302-310” that “may interface with global network 320,” and that “may be implemented as part of an all software application which executes on a computer architecture similar to that described with reference to FIG. 1.” Rohall, [0051]-[0052]. Accordingly, servers (e.g., servers 302-310 in Rohall’s annotated Fig. 3) included in Rohall’s electronic mail system of Rohall correspond to “*computerized messaging server[s] in an electronic messaging system*”). The network computer architecture depicted in Figure 1 and the network environment depicted in Figure 3 of Rohall are reproduced below:

³ Although Rohall’s Fig. 1 is captioned as “Prior Art,” Rohall describes it as “a computer systems [sic] suitable for use with the present invention.” Rohall, [0020].



Rohall, Figure 3 (annotated).



Rohall, Figure 1.

131. Therefore, Rohall discloses “[a] *computerized messaging server in an electronic messaging system.*”

132. Additionally, Rohall’s “electronic messaging system” uses “[a] *computer program product having a non-transitory computer-readable medium having embodied thereon program code*” recited in claim 10. For example, Rohall explains that “[a] **software implementation of the above-described embodiments may comprise a series of computer instructions either fixed on a tangible medium, such as a computer readable media, e.g. diskette 142, CD-ROM 147, ROM 115, or fixed disk 152 of FIG. 1A [sic: FIG. 1], ... over a medium [that] can be either a tangible medium, including but not limited to optical or analog communications lines ... The series of computer instructions [which] embod[y] all or part of the functionality previously described herein with respect to the invention ... can be written in a number of programming languages for use with many computer architectures or operating systems.**” Rohall, [0133] (emphases added).

133. Rohall’s “*electronic messaging system*” also teaches “[a] *computer-implemented method of managing messages.*” Rohall describes that “**the methods of the invention [relating to a technique for effectively reviewing and processing electronic mail and electronic mail threads] may be achieved in either all**

software implementations, using the appropriate processor instructions, or in hybrid implementations which utilize a combination of hardware logic and software logic to achieve the same results.” Rohall, [0134] and [0005]. Accordingly, Rohall’s “*electronic messaging system*” provides “[a] *computer-implemented method of managing messages.*”

134. Therefore, Rohall discloses elements [1.0], [10.0], [13.0].

b. Messaging Module Controlling a Message Database

[1.1] a messaging module adapted to control a message database storing messages sent among users of the messaging system,

[10.1] a messaging module adapted to control a message database storing messages sent among users of the messaging system to at least one other end user of the messaging system,

[13.2] receiving a message sent by one end user of the messaging system to at least one other end user of the messaging system,

135. Elements [1.1], [10.1], [13.2] generally relate to “*a messaging module adapted to control a message database storing messages sent among users of the messaging system.*” Rohall discloses these elements.

136. Initially, the 911 Patent describes that a “‘message’ refers to a data communication sent by one end-user to one or more end-users of the messaging system or another messaging system.” Ex. 1001, 3:4–6. Likewise, Rohall teaches a

“system ... for sending and receiving electronic mail documents,” where document refers to “a document, database, electronic mail message code, a ‘Note’ or any file which is accessible and storable by a computer system.” Rohall, [0014], [0050].

137. First, Rohall discloses “*a messaging module*.”

To implement the primary functionality of the present invention in a Lotus Notes environment, **a module, referred to hereafter as Notes Mail Agent 230** interacts with the existing functionality, routines or commands of Lotus Notes client application and/or a Lotus “Domino” server, many of which are publicly available.


Rohall, [0054] (emphasis added).

138. The “Mail Agent module⁴” in Rohall corresponds to “*a messaging module*” of the 911 Patent. Rohall, [0073]. For example, Rohall explains that the “Mail Agent module” can be implemented as “a back-end module for processing electronic **mail messages**.” Rohall, [0073]. Rohall further explains that the primary function of the Mail Agent module (“*a messaging module*”) is “**to create a shadow document from an original document ... [such as] an electronic mail message**.” Rohall, [0055], emphasis added. Rohall provides an example in which the Mail

⁴ Rohall alternatively refers to the “Mail Agent module” as “Mail Agent.” See Rohall, [0054]-[0055]

Agent module is implemented as a “Notes Mail Agent” or “Lotus Notes Mail Agent.” *Id.*, [0054]-[0055].

139. Rohall’s Figure 2 (reproduced below) depicts the Mail Agent module (“a messaging module”), e.g., referenced in Figure 2 as Mail Agent Module 230 and annotated in green.



Rohall, Figure 2 (annotated).

140. Additionally, Rohall’s Mail Agent module (“a messaging module”) is “adapted to control a message database storing messages sent among users of the messaging system.” For example, Rohall discloses that the Mail Agent module (“a

messaging module”), **“invoked upon the sending of an electronic mail message,”** **“creates a shadow document from ... [the] electronic mail message,”** and **“interacts with one or more documents 260 in databases 250”**⁵ (*“a message database”*) for example, with regards to storing shadow documents. Rohall, [0054]-[0055]. Rohall provides an example in which databases 250 (*“a message database”*) corresponds to a Notes mail database.

Mail Agent 230 creates a shadow document from an original document by generating a file containing data related to the document. ... **[S]hadow documents are stored as documents in a Lotus Notes database Specifically, shadow documents are stored in a Notes mail database.** ... The shadow documents are generated transparently to the actual user sending or receiving the electronic message.

Rohall, [0055]-[0056]. Rohall uses surrogate document to describe shadow documents: “Using the shadow document generation techniques described herein, a surrogate document including meta data such as size, date, heading information and a pointer to the pointer actual data, is generated by Notes Mail Agent 230.” Rohall, [0096]. “Mail agent 230 upon receiving data associated with a particular user

⁵ Although Rohall references databases 250 (presumably in connection with Fig. 2), Rohall’s Fig. 2 does not depict databases 250.

generates a shadow or circuit [sic:surrogate] document, in a manner as described herein and transmits the surrogate document to inbox 902.” Rohall, [0097].

141. In view of the above teachings of Rohall, and further because Rohall discloses shadow documents are created from electronic mail messages sent among users upon the sending of electronic mail messages in a manner such that their creation is transparent to users of the electronic mail system and further shadow documents are stored as documents in a Lotus Notes database, therefore, shadow documents are internal surrogates of electronic mail messages representing relationships between messages within Rohall’s architecture. *See* Rohall, [0055]-[0056]. And, further, because a shadow document includes a pointer to the original document / actual message itself, shadow documents are not copies of messages. *See* Rohall, [0057] and [0100]. Therefore, in Rohall, “original electronic mail messages sent among users” correspond to the claimed “*messages sent among users*,” which are internally represented within Rohall’s system as “shadow documents” (“*at least one of the sent messages . . . comprising a message container*”) that include data structures having pointers to other messages in a conversation thread such as children documents or shadows thereof (“*containing relational references pointing to a plurality of sent submessages.*”)

142. Indeed, the use of “shadow documents” representing relationships between messages appears to be known in IBM patents. *See, e.g.*, Ex.1011, [0061]-

[0064], Ex.1012, [0059]-[0062]. Further, as early as 1997, a technical paper described using “shadow documents” in connection with developing educational course materials for distance learning. *See, e.g.*, Ex.1013, 14-15. Furthermore, because Rohall discloses that shadow documents are stored in databases 250 (“*a message database*”) such as a Notes mail database, Rohall discloses “*a message database storing messages sent among users of the messaging system.*” And moreover, in view of Rohall’s disclosure that the Mail Agent module (“*a messaging module*”) is “invoked upon the sending of an electronic mail message” and “interacts with one or more documents 260 in databases 250” (“*a message database*”), Rohall’s Mail Agent module is therefore “*a messaging module adapted to control a message database storing messages sent among users of the messaging system to at least one other end user of the messaging system.*”

143. Rohall elaborates on the use of shadow documents for defining a parent/child hierarchy with respect to electronic messages sent in electronic mail systems:

The data maintained in a shadow document defines the parent/child relationships among original documents and their respective shadow documents. In the illustrative embodiment, a new electronic mail message is considered a parent document and serves as the root from which a new shadow tree may be derived, as explained hereinafter. Any replies to the original electronic mail message is/are considered a child/children document(s). Within a conversation thread, and a hierarchical tree that represents

such thread, children documents derive from a common root document. Accordingly, a parent/child tree hierarchy representing a conversation thread terminates at one extreme with a root document, or a shadow document thereof, and, at the other extreme, with one or more children documents, or shadows thereof, as the leaves of the tree.

Rohall, [0056]. As such, “[t]he conversation thread, that may include a parent/child hierarchy of a plurality of original electronic mail documents, summaries of electronic mail documents, shadow documents of original electronic mail documents, or any combination thereof, may be generated dynamically and the data structure representing the thread forwarded to another electronic mail user.” Rohall, [0125].

144. In addition, because the “*messages sent among users of the messaging system*” encompass “*messages sent ... to at least one other end user of the messaging system*” and also “*receiving a message sent by one end user of the messaging system to at least one other end user of the messaging system,*” Rohall discloses element [10.1] and [13.2].

145. Therefore, Rohall discloses elements [1.1], [10.1], [13.2].

c. Sent Message comprising a Message Container

[1.2] at least one of the sent messages stored in the message database comprising a message container containing relational references pointing to a plurality of sent submessages stored externally to the message container in the message database,

[10.2] at least one of the sent messages comprising a message container containing relational references pointing to a plurality of sent submessages stored in the message database externally to the message container,

[13.3] the sent message comprising a message container containing relational references pointing to a plurality of sent submessages stored in a message database externally to the message container,

146. Breaking down the language of Elements [1.2], [10.2], and [13.3], these elements generally relate to “*at least one of the sent messages stored in the message database comprising a message container containing relational references pointing to a plurality of sent submessages,*” and that “*the plurality of sent submessages [is] stored externally to the message container,*” and further that “*the plurality of sent submessages [is] stored ... in the message database.*” Rohall discloses these elements.

147. Initially, the 911 Patent states that “[t]he messaging system treats a set of related messages, such as an email string between two or more people, as a message container (200) having relational references to one or more submessages (210, 212, 214).” Ex.1001, 1:65–2:2. Also, “[t]he container itself does not contain content, but rather points to submessages and/or attachments in which content resides.” *Id.*, 4:62–64. The 911 Patent describes that a submessage is defined as the

message that the end user actually composes: “When an end-user composes and sends a message, she is actually composing a submessage, and then sending a message 200 containing a reference to the submessage 200 to other end-users.”

Ex.1001, 5:10–13.

148. Rohall discloses “*at least one of the sent messages stored in the message database comprising a message container containing relational references pointing to a plurality of sent submessages.*” Rohall explains that the Mail Agent module is “invoked upon the sending of an electronic mail message,” “creates a shadow document from ... [the] electronic mail message,” and “interacts with one or more documents 260 in databases 250” (“*a message database*”) for example, with regards to storing shadow documents. Ex.1005, [0054]-[0055]. Rohall provides an example in which databases 250 (“message database[s]”) corresponds to a Notes mail database and identifies that “shadow documents are stored as documents in a Lotus Notes database”. Rohall, [0056] and [0129]. Rohall also explains that a new electronic mail message is considered a parent document and **serves as the root from which a new shadow tree may be derived.**” Ex.1005, [0056]. Rohall further explains that “Mail Agent 230 creates a shadow document from an original document by generating a file containing data related to the document.” Ex.1005, [0055]. Thus, in Rohall, electronic messages (e.g., “*sent messages*”) are represented internally or characterized by “shadow documents” (“*at least one of the sent*

messages . . . comprising a message container”) that include data structures having pointers to other messages in a conversation thread such as children documents or shadows thereof (“*containing relational references pointing to a plurality of sent submessages.*”)

149. Further, Rohall provides details of data structures used to maintain data associated with a conversation thread of an electronic mail document, e.g., in connection with sending and receiving electronic mail documents. Specifically, Rohall teaches:

a data structure stored in the memory and usable to maintain data associated with a conversation thread of an electronic mail document, the data structure comprising: (a) data identifying the electronic mail document; (b) data identifying a root document of the conversation thread; and (c) data identifying any of the child documents of the root document.

Rohall, [0017]. Rohall explains: “[t]he data maintained in a shadow document defines the parent/child relationships among original documents and their respective shadow documents.” Rohall, [0056]. Rohall further explains that the meta data in shadow documents may include such logistical information as sender, receiver, original size, subject, date, any carbon copy recipients, etc. associated with the document. **In addition, key words or summaries of the content of the document or any attachments may likewise be included.** Ex.1005, [0058]. Indeed, Rohall

explains the benefits of implementing shadow documents using data structures:
“[s]uch implementation may be **more efficient.**” Ex.1005, [0064].

150. Rohall elaborates on the structure and content of its data structure. For example, Rohall explains that its shadow document data structures include pointers to other messages in a conversation thread (“*relational references pointing to*”), such as children documents or shadows thereof (“*a plurality of sent submessages*”). Specifically, with reference to Fig. 4, Rohall provides examples of the data maintained in a shadow document:

[S]hadow document 400 comprises an Original Document Identified (ID) 402, a Parent Document ID 404, an optional Root Document ID 406, **zero or more Child Document IDs 408 a-n**, and optional Meta Data fields 410 a-n. ... **Child Document IDs 408 a-n may comprise a list of pointers to the immediate children documents, whether shadow or original, in the tree hierarchy.**

Rohall, [0057] (emphases added). Additionally, with reference to Fig. 17B, Rohall describes:

[D]ata structure 1720 comprises an Original Document Identified (ID) 1722, Root Document ID 1724, zero or more Child Document IDs 1726A-N, and Meta Data fields 1728A-N. ... Child Document IDs 1726A-N may comprise a list of pointers to the immediate children documents, whether shadow or original, in the tree hierarchy.

Rohall, [0129] (emphases added). In sum, Rohall explains that “[t]he data maintained [by data structures] in a shadow document defines the parent/child

relationships among original documents and their respective shadow documents.” Rohall, [0056].

151. Indeed, Rohall explains the benefits of implementing shadow documents using data structures. “**Such implementation may be more efficient** in an environment where the thread trees can be recomputed quickly relative to storing and retrieving the thread trees from a memory.” Rohall, [0064]. .

152. In Rohall, “children documents or shadows thereof” constitute “*a plurality of sent submessages*.” Rohall discloses a “system ... for sending and receiving electronic mail documents,” where “document” refers to “a document, database, electronic mail message code, a ‘Note’ or any file which is accessible and storable by a computer system.” Rohall, [0014], [0050]. Further, Rohall defines “children documents” as follows:


Any replies to the original electronic mail message is/are considered a child/children document(s). Within a conversation thread, and a hierarchical tree that represents such thread, children documents derive from a common root document. Accordingly, **a parent/child tree hierarchy representing a conversation thread** terminates at one extreme with a root document, or a shadow document thereof, and, at the other extreme, with **one or more children documents, or shadows thereof**, as the leaves of the tree.

Rohall, [0056], emphases added. Therefore, in Rohall, “children documents” (or, as Rohall describes, shadows thereof) are essentially—electronic mail messages sent

as replies to an original message and arranged in a parent/child hierarchy with respect to the original message.


153. With reference to Figure 6C⁶ (reproduced below) Rohall explains the relationship between shadow documents created from electronic mail messages sent among users and “children documents or, shadows thereof. This example shows a conversation thread among users Al, Bob, and Charlie. There, a message sent from Al to Bob and Charlie is the “*sent message*” and serves as an original document (original document 602B). Documents 604B, 606B, and 608B are replies or replies-to-replies of the sent message and are therefore child documents (“*plurality of sent submessages*”) of original document 602B. Shadow documents 602C (red), 604C (blue), 606C (green), 608C (yellow) created from their corresponding original documents 602B (red), 604B (blue), 606B (green), 608B (yellow) are representations of their corresponding original documents within Rohall’s architecture. *See generally* Rohall, [0068].

⁶ While Rohall uses Figure 6C specifically in the context of deletion of an original document, the explanation setting forth the relationship between “shadow documents created from electronic mail messages” and “children document” still applies.




Rohall, Figure 6C, (shadow documents created from their respective original document are annotated using the same color).

154. Moreover, in Figures 4 and 17B, (reproduced below) which relate to Rohall's shadow document, its data structure, and their use in implementing a conversation thread, Rohall shows multiple children document IDs, which therefore identify multiple children documents.



Rohall, Fig. 4.



Rohall, Fig. 17B. Therefore, in Rohall, “children documents” (or, shadows thereof) constitute “*a plurality of sent submessages.*”

155. Moreover, Rohall discloses building “conversation threads” using shadow document data structures that include pointers. Like the 911 Patent, which describes that “[t]he messaging system treats a set of related messages, such as an email string between two or more people, as a message container (200) having relational references to one or more submessages (210, 212, 214),” Rohall explains that its system builds “conversation threads” comprising electronic mail messages in a parent/child hierarchy. Ex.1001, 1:65-2:2. Rohall explains:

The conversation thread, that may include a parent/child hierarchy of a plurality of original electronic mail documents, summaries of electronic mail documents, shadow documents of original electronic mail documents, or any combination thereof, may be generated dynamically and the data structure representing the thread forwarded to another electronic mail user.

Rohall, [0013] (emphases added). Rohall further elaborates on how “conversation threads” is built using a Tree Builder (such as Tree Builder 236).

Given the content of shadow documents and their relationship to the original or root document, an algorithm in Tree Builder 236 can be used to traverse the chain of pointers or references to the parent of each shadow document, and, once the root has been identified, to then recursively traverse all references to each child document within the conversation thread. In this manner, a complete tree representing the conversation thread may be determined from the data collected by Tree Builder 236.

Rohall, [0062] (emphasis added). Further, Rohall describes in detail the method of constructing “a conversation thread”. *See, e.g.*, Rohall, [0063], Fig. 5B disclosing “Tree Builder 236 resolves the child document IDs 408 a-n in the document identified by the pointer within the Root Document ID field, i.e. the root document, as well as each of their respective child documents, in a recursive manner, as will be understood by those reasonably skilled in the arts, until the Child Document ID fields in all child documents are null, indicating that the leaf nodes within the conversation thread tree have been identified, as illustrated by steps 528”. Therefore, Rohall discloses “*at least one of the sent messages stored in the message database comprising a message container containing relational references*” by building “conversation threads” using shadow document data structures that include pointers.

156. Indeed, Rohall discloses “shadow documents are stored as documents in a Lotus Notes database” and “[a] **Notes database ... acts as a container in which data Notes ...** compris[ing] user defined documents and data” may be grouped. Rohall, [0049]. Rohall notes that a “document” such as a user-defined document “may refer to a document, database, electronic mail message code, a “Note” or any file which is accessible and storable by a computer system.” Ex.1005, [0050]. Thus, for example, a reply by a user to an original email message can be added to the database (e.g., Notes database) as a children document or an additional submessage sent by the user. Therefore, Rohall discloses “*at least one of the sent messages stored*

in the message database comprising a message container containing relational references pointing to a plurality of sent submessages.”

157. Moreover, to the extent one asserts that Rohall does not expressly or implicitly disclose “*at least one of the sent messages stored in the message database comprising a message container . . .*” as demonstrated above, it would have been obvious in view of Rohall’s express teachings related to hybrid conversation trees, which include original documents and shadow documents. Ex.1005, [0068] (“**a hybrid tree consisting of shadow documents 602C-604C and original documents 606D and 608D**”). Such a modification, for example, can provide the benefit of faster searches in an email search application where emails containing a certain term are to be searched. Thus, having the original email included with its shadow document would reduce the access time required to retrieve the original email from memory facilitating faster searches. A POSITA would have had a reasonable expectation of success in implementing the modification above because it is analogous to “caching,” a well-known technique providing faster access to data with the goal of improving application and system performance. For example, Markatos explains that “[t]o reduce the overhead imposed by frequent document requests, we propose the notion of caching a World Wide Web server’s documents in its main memory.” Ex.1021, Abstract. Specifically, Markatos proposes “caching of local documents in a Web server’s main memory, so that they can be sent faster

to clients that request them” and thereby, Markatos’s system “reduces a server’s load by keeping the most frequently accessed documents in the server’s main memory, thus avoiding disk accesses as much as possible.” Ex.1021, 894 and 904.

158. Markatos further explains:

Some of the well known caches include the Harvest hierarchical Cache [4], the Lagoon system [3], the push-caching approach at Harvard [5], the caching work at Boston University [1,2], the Hensa Archive [8], etc. All these approaches deal with caching of remote documents on a disk that is close to the clients that request the document. Several of them are based on proxy servers that service a community of users ranging from one building to a whole country. When one user asks for a document, the proxy caches it, so that next accesses to the same document by nearby users will receive the cached copy.

Ex.1021, 904.

159. Moreover, Rohall discloses this limitation consistent with Patent Owner’s apparent interpretation of this claim limitation in its PIC in which Patent Owner has mapped this claim limitation to “defin[ing] conversations based on the Message-ID value of the first email message in a thread” in Microsoft Exchange Online. PIC, 7.

Microsoft Exchange Online

Microsoft Exchange Online meets this limitation.

In Microsoft Exchange Online, at least one of the sent messages stored in the message database comprises a message container containing relational references pointing to a plurality of sent submessages stored externally to the message container in the message database.

The plurality of submessages refers to the inventive object-based messaging system discussed above. Each submessage is a unique object.

“In the context of Exchange, conversations are a way to group and manage a related set of email messages. They can also provide a way to view related messages. Exchange defines conversations based on the Message-ID value of the first email message in a thread. All replies and related messages reference the original message's Message-ID header in their References and In-Reply-To headers.

Additionally, inside the SOAP envelope, for each message received in a mailbox, Exchange sets specific properties and elements.”

- <https://learn.microsoft.com/en-us/exchange/client-developer/exchange-web-services/how-to-work-with-conversations-by-using-ews-in-exchange>

PIC, 7.

160. In addition, Rohall discloses that the “*plurality of sent submessages [is] stored externally to the message container.*” Specifically, Rohall discloses that “children documents or shadows thereof” (“*plurality of sent submessages*”) are stored externally to” the shadow document. With reference to Figures 4 and 17B, Rohall describes:

[S]hadow document 400 comprises an Original Document Identified (ID) 402, ... zero or more **Child Document IDs 408 a-n ...[that] comprise a list of pointers to the immediate children documents**, whether shadow or original, in the tree hierarchy.

[D]ata structure 1720 comprises an Original Document Identified (ID) 1722, ... zero or more **Child Document IDs 1726A-N ... [that] comprise a list of pointers to the immediate children documents**, whether shadow or original, in the tree hierarchy.

Rohall, [0057], [0129] (emphases added). In the above disclosures, Rohall teaches shadow documents represented by data structures as specifically comprising a list of pointers to the children documents, and therefore, Rohall discloses that “children documents or shadows thereof” (“*plurality of sent submessages*”) are “*stored externally to*” the shadow document. Notably, in the above disclosures, Rohall describes that while for the original document, its ID is included in the shadow document, for children documents (stored outside the shadow document and identified by Child Document IDs)—a list of pointers to the children documents is necessary. Because a pointer is a reference that provides an unambiguous location of a resource or object (*see* Section VIII.C), a POSITA would have understood that Rohall’s description of shadow documents including “pointers” to children documents would be understood by a POSITA to be describing values that “point” to the “location” of the children document data, which makes clear that “children documents or shadows thereof” (“*plurality of sent submessages*”) are “*stored externally to*” the shadow document. This is consistent with the 911 Patent’s description of “*relational references*,” which also point to the **location** of the message content. “A message can be thought of as a container with **relational references**. The container itself **does not contain content**, but rather **points to submessages and/or attachments in which content resides.**” Ex.1001, 4:63-63.

“Using the shadow document generation techniques ..., **a surrogate document**

including meta data such as size, date, heading information and **a pointer to the pointer actual data**, is generated by Notes Mail Agent 230.” Rohall, [0096]; see also background describing how pointers work in Section VIII.C. Indeed, Rohall discloses “*the plurality of sent submessages stored externally to the message container*” because Rohall teaches that “[t]he **Mail agent 230** processes each received item and ... **stores** either the complete document or **a surrogate document, complete with a pointer to the document location, within the folder or mail database.**” Rohall, [0100]. Rohall uses surrogate document to describe shadow documents: “Mail agent 230 upon receiving data associated with a particular user generates a shadow or circuit [sic:surrogate] document, in a manner as described herein and transmits the surrogate document to inbox 902.” Rohall, [0097]. Therefore, in view of the above-mentioned disclosures, Rohall discloses “*the plurality of sent submessages stored externally to the message container.*”

161. Moreover, for the additional reasons explained below, Rohall suggests or teaches that “*the plurality of sent submessages [is] stored externally to the message container.*” For example, Rohall teaches that a shadow document identifying any parent and/or child documents of the subject document “is stored in an electronic mail database which may then be replicated” and further that “[a]ll other copies of database 350 within the network are replica copies of the master copy,” Rohall, [0052], [0059]. Therefore, by storing messages distributed across


databases containing replicas of a master database, Rohall suggests or teaches that the “*the plurality of sent submessages [is] stored externally to the message container*” for this additional reason.

162. Indeed, Rohall discloses that “*the plurality of sent submessages [is] stored ... in the message database.*” Rohall discloses storing “children documents” in databases 250 (“a message database”) such as a Notes mail database implemented using the Lotus Notes/Domino architecture, which is built on the premise of databases and replication thereof.

A Notes database, referred to hereafter as simply a “database”, acts as a container in which data **Notes [comprising user defined documents and data]** and design Notes [comprising application elements such as code or logic that make applications function] may be grouped. ... In Notes, **every database has a master copy which typically resides on the server or user platform where the database was created. All other copies of the database are replicas of the master copy.** Replicas of databases may be located remotely over a wide area network, which may include as a portion thereof one or more local area networks.

Rohall, [0048]-[0049] (emphases added). “[A] shadow document ... identifies any parent and/or child documents of the subject document” and further that it “is stored in an electronic mail database which may then be replicated ... in the Notes environment.” Rohall, [0059]. “Thereafter, the shadow document is stored in memory, which, in the illustrative embodiment, is a mail database.” Rohall, [0060].

163. With reference to Fig. 3, Rohall explains that “the copy of database 350 associated [which includes document 360] with server 310 is designated as the ‘master’ copy of database 350,” where “[a]ll other copies of database 350 within the network are replica copies of the master copy.” Rohall, [0052].




Rohall, Figure 3 (annotated). *See also* Rohall, [0050] (disclosing that “[t]he Notes Storage Facility (NSF) architecture defines the manner in which documents and databases are created, modified and replicated among Notes servers across a computer network.”). Thus, by teaching that a database stores user defined documents and data, including e.g., “children documents” (or, shadows thereof) and further, by teaching that content stored in databases is replicated, Rohall teaches that every database in its system has the same set of “children documents” or shadows

thereof. And, as such, Rohall discloses that “children documents” (or, shadows thereof) constitute “*the plurality of sent submessages*” and are “*stored ... in*” databases 250 (“*message database*”).

164. Therefore, Rohall teaches “at least one of the sent messages stored in the message database comprising a message container containing relational references pointing to a plurality of sent submessages” and that “the plurality of sent submessages [is] stored externally to the message container,” and further that “the plurality of sent submessages [is] stored ... in the message database.”

165. Moreover, Rohall’s disclosures of the “*message container ...*” limitation using replica databases that store copies of the messages stored in the master is consistent with Patent Owner’s statements during prosecution in which Patent Owner admitted that “the content of [the submessages] is stored elsewhere, such as on the messaging server 112, or in local caches 118 [storing local copies of messages held in the message store database 120⁷], as depicted in FIG. 1.” Ex.1002, 243-244 (citing Fig. 1, paragraph [0026] of the as-filed application).

⁷ There is a typo in 4:21 in the 911 Patent Specification (or, equivalently paragraph [0026] of the as-filed application)—the “message store database” should be labeled “120” and not 118. Fig. 1 of the 911 Patent confirms this.



Ex.1001, Fig. 1.

166. Therefore, Rohall discloses elements [1.2], [10.2], [13.3].

- d. Sent submessages Are a Reply-To or a Forward-Of
[1.3] at least one of the sent submessages⁸ being one of a reply to
and a forward of other sent submessages of the sent message;
and

⁸ It appears that reciting “*the sent submessages*” here instead of “*the plurality of sent submessages*” is a typographical error.

[10.3] at least one of the sent submessages⁷ being one of a reply to and a forward of other sent submessages of the sent message; and

[13.4] the plurality of sent submessages including a first sent submessage and a second sent submessage, at least one of the sent submessages being one of a reply to and a forward of another sent submessage of the plurality of sent submessages; and

167. Elements [1.3], [10.3], [13.4] generally relate to “*at least one of the sent submessages being ... a reply to the sent message*” or “*a forward of other sent submessages.*” Rohall discloses these claim elements.

168. Rohall discloses “*at least one of the sent submessages being ... a reply to the sent message.*” As explained above with respect to element [1.2], in Rohall, “children documents or shadows thereof” constitute “*a plurality of sent submessages.*” And, as explained below, Rohall also teaches that “a child/children document(s) or shadows thereof” is/are “*a reply to the sent message.*”

Any replies to the original electronic mail message is/are considered a child/children document(s). Within a conversation thread, and a hierarchical tree that represents such thread, children documents derive from a common root document. Accordingly, **a parent/child tree hierarchy representing a conversation thread** terminates at one extreme with a root document, or a shadow document thereof, and, at the other extreme, with

one or more children documents, or shadows thereof,
as the leaves of the tree.

Rohall, [0056] (emphases added). Rohall further explains how to determine if a particular message (which can be “*one of the sent submessages*”) is “*a reply to the sent message.*”

Within a Lotus Notes electronic mail domain, it is possible to determine with existing object methods to which message an incoming message is a reply. If the incoming message is from the Internet or from a non-Notes environment, a conventional algorithm can be used to determine if the message is a new message or a reply to an existing message by determining if the message has an “In-Reply-To” header, or whether the subject lines of the message match an existing message or shadow document.

Id., [0060] (emphasis added).

169. Using these methods, Rohall explains that “a complete conversation tree” can be created to include reply messages:

The reader can appreciate that there are alternative ways to maintain and/or compute threads within an electronic mail database, e.g., **the use of In-Reply-To headers where a document refers to its parent.** The shadow documents disclosed herein provide a complete conversation tree, the way that a discussion database would have complete thread trees.

Id., [0091] (emphasis added). See also Rohall, [0093] (disclosing an “electronic mail box,” that includes “a map 912 illustrating other messages in the conversation thread—the Ccs, **the Reply Tos**, the forwards”). Therefore, Rohall discloses to “*at least one of the sent submessages being ... a reply to the sent message.*”

170. Rohall also teaches “*at least one of the sent submessages being ... a forward of other sent submessages.*”

171. Rohall teaches that a submessage is a forward of other sent submessages of the sent message. Paragraphs [0013] and [0125] in Rohall describe, in relevant parts, that “**an electronic mail conversation thread may be handled or processed as a single entity with functions similar to those typically applied to an individual electronic mail message, such as including any of a select, drag and drop, copy, paste, delete, *forward*, print, move, file, archive, or operation on the data structure representing all ancestor documents in the conversation thread.**” Rohall, [0013], [0125] (emphases added). *See also* Rohall, [0093] (disclosing an “electronic mail box,” that includes “a map 912 illustrating other messages in the conversation thread—the Ccs, the Reply Tos, **the forwards**”) (emphasis added). Therefore, Rohall discloses “at least one of the sent submessages being ... a forward of other sent submessages.”

172. Additionally, as explained above with respect to element [1.2], in Rohall, “children documents or shadows thereof” constitute “*a plurality of sent submessages.*” And, therefore, “*a plurality of sent submessages*” encompasses “a first sent submessage and a second sent submessage,” as required in element [13.3].

173. Therefore, Rohall discloses elements [1.3], [10.3], [13.4].

e. Governance Module controlling a Governance Policy

[1.4] a governance module adapted to control a governance policy database storing a compliance policy describing rules applicable to the sent submessages,

[10.4] a governance module adapted to control a governance policy database storing a compliance policy describing rules applicable to the sent submessages,

[13.1] defining a compliance policy describing rules applicable to sent submessages of sent messages in the electronic messaging system;


174. Elements [1.4], [10.4], [13.1] generally relate to “*a governance module adapted to control a governance policy database storing a compliance policy describing rules applicable to the sent submessages.*”

175. Although Rohall does not expressly disclose these elements, Schiavone combined with Rohall (see Section X.B) discloses these elements because Schiavone teaches a recipient compliance engine identifying rules and referencing a rule base, as explained further below.

176. The 911 Patent describes that the “governance module 422 (also called a “compliance module”) applies governance policies (also called “compliance policies”) to the messaging system and controls the governance policy database


424.” Ex.1001, 7:62–65. For example, “the governance module 422 specifies compliance policies in terms of populations, rules, and jobcodes.” *Id.*, 8:12–14.

177. With reference to Figure 4 (reproduced below), the 911 Patent states that the “*governance policy database*” is included in the “*governance module*” which is included in the “*messaging server*.” “[T]he **messaging server 112** **includes** a messaging module 410, an auditing module 412, a security module 414, **and a governance module 422** ... [and further that] [t]hese modules respectively **contain** a message database 416, an audit information database 418, a security database 420, and **a governance policy database 424.**” Ex.1001, 5:45-50.



Ex.1001, Figure 4. With reference to Figure 5 (reproduced below), the 911 Patent notes that the “*governance module*” is inside the client module. *Id.*, 14:44-46 (“In

one embodiment, the client module 510 includes a governance module 516 for storing one or more governance policies received from the messaging server 112”).



Ex. 1001, Figure 5.

178. Schiavone discloses “a governance module adapted to control a governance policy database storing a compliance policy describing rules applicable to the sent submessages.”


179. First, Schiavone discloses “a governance module adapted to control a governance policy database storing a compliance policy describing rules”:

[A] **recipient compliance engine 154** [that] is configured to identify a rule associated with the message, including referencing a rule base 155 as necessary, and to process the rule, e.g., by referring to public preference data 120 a, private profile data 120 b and/or private commerce data 120 c stored in the recipient client’s data store 120. [T]he recipient compliance engine may reference private profile

data and modify the message in accordance with the rule as a function of the recipient's profile data. The recipient compliance engine 154 may store transaction data in a transaction data data [sic] store 158, which may be reference by a third party such as trusted authority 200's accounting engine 220 for monitoring such transactions, e.g., for advertising revenue purposes.

Schiavone, [0064]. Thus, by **“identify[ing] a rule associated with the message, including referencing a rule base 155⁹ and ... process[ing] the rule,”** the “recipient compliance engine” in Schiavone (e.g., included in a “*governance module*”) serves the claimed functionality of “*control[ling]*” a database (“*a governance policy database*”) including (“*storing*”) “rule base 155” (“*[one or more] compliance polic[ies] describing rules*”). Ex.1006, [0064] (emphasis added). With reference to Schiavone's Fig. 6 (reproduced below), the “*governance module*” (shown with a solid blue bounding box) controls the “*governance policy database*” (shown with a dashed red bounding box), which is shown as including transaction data store 158, rule base 155, a public preference data store 120a, a private profile data store 120b, and a private commerce data store 120c.

⁹ Schiavone alternatively refers to “rule base 155” as “data store of rules 155” in [0030].



Schiavone, Fig. 6 (annotated).

180. Moreover, Schiavone describes that its methods of rule-based processing of electronic mail messages involves rules described in accordance with “a compliance policy” because Schiavone teaches that “the rule is obtained from a

data store of rules 155” and “[a] common set of rules [that] is used by two or more users of the system. The rule may be associated with messages by default, in an automated fashion, or as the result of a sender’s specification.” Schiavone, [0007] and [0030]. As such, Schiavone’s rule-based processing “provides for widespread usage of a common, standard set of rules.” Schiavone, [0035].

181. Schiavone’s “rule base,” *e.g.*, rule base 155 in Fig. 6 (“*a compliance policy describing rules*”) may be on a client device or a server. Schiavone states that the “rule base” can be “stored in a memory of the recipient’s communications device ... [and] ... periodically updated by a third party trusted intermediary responsible for **defining rules**, maintaining and updating rules, approving proposed rules, and distributing such rules to senders and recipients known to the trusted intermediary” or, that the “rule base” “may also or alternatively be stored at the commercial mailer, trusted authority and/or another third party and referenced as necessary via a communications network.” Schiavone, [0030], [0066]. See also, Schiavone, [0029] (describing an example of a “COPPA” rule that complies with U.S. laws). Further, in view of Schiavone’s afore-mentioned teachings of “defining rules,” Schiavone also discloses “*defining a compliance policy describing rules.*” Therefore, Schiavone teaches providing rules described in accordance with “*a compliance policy.*”

182. Additionally, claim 1 expressly requires that the “*messaging server*” includes the “*governance module*.” (See elements [1.0], [1.4] in Claim Listing.) Schiavone’s Fig. 6 embodiment describes several engines and data stores such as the “recipient compliance engine,” “transaction data,” “rule base,” “public preference data,” “private profile data,” and “private commerce data” (e.g., collectively the “*governance module*”) as a client-side module analogous to the 911 Patent’s description of “**the client module 510 include[s] a governance module 516 for storing one or more governance policies** received from the messaging server 112.” Ex.1001, 14:44-46. See also Schiavone, [0064] (disclosing “the recipient’s communications device 150 store[ing] and run[ning] ... a recipient compliance engine 154”).

183. And it would have been obvious to a POSITA that the engines and data stores of Schiavone can be implemented on a server (such as the servers in Rohall’s system) based on Schiavone’s disclosures that (1) “recipient compliance engine” may reside at “the commercial mailer, trusted authority and/or another third party,” instead of at a client device, and (2) “**mail server software ... is well known in the art.**” Schiavone, [0066], [0064], [0030], emphasis added. Schiavone also explains that in “an alternate embodiment **the rule base 155 may also or alternatively be stored at the commercial mailer, trusted authority and/or another third party** and referenced as necessary via a communications network,” and “[i]n another

alternate embodiment, **a recipient compliance engine may reside at** commercial mailer, **trusted authority and/or another third party.**” Schiavone, [0066] (emphases added). Schiavone defines “**a third party trusted intermediary**” as a “**central source**” that is “responsible for defining rules, maintaining and updating rules, approving proposed rules, and distributing such rules to senders and recipients known to the trusted intermediary.” Schiavone, [0066], [0030]. From these teachings, a POSITA would have recognized that there are only two possible ways of implementing engines and data stores: at the server-side or at the client-side, and moreover, both implementations were known in the prior art. Schiavone, [0066].

184. Accordingly, a POSITA would have understood that server-side and client-side implementations of Schiavone’s engines and data stores represent a finite number of identified, predictable solutions to the design need of implementing engines and data stores. Thus, it would have been obvious to a POSITA that Rohall’s messaging system server, as modified by Schiavone’s teachings of rule-based processing of electronic mail messages would additionally include Schiavone’s teachings of “recipient compliance engine,” “transaction data,” “rule base,” “public preference data,” “private profile data,” and “private commerce data” collectively serving the claimed functionality of the “*governance module*”.

185. Therefore, Schiavone discloses “*a governance module adapted to control a governance policy database storing a compliance policy describing rules.*”

186. In addition, Schiavone discloses that “*rules [stored in the governance policy database]*” are applicable “*to the sent submessages.*” For example, Schiavone states that “**the electronic mail message is processed based upon reference to a rule applicable to that particular message.** The rule includes a conditional instruction conditioned upon a result of a logic or other function of recipient profile data.” Schiavone, Abstract (emphasis added). Schiavone further explains “[t]he **rule may be associated with messages by default, in an automated fashion, or as the result of a sender’s specification, e.g., by including a mail-type specifier in the message.**” Schiavone, [0007] (emphasis added). Schiavone also discusses its process flow, stating that “an incoming **electronic mail message is received**, and a message type specifier associated with the electronic mail message is identified,” and “[a]s shown at step 28 in FIG. 2, **a rule associated with the message type specifier is identified.**” Schiavone, [0028]–[0030] (emphases added). Indeed, Schiavone is directed to “a method and apparatus for rule-based processing of electronic mail messages, e.g. to negotiate transactions, enable automated verification of compliance with privacy, legal, regulatory and/or best practices standards, ...” Schiavone, [0001]. Furthermore, and as described above (see §VII.A.1), a POSITA would have been motivated to modify Rohall’s messaging system architecture to additionally include Schiavone’s above-mentioned teachings.

187. Therefore, Rohall/Schiavone render obvious elements [1.4], [10.4], [13.1].

f. Governance module determines rules described by Compliance Policy


[1.5] the governance module further adapted to determine rules described by the compliance policy that are applicable to ones of the plurality of sent submessages,

[10.5] the governance module further adapted to determine rules described by the compliance policy that are applicable to ones of the plurality of sent submessages,

188. Elements [1.5], [10.5] generally relate to “the governance module ... determin[ing] rules described by the compliance policy that are applicable to ones of the plurality of sent submessages.”

189. Although Rohall does not expressly disclose these elements, as explained earlier in connection with element [1.4], Schiavone teaches “*a governance module*” and for the reasons below, elements [1.5], [10.5] are taught by Rohall as modified by the teachings of Schiavone (see §X.B).

190. For example, Figure 2 of Schiavone (shown below) describes receiving an incoming message and identifying a rule (“*determin[ing] rules*”) “obtained from a data store of rules 155” (“*described by the compliance policy*”). Schiavone, [0010], [0030], Fig. 2 (blocks 24 and 28).



Ex. 1006, Fig. 2 (annotated).

191. Schiavone provides examples of how “to determine rules described by the compliance policy that are applicable to ones of the plurality of sent submessages.” For instance, “a sender could select several different rule sets to

automate their compliance with selected US federal and local laws and regulations and, if desired, varying laws and regulations enacted by the European Union and elsewhere globally” and that “[t]he rule controls how the e-mail message will be processed.” Schiavone, [0029]-[0030]. In one example, “consider that the rule for COPPA ... and that the mail message should be delivered to a U.S. recipient only if the recipient is over the age of 13.”¹⁰ *Id.*, [0031]. If “the recipient lives in the US and is 45 years of age,” then “it is determined that the recipient’s profile data is in compliance, i.e., compatible, with the applicable rule in step 32 and the electronic mail message is delivered to the recipient, e.g. placed in the recipient’s electronic “inbox” and/or displayed to the recipient, as shown at step 34.” *Id.*, [0033].

192. Furthermore, and as described above (*see* §X.B), a POSITA would have been motivated to modify Rohall’s messaging system architecture to additionally include Schiavone’s above-mentioned teachings.

193. Therefore, Rohall/Schiavone render obvious elements [1.5], [10.5].

¹⁰ COPPA refers to the “Children’s Online Privacy Protection Act, a law applicable in the United States (“US”) which forbids delivery of commercial electronic mail communications to children aged thirteen (13) and under unless a parent has given prior consent.” *Id.*, [0029].

g. Rules of the compliance policy selectively apply to submessages

[1.6] wherein different rules of the compliance policy are applicable to different ones of the plurality of sent submessages;

[10.6] wherein different rules of the compliance policy are applicable to different ones of the plurality of sent submessages;

[13.5] selectively applying rules of the compliance policy to the sent submessages

194. Elements [1.6], [10.6], [13.5] generally relate to “different rules of the compliance policy ... applicable to different ones of the plurality of sent submessages.”

195. Although Rohall does not expressly disclose these elements, they are taught by Rohall, as modified by Schiavone (*see* §X.B) because, as explained below, Schiavone teaches an example in which email messages sent to US recipients are delivered to recipients in accordance with laws for US residents, whereas email messages sent to recipients in France are delivered to recipients in accordance with laws for French residents.

196. First, as explained in connection with element [1.4], Schiavone teaches providing rules described in accordance with “*a compliance policy*.” Furthermore, Schiavone provides an example describing how different rules that apply to different messages:

In the example of FIG. 2, consider that the recipient's profile data indicates that the recipient lives in the US and is 45 years of age. Accordingly, it is determined that the recipient's profile data is in compliance, i.e., compatible, with the applicable rule in step 32 and the electronic mail message is delivered to the recipient, e.g. placed in the recipient's electronic "inbox" and/or displayed to the recipient, as shown at step 34. The transaction then ends, as shown at step 37. For example, the same message sent to a 10 year old in France could result in a determination of compliance and delivery of the message because the rule (and COPPA law) is applicable only to residents of the US. Compliance with local laws is thereby ensured.

Schiavone, [0033]. Accordingly, a sender could select **several different rule sets to automate their compliance with selected US federal and local laws and regulations and, if desired, varying laws and regulations enacted by the European Union and elsewhere globally.**" *Id.*, [0029] (emphasis added).

197. Additionally, Schiavone teaches different rules that apply to different messages, stating: "The **rules** include logic for processing of incoming **messages**," and a "common set of rules is used by two or more users of the system." *Id.*, [0007] (emphases added). Further, "rules may be implemented for a wide variety of applications, **including compliance with applicable laws**," and "[a]ccordingly, a mail message can be automatically processed and consequently filtered or routed based on a patchwork of applicable laws, regulations and best practices applicable to the specific sender, receiver and content of the e-mail message." *Id.*, [0016].

198. As another example, Schiavone teaches applying rules to messages to customize their content, stating that “an electronic mail message may be modified before it is sent, or before it is received, e.g., by adding or deleting certain content, by modifying the message’s color, length, etc., or by assigning values to variables within the mail message.....Accordingly, recipients receive less unwanted mail (e.g. spam) and the mail received is more interesting to the recipient because it is customized to the recipient’s preferences.” Schiavone, [0017]. Thus, for example, by customizing or otherwise modifying certain email messages, Schiavone teaches “*different rules of the compliance policy are applicable to different ones of the plurality of sent submessages.*” Furthermore, and as described above (see §X.B), a POSITA would have been motivated to modify Rohall’s messaging system architecture to additionally include Schiavone’s above-mentioned teachings.

199. Therefore, the Rohall/Schiavone combination renders obvious elements [1.6], [10.6], [13.5].

h. Entities in a Population Structure

[1.7] wherein the messaging system is utilized by a set of entities in a population structure and

[10.7] wherein the messaging system is utilized by a set of entities in a population structure and

**[13.6] wherein the messaging system is utilized by a set of entities
in a population structure and**

200. Elements [1.7], [10.7], [13.6] generally relate to “*the messaging system ... utilized by a set of entities in a population structure.*” Rohall discloses these elements.

201. Initially, the 911 Patent describes “*entities*” and “*population structure*” as follows: “[T]he populations represent the universe of end-users and/or other entities that use the messaging system.” Ex.1001, 8:14–15. Further, “the compliance policy population structure associates three types of attributes with an end-user,” including (1) domain and group to which the end-user belongs, where “all end-users assigned to the same domain to have the same properties”; (2) “allowable actions” that “define the activities that the end-user can participate in using the messaging system”; and (3) “attributes that relate to the specific end-user... such as the end-user’s line of reporting, human resources advisor, time spent with the company, etc.” *Id.*, 8:20–9:2.

202. Rohall describes an electronic mail system (“*messaging system*”) is “*utilized by*” “users/people” (“*a set of entities*”) “in a group or organization” (“*a population structure*”). Rohall identifies that:

Electronic mail has become one of the most widely used business productivity application. ... A problem facing people in organizations is persistence. Too often projects

get started only to lose momentum because of the effort required to track actions, activity and progress. ... This time depletes the time spent performing productive work. In addition, many current electronic mail applications do not interact with other productivity tools of the user, such as calendar programs and collaborative meeting applications.

Rohall, [0006]-[0007]. Accordingly, Rohall teaches “an improved inbox or viewer for electronic mail that allows for greater integration of functions to enhance usability and productivity.” *Id.*, [0012]. Rohall “recognizes that there are specific domains in which identifying features, such as names, dates, and company names, product names, becomes useful,” and that “commercially-available feature extraction software extracts relevant features in documents, including names, numbers, and names of organizations and products.” *Id.*, [0084].

203. Rohall discloses several examples of how its “*messaging system*” is “*utilized by a set of entities in a population structure.*” For example, Rohall states that “Mail Agent 230 first detects the occurrence of a triggering event as illustrated by decisional step 500,” where “[s]uch event may include the sending or receipt of an electronic message, or, alternatively a request to delete an electronic message.” *Id.*, [0060]. Additionally, “the user may designate specific electronic mail types by name, keywords or senders as ‘Urgent’ and have a folder or Notes mail database created into which these urgent documents may be stored until accessed.” *Id.*, [0107]. In another example, “Mail Agent 230 performs the functionality necessary

to dynamically generate conversation threads, store the conversation thread representations and forward the same to other electronic mail users.” *Id.*, [0126].

204. Therefore, Rohall discloses elements [1.7], [10.7], [13.6].

i. Rule Elements of Compliance policy

[1.8] wherein a rule of the compliance policy includes one or more rule elements describing whether the rule is applicable to a sent submessage,

[10.8] wherein a rule of the compliance policy includes one or more rule elements describing whether the rule is applicable to a sent submessage,

[13.7] wherein a rule of the compliance policy includes one or more rule elements describing whether the rule is applicable to a sent submessage,


205. Elements [1.8], [10.8], [13.7] generally relate to “a rule of the compliance policy include[ing] one or more rule elements describing whether the rule is applicable to a sent submessage.”

206. Although Rohall does not expressly disclose these elements, they are taught by Rohall, as modified by Schiavone (*see* §X.B) because, as explained below, Schiavone teaches an example in which a recipient’s age is used as a rule element for describing whether the COPPA rule is applicable to a message.

207. First, as explained in connection with element [1.4], Schiavone teaches providing rules described in accordance with “*a compliance policy.*”

208. In addition, Schiavone discloses “*a rule of the compliance policy include[ing] one or more rule elements describing whether the rule is applicable to a sent submessage.*” Specifically, Schiavone teaches an example in which US-based, COPPA laws is applicable to an email if the recipient is under the age of 13, whereas it would be inapplicable (i.e., does not apply to) an email to a recipient that is 45 years of age. Schiavone, [0031], [0033].

209. As another example, Figure 2 of Schiavone teaches “*rule elements describing whether the rule is applicable to a sent submessage*” based on the “recipient profile data” (“one or more *rule elements*”) “*describing whether the rule*” is “in compliance with” (“*applicable to*”) “the electronic mail message” (“*a sent submessage*”).



Schiavone, Fig. 2; See also Schiavone, [0032]-[0033]. “If the recipient’s profile data is determined not to be in compliance with the applicable rule in step 32, then the electronic mail message is rejected.” Schiavone, [0034]. Further, Schiavone also states that **“a mail message can be automatically processed and consequently filtered or routed based on a patchwork of applicable laws, regulations and best practices applicable to the specific sender, receiver and content of the e-mail message,”** e.g., “a message may be rejected, and not delivered to the intended recipient, if a recipient’s profile data is incompatible with applicable laws for receiving the message, as specified in a rule associated with the message.”

Schiavone, [0016] (emphasis added). *See also* example of COPPA rule including age-related restrictions (“one or more *rule elements*”). Schiavone, [0031]-[0034].

210. Therefore, the Rohall/Schiavone combination renders obvious these claim elements.

j. Rule Elements selected from Source/Target/Rank

[1.9] at least one of the rule elements selected from the group consisting of: a source element describing a population structure entity that is a source of the sent submessage; a target element describing a population structure entity that is a recipient of the sent submessage; and a rank describing a ranking of the rule relative to other rules in the compliance policy.

[10.9] at least one of the elements selected from the group consisting of: a source element describing a population structure entity that is a source of the sent submessage; a target element describing a population structure entity that is a recipient of the sent submessage; and a rank describing a ranking of the rule relative to other rules in the compliance policy.

[13.8] at least one of the elements selected from the group consisting of: a source element describing a population structure entity that is a source of the sent submessage; a target element describing a population structure entity that is a recipient of the sent submessage; and a rank describing a ranking of the rule relative to other rules in the compliance policy.

211. The Rohall/Schiavone combination renders these elements obvious.

212. Although Rohall does not expressly disclose these elements, they are taught by Rohall, as modified by Schiavone (*see* §X.B) because Schiavone teaches “*a source element describing a population structure entity that is a source of the sent submessage [and] a target element describing a population structure entity that is a recipient of the sent submessage.*”

213. For example, Schiavone states that “rules may be implemented for a wide variety of applications, including compliance with applicable laws,” where “a mail message can be automatically processed and consequently filtered or routed based on ... best practices **applicable to the specific sender, receiver** and content of the e-mail message.” Schiavone, [0016] (emphasis added). Thus, in the above-mentioned disclosures, the “sender” (“*a source element*”) is a “*rule element describing*” “users/people in a group or organization” (“*a population structure entity*”) “*that is a source of*” “an electronic mail message” (“*the sent submessage*”). And similarly, the “recipient” corresponds to the “*target element.*” *See also* Schiavone, [0016] (describing that “a message may be rejected, and not delivered to the intended recipient, **if a recipient’s profile data is incompatible** with applicable laws for receiving the message, as specified in a rule associated with the message”). Schiavone also teaches that “the rule applies only to recipients living in the US.” Schiavone, [0031].

214. In addition, Schiavone’s disclosure of “*source element*” and “*target element*” are consistent with the description of “source” as “the sender of a submessage” and a “target” as “the recipient of a submessage” in the 911 Patent. *See generally* Ex. 1001, 9:45–12:67.

215. Therefore, the Rohall/Schiavone combination renders these elements obvious.

2. Dependent Claims

a. Jobcode

[2.0] The messaging server of claim 1, wherein the compliance policy describes a population structure of end-users of the messaging system, a set of jobcodes with which the sent submessages can be classified, and a set of rules applicable to the population structure, jobcodes¹¹, and sent submessages¹².

216. The combination of Rohall and Schiavone renders this element obvious.

¹¹ It appears that the use of “*jobcodes*” here instead of “*a set of jobcodes*” is a typographical error.

¹² It appears that the use of “*sent submessages*” here instead of “*the sent submessages,*” as used in the earlier instance in this claim, is a typographical error.

217. Claim 2 requires that the “*compliance policy describes*”:

- (1) “a population structure of end-users of the messaging system,”
- (2) “a set of jobcodes with which the sent submessages can be classified,” and
- (3) “a set of rules applicable to the population structure, jobcodes, and sent submessages.”

218. As discussed above in connection with elements [1.7], [1.4], and [1.5], Rohall’s “messaging system,” as modified by Schiavone’s teachings of rule-based processing of electronic mail messages (*see* §X.B), renders obvious an “a group or organization” (“*a population structure*”) “*of*” “users/people” (“*end-users*”) “of the” “electronic mail system” (“*messaging server*”) and “*rules described by the compliance policy applicable to the ... sent submessages.*”

219. First, the combination of Rohall and Schiavone renders obvious that “*the compliance policy describes a population structure of end-users of the messaging system.*” As explained earlier, the Rohall/Schiavone combined system renders obvious “a group or organization” (“*a population structure*”) and “*rules described by the compliance policy applicable to the ... sent submessages*” and additionally, because the “*compliance policy*” in the Rohall/Schiavone combined system is defined in accordance with the group/organization where such a policy applies, it would have been obvious to a POSITA that “the compliance policy” applies to, or otherwise “*describes [the] population structure.*”

220. Second, the Rohall/Schiavone combined system renders obvious that “*the compliance policy*” “*describes ... a set of jobcodes with which the sent submessages can be classified.*” The 911 Patent explains that “jobcodes describe which rules are applicable in which circumstances, and can represent exceptions to the standard rules,” and that jobcodes “allow submessages to be tied to external business processes,” e.g., “jobcodes allow messages to be associated with a business deal such as a merger, with a patient at a hospital, with particular clients of a law firm.” Ex.1001, 8:17–19, 12:35–40. The 911 Patent does not appear to define what “classify” means and gives only one example where the governance module scans the message database at a regular interval to “to classify the message components and other entities and delete entities that have reached the end of their retention periods.” *Id.*, 13:65–14:2.

221. Like the 911 Patent’s description of “*jobcodes*” as “allow[ing] messages to be identified with specificity such as “a business deal such as a merger, with a patient at a hospital, [and] with particular clients of a law firm,” analogously, Schiavone discloses use of recipients’ “personal/demographics information,” which can include private profile information such as “name, address, telephone number, age, sex, race and social security number,” for identifying “a specific transaction, specific sender, or specific message.” Furthermore, Schiavone teaches that use of privacy information such as “personal/demographics information” of recipients

provide “enhanced preference-based filtering, targeting and/or compliance” of the messaging system/platform, and hence, “personal/demographics information” of recipients constitute an example of “*a set of jobcodes*,” e.g., overlaid on a primary file record (such as an e-mail address) and used in a manner such that specific messages sent to recipients (such as those sent using Rohall’s system and including messages in a conversation thread represented using children documents or shadows thereof) (“*the sent submessages*”) “*can be*” “filtered” (“*classified*”). Schiavone, [0004], [0032]. Therefore, Schiavone’s teachings related to “personal/demographics information” of recipients serve as “*a set of jobcodes with which the sent submessages can be*” filtered (“*classified*”).

222. Third, the combination of Rohall and Schiavone renders obvious “*a set of rules applicable to the population structure, jobcodes, and sent submessages*.” As explained earlier, in the Rohall/Schiavone combined system, “*the compliance policy*,” “*describes a population structure*,” “*describes rules applicable to the sent submessages*,” and further “*describes ... a set of jobcodes with which the sent submessages can be classified*,” and therefore, the Rohall/Schiavone combination renders obvious “*a set of rules applicable to the population structure, jobcodes, and sent submessages*.”

223. Therefore, the Rohall/Schiavone combination renders this element obvious.

b. Jobcode Rules applicable to all sent submessages

[3.0] The messaging server of claim 2, wherein a jobcode describes a set of jobcode rules applicable to all sent submessages classified with the jobcode.

224. The Rohall/Schiavone combination renders this element obvious.

225. As discussed above in connection with claim 2, Rohall’s “messaging system,” combined with Schiavone’s teachings of rule-based processing of electronic mail messages (see §X.B), renders obvious that “*the compliance policy*” “*describes ... a set of jobcodes with which the sent submessages can be classified, and a set of rules applicable to the ... jobcodes, and sent submessages.*” And, additionally, by teaching use of “personal/demographics information” (“*jobcodes*”) for “providing enhanced preference-based filtering, targeting and/or compliance,” which a POSITA would have understood as examples of rule-based processing of electronic mail messages, Schiavone’s use of “personal/demographics information” (“*jobcodes*”) constitute as “*describe[ing] a set of jobcode rules applicable to ... sent submessages classified with the jobcode*” Schiavone, [0032]. **“This invention relates generally to ... rule-based processing of electronic mail messages, e.g. to negotiate transactions, enable automated verification of compliance with privacy, legal, regulatory and/or best practices standards, to perform spam**

filtering, anonymous profiling, target marketing through dynamic offers, and prioritization and standardization of e-mail to transactions.” Schiavone, [0001].

226. And moreover, Schiavone does not teach any restrictions on the number of sent submessages to which its teachings of rule-based processing of electronic mail messages apply. Rather, Schiavone teaches using its rule-based processing of electronic messages for enabling “nearly limitless transaction types” and further describes that “rules can be added, deleted and/or revised easily without a need for revisions to the data stored in the private profiles of individual users.” Schiavone, [0035]. Therefore, in the Rohall/Schiavone combined system, a POSITA would have found it obvious that “[*the*] set of jobcode rules” is “*applicable*” nearly limitlessly “*to all sent submessages*” that are “*classified with the jobcode.*”

227. Therefore, the Rohall/Schiavone combination renders this element obvious.

c. Associate a jobcode with a sent submessage

[4.0] The messaging server of claim 2, wherein the messaging module is further adapted to receive an indication of a jobcode and of a sent submessage from an end-user, and is further adapted to associate the jobcode with the sent submessage responsive to receipt of the indication of the jobcode from the end-user.

228. The Rohall/Schiavone combination renders this element obvious.

229. As explained above in connection with elements [1.1] and [2.0], Rohall discloses “*the messaging module*” and Schiavone’s use of “personal/demographics information” (such as a name and a zip code) of recipients serve as “*a set of jobcodes.*”

230. Moreover, Rohall discloses “*the messaging module ... receiv[ing] an indication of ... a sent submessage from an end-user.*” Rohall’s invention is directed to a method comprising “detecting the occurrence of an triggering event related to an electronic mail document.” Rohall, [0018]. “Mail Agent 230 first detects the occurrence of a triggering event ... [which] may include the sending or receipt of an electronic message.” Rohall, [0060].

231. Schiavone explains that conventional techniques “do not contain a significant amount of, if any, personal/demographics information (often, they contain only an e-mail address, not even a name or zip code).” Schiavone, [0004]. Schiavone further explains that “[a]lternatively, a list broker or other intermediary has additional demographic information available regarding all or a portion of the people whose e-mails are contained on its lists.” Schiavone, [0004]. And therefore, this additional demographic information, for example included in a recipient’s private profile may include “information such as name, address, telephone number, age, sex, race and social security number. This private and public preference data

may be compiled, for example, during the recipient's completion of a questionnaire." Schiavone, [0022].

232. From these teachings, in the combined Rohall/Schiavone system, the recipient's answers providing his/her personal/demographic information in the questionnaire of Schiavone would enable Rohall's Mail Agent module ("*messaging module*") "*receiv[ing] an indication of a jobcode ... from an end-user.*" Furthermore, in the combined Rohall/Schiavone system, Rohall's teachings of "detect[ing] the occurrence of a triggering event" such as "the sending or receipt of an electronic message" for a recipient (Ex.1005, [0060]), e.g., whose personal/demographic information is stored in a recipient's private profile, as taught in Schiavone, would result in Rohall's Mail Agent module ("*messaging module*") "*associate[ing] the*" "personal/demographic information" ("*jobcode*") *with the*" "sending of a message" ("*sent submessage*") "responsive to receipt of" the recipient's answers providing his/her personal/demographic information ("*the indication of the jobcode from the end-user*").

233. Therefore, the Rohall/Schiavone combination renders this element obvious.

d. Communication and Readership Rules

[5.0] The messaging server of claim 1, wherein the rules described by the compliance policy include one or more types of rules from

the set consisting of: communication rules describing which population structure entities can communicate with other population structure entities via the messaging system; and readership rules describing which population structure entities can read messages from other population structure entities via the messaging system.

234. The Rohall/Schiavone combination renders this element obvious.

235. The 911 Patent explains that “[c]ommunication rules determine which end-users, domains, and groups are allowed to communicate with other end-users, domains, and groups” and “[r]eadership rules determine which end-users, domains, and groups are allowed to read content written by other end-users, domains, and groups when such content is forwarded to them in some manner.” Ex.1001, 9:37-39, 10:10-13. The 911 patent specification makes clear that “message” refers to “a data communication sent by one end-user to one or more end-users of the messaging system or another messaging system,” and can include different types of media, such as “emails, Short Message Service (SMS) messages, Instant Messages (IMs), Multi-Media Message (MMS) and/or other types of messages.... [and] media files, such as discrete and/or streaming audio and/or video, still images, etc.” Ex.1001, 3:4-16.

236. As discussed above in connection with elements [1.7], [1.4], and [1.5], Rohall’s “messaging system” and “users in a group or organization” (“*population*

structure entities”), as modified by Schiavone’s teachings of “*rules described by the compliance policy*.”

237. Additionally, Schiavone provides an example of a company-developed rule which “may require internal approval (e.g., by clicking an appropriate button or typing OK, etc.) from a purchasing agent, a supervisor, an accountant, and a vice-president before a purchase order is forwarded to a vendor,” “although the mail message is prepared by the purchasing agent and addressed to the vendor, the electronic mail message will be processed according to the rule and forwarded to each individual for approval before the message is able to be sent to the vendor” and therefore serves as an example of a “*communication rule*” because it “*describe[es] which population structure entities can communicate with other population structure entities via the messaging system*.” Schiavone, [0036]. A POSITA would have found it obvious that Rohall’s electronic messaging system, which is utilized by “users/people” (“*the set of entities*”) in “a group or organization” (“*a population structure*”) would be additionally modified to include Schiavone’s teachings of “a company-developed rule” (“*a “communication rule”*”) “*describing*” “specific personnel holding positions of rank and involved in approving a purchase order” (“*which population structure entities can communicate with other population structure entities*”) “*via*” “Rohall’s electronic messaging system” (“*the messaging system*”).

238. In addition, Schiavone teaches that its “invention can be used for filtering, ... by allowing receipt of a message only if certain laws or standards are complied with.” Schiavone, [0007]. For example, as discussed above, Schiavone discloses that the rule, COPPA, “forbids delivery of commercial electronic mail communications to [a group comprising] children aged thirteen (13) and under unless a parent has given prior consent.” Schiavone, [0029]. This rule therefore meets the “*readership rule*” element because it describes “recipients under the age of thirteen having parental consent” (“*which population structure entities*”) “*can read messages from*” “commercial mailers” (“*other population structure entities*”) “*via the messaging system*” (e.g., Rohall’s messaging system).

239. Schiavone also explains that “[f]or a web browser request having an associated rule (e.g., a URL pointing to an adult content website may be associated with an “OVER 21 YEARS OF AGE” rule), and the rule will be checked against the outgoing browser request and blocked if the recipient profile data is not in compliance with the rule.” *Id.*, [0042]. Given that the 911 Patent explains that a “message” can include various types of media, this exemplary rule of Schiavone meets the claimed “*readership rule*” of the 911 Patent because any end-user under 21 years of age cannot read messages from other population structure entities via the messaging system.

240. Therefore, the Rohall/Schiavone combination renders this element obvious.

e. Validity and Retention Rules

[6.0] The messaging server of claim 1, wherein the rules described by the compliance policy include one or more types of rules from the set consisting of: validity rules describing a time period for which a sent submessage or entity in the population structure is useable; and retention rules describing a time period after which a sent submessage or entity in the population structure should be deleted from the messaging system

241. The Rohall/Schiavone combination renders this element obvious.

242. As discussed above in connection with elements [1.7], [1.4], and [1.5], Rohall's "messaging system," as modified by Schiavone's teachings of rule-based processing of electronic mail messages (*see §X.B*) renders obvious "a group or organization" ("*population structure*") and "*rules described by the compliance policy.*" Schiavone teaches "*validity rules describing a time period for which a sent submessage ... is useable*" and Schiavone also teaches "*retention rules describing a time period after which a sent submessage . . . should be deleted from the messaging system.*"

243. Additionally, Schiavone teaches "*validity rules describing a time period for which a sent submessage or entity in the population structure is useable.*"

Schiavone states that “a recipient’s profile data may contain time-based limitations for receiving electronic mail, which may be useful for filtering purposes.” Schiavone, [0067]. For example, “a recipient may include in the recipient profile data that for a period of two months, e.g. after purchasing a house, the recipient wishes to receive sale offers relating to purchases of new homes, or home furnishings, etc.,” and “[t]he rule can be used to reject or discard the electronic mail message, e.g. before viewing by a recipient, if the appropriate information is not contained in the recipient’s profile data.” *Id.* Under this rule, any submessage dated more than two months after the recipient purchases a house, the submessage is rejected, i.e., unusable.

244. As another example, Schiavone teaches “work flow processing transactions, e.g., by providing a rule requiring approval and/or handling of a given electronic mail message by more than one recipient and/or according to a **predetermined deadline schedule.**” Schiavone, [0036] (emphasis added). A POSITA would have understood from this teaching that the approval message is usable only within a certain time period, and therefore, the approval rule meets the “*validity rules describing a time period for which a sent submessage ... is useable.*”

245. Therefore, the Rohall/Schiavone combination renders this element obvious.

f. Prior compliance policy active at a prior time

[7.0] The messaging server of claim 1, wherein the governance policy database further stores at least one prior compliance policy active at a time prior to a time at which the compliance policy was active, and wherein the sent submessages of the stored messages are associated with a compliance policy that was active at a time that content of the sent submessages was created.

246. The Rohall/Schiavone combination renders this element obvious.

247. As discussed above in connection with elements [1.4], Rohall’s “messaging system,” as modified by Schiavone’s teachings of rule-based processing of electronic mail messages (see §X.B) renders obvious “*a governance policy database storing a compliance policy ... applicable to the sent submessages.*”

248. Moreover, Schiavone teaches providing “at least one prior compliance policy active at a time prior to a time at which the compliance policy was active” and further that “wherein the sent submessages of the stored messages are associated with a compliance policy that was active at a time that content of the sent submessages was created.” Because the “compliance policy [is] applicable to the sent submessages,” in the Rohall/Schiavone combined system, a POSITA would have understood that “*the sent submessages ... are associated with a compliance policy.*” Furthermore, as described above in connection with element [1.2], in Rohall, “children documents or shadows thereof “*constitute “a plurality of sent*

submessages,” and therefore, children documents (or, their shadows) are “*of the*” original messages (“*the stored messages*”).

249. Schiavone describes an example in which “a recipient’s profile data may contain time-based limitations for receiving electronic mail, which may be useful for filtering purposes. For example, a recipient may include in the recipient profile data that for a period of two months, e.g. after purchasing a house, the recipient wishes to receive sale offers relating to purchases of new homes, or home furnishings, etc. The rule can be used to reject or discard the electronic mail message, e.g. before viewing by a recipient, if the appropriate information is not contained in the recipient’s profile data.” Schiavone, [0067]. A POSITA would have understood from this example that the recipient’s profile data indicates “*a time that content of the sent submessages was created,*” and that the “*prior compliance policy*” (active for a 2-month period after purchasing the house) is “*a compliance policy that was active at a time that content of the sent submessages was created*” and further that the “*prior compliance policy*” was “*active at a time prior to a time at which the compliance policy was active.*”

250. Therefore, the Rohall/Schiavone combination renders this element obvious.

g. Rules are associated with submessages

[8.0] The messaging server of claim 7, wherein at least one of the messages stored in the message database has a first sent submessage associated with a first compliance policy and a second sent submessage associated with a second compliance policy different from the first compliance policy, and wherein the governance module applies the rules described by the first compliance policy to the first sent submessage and applies the rules described by the second compliance policy to the second sent submessage.

251. The Rohall/Schiavone combination renders this element obvious.

252. As discussed above in connection with elements [1.4], [1.5], and [1.6], Rohall's "messaging system," as modified by Schiavone's teachings of rule-based processing of electronic mail messages (see §X.B), renders obvious "*a governance module adapted to control a governance policy database storing a compliance policy describing rules applicable to the sent submessages,*" "*the governance module further adapted to determine rules described by the compliance policy that are applicable to ones of the plurality of sent submessages*" and further that "*different rules of the compliance policy are applicable to different ones of the plurality of sent submessages,*" which therefore suggests or teaches that "*the governance module applies the rules described by the [first / second] compliance policy to the [first / second] sent submessage.*" Specifically, as explained in [1.4], Schiavone's "rule base" corresponds to "*a compliance policy describing rules.*" A

POSITA would have understood that an organization or a group can have more than one compliance policy and therefore “rule base 155” can store one or more “*compliance polic[ies] describing rules*,” which suggests or teaches “*a first compliance policy*” and “*a second compliance policy*.” See Arnold, [0028], [0030] (disclosing “explicitly defined business policies” within an enterprise and “needs to accommodate different policies in different environments”).

253. In addition, Schiavone teaches that “[a] **rule may be associated with messages** by default, in an automated fashion, or **as the result of a sender’s specification**, e.g., by including a mail-type specifier in the message.” Schiavone, [0007]. A POSITA would have understood from this disclosure that a sender’s specification (e.g., included as part of a group or an organization’s compliance policy) of his/her sent message (e.g., a reply or forward) can be different from one sender to another, and therefore, would result in different compliance policies. And, consequently, in the Rohall/Schiavone combination, “*a [first / second] sent submessage [is] associated with a [first / second] compliance policy*” and further “*the rules described by the [first / second] compliance policy [apply] to the [first / second] sent submessage*.”

254. Therefore, the Rohall/Schiavone combination renders this element obvious.

h. Communication, Readership, Action, and Search Rules

[11.0] The computer program product of claim 10, wherein the rules described by the compliance policy include one or more types of rules from the set consisting of: communication rules describing which population structure entities can communicate with other population structure entities via the messaging system; readership rules describing which population structure entities can read messages from other population structure entities via the messaging system; action rules describing actions a population structure entity can perform on a message component in the -messaging system; and search rules describing which population structure entities can view which message components when searching the messaging components stored by the messaging system.

255. *See* element [5.0].

i. Validity and Retention Rules

[12.0] The computer program product of claim 10, wherein the rules described by the compliance policy include one or more types of rules from the set consisting of: validity rules describing a time period for which a sent submessage or entity in the population structure is useable; and retention rules describing whether a sent submessage or entity in the population structure should be deleted from the messaging system.

256. *See* element [6.0].

j. Group and Rules of Entities

[15.0] The computer-implemented method of claim 13, wherein the compliance policy describes the set of entities in the population structure, the method further comprising: defining a group containing a set of entities, the entities¹³ in the group defined by their locations¹⁴ in a hierarchy defined by the hierarchical population structure¹⁵; and defining a set of rules applicable to the entities in the group.

257. The combination of Rohall and Schiavone combination renders this element obvious.

258. Initially, the 911 Patent states that “[t]he first type of end-user attribute specifies the domain and groups to which the end-user belongs.” Ex.1001, 8:20-22. And “[u]nlike the domains, groups are non-mutually exclusive sets of end-users that

¹³ It appears that the use of “*the entities*” here instead of “*the set of entities*” is a typographical error.

¹⁴ For my analysis, I have understood the term “*their locations*” to mean “*locations of the set of entities.*”

¹⁵ The term “**the** hierarchical population structure” first appears in dependent claim 15. There appears to be an antecedent basis issue in usage of “**the** hierarchical population structure.”

can span multiple domains, [and] [t]he use of groups allows users in different domains to be logically related.” *Id.*, 8:36-40. Examples of groups include geography, seniority, and project based. *Id.*, 8:42–52. The 911 Patent explains that “a population structure entity is specified relative to a position in the domain (or another) hierarchy. For example, a group can be defined that includes all end-users above the rank of vice-president or all end-users that are under the management of a particular manager.” Ex.1001, 9:7-12.

259. As explained above in connection with element [1.7], Rohall discloses “*the set of entities in the population structure.*” Furthermore, as described in the 911 Patent, which appears to refer to “locations of the set of entities in the group” as a position in a domain hierarchy, such as a position within a company, Schiavone teaches “*the entities in the group defined by their locations in a hierarchy defined by the hierarchical population structure.*” Specifically, Schiavone describes an example involving an email requiring approval by the “specific personnel holding positions of rank and involved in approving a purchase order” (“*group defined by their locations in a hierarchy*”) “*defined by the*” organizational chain-of-command (“*hierarchical population structure*”) before it is sent to the vendor. Schiavone explains: “a company-developed rule may require internal approval (e.g., by clicking an appropriate button or typing OK, etc.) from a purchasing agent, a supervisor, an accountant, and a vice-president before a purchase order is forwarded to a vendor,”

and “the electronic mail message will be processed according to the rule and forwarded to each individual for approval before the message is able to be sent to the vendor.” Schiavone, [0036]. In this example, the requirement of purchase order approval by the specific personnel involved in the approving process constitutes “*a set of rules applicable to the entities in the group.*” A POSITA would have found it obvious that Rohall’s electronic messaging system, which discloses “users/people” (“*the set of entities*”) in “a group or organization” (“*a population structure*”) would be additionally modified to include Schiavone’s teachings of “specific personnel holding positions of rank and involved in approving a purchase order” (“*group defined by their locations in a hierarchy*”) “*defined by the*” organizational chain-of-command (“*hierarchical population structure*”).

260. Therefore, the Rohall/Schiavone combination renders this element obvious.

k. Communication, Readership, Action, and Search Rules

[18.0] The computer-implemented method of claim 13, wherein the rules described by the compliance policy include one or more types of rules from the set consisting of: communication rules describing which population structure entities can communicate with other population structure entities via the messaging system;

readership rules describing which population structure entities can read sent messages from other population structure entities via the messaging system;

action rules describing actions a population structure entity can perform on a sent submessage in the messaging system; and

search rules describing which population structure entities can view which sent submessages when searching the sent submessages stored by the messaging system.

261. See element [5.0] meeting “communication rules” and “readership rules” and element [20.0] meeting “*action rules*.”

262. Therefore, the Rohall/Schiavone combination renders this element obvious.

I. Action Rules

[20.0] The computer-implemented method of claim 13, wherein the rules described by the compliance policy include action rules describing actions an entity can perform on a sent submessage in the messaging system, and wherein the action rules describe a number of times within a given time period that the entity can perform a specified action.

263. The Rohall/Schiavone combination renders this element obvious.

264. Initially, the 911 Patent provides an example where “[a]ction rules can also perform action budgeting, where a given action is permitted only up to a

specified number of times within a time period (e.g., five times a day).” Ex.1001, 10:45–48.

265. Schiavone teaches this limitation, for example, through its disclosures regarding auction rules:

The auction rule could specify that a message is sent to each of the other individuals every time any individual places a bid by e-mailing the seller, although the message is addressed only to the seller, and sending an automated e-mail message (“YOU WON!”) to the highest bidder after expiration of a predetermined auction period (e.g. a predetermined length of time or number of bids).

Schiavone, [0037] (emphases added). Here, the auction rule describes (1) actions an entity can perform—i.e., placing bids—on a sent submessage, and (2) the action a number of times within a given time period—i.e., “a predetermined length of time or number of bids.” Schiavone therefore teaches this claim.

266. Therefore, the Rohall/Schiavone combination renders this element obvious.

m. Group and Rules of Entities

[21.0] The computer-implemented method of claim 13, further comprising:

defining a group of entities in the population structure;

defining group rules applicable to sent submessages related to the group of entities; and

applying the group rules to a sent submessage related to the group of entities.

267. See elements [1.7] and [15.0], which meet the “*defining a group of entities in the population structure*” element of claim 21.

268. Additionally, Schiavone teaches “*defining group rules applicable to sent submessages...*” and “*applying the group rules to a sent submessage...*” Furthermore, and as discussed above, Schiavone explains that “**a rule may be constructed to determine compliance with any applicable law, rule, regulation, code, best practice**, etc. of any region, e.g., local, state, federal, foreign, etc. (“regional law”).” Schiavone, [0029].

269. Schiavone teaches an embodiment that describes and applies its auction rules to sent messages:

The ***auction rule*** could specify that ***a message is sent to each of the other individuals*** every time any individual places a bid by e-mailing the seller, although the message is addressed only to the seller, and sending an automated e-mail message (“YOU WON!”) to the highest bidder after expiration of a predetermined auction period (e.g. a predetermined length of time or number of bids).

Schiavone, [0037] (emphases added). In this example, the group consists of the auction participants, and the rules apply only to the participants when the email message(s) is/are sent to seller and/or bidders.

270. As another example, Schiavone teaches an embodiment relating to a “rule [that] could be constructed to facilitate viral group purchasing.” Schiavone, [0039]. For example, “a vendor could send an e-mail stating ‘I’ll sell these CD players for \$39 each if you can get 20 buyers’ and specifying a rule requiring the e-mail message to await 20 buyers’ indications of agreement to purchase for \$39 before automatically *sending a reply message to the vendor (and each of the other purchasers)* indicating that an aggregate of 20 buyers had been obtained, including their shipping address and credit card information gathered from their personal profile data, etc.” *Id.* (emphasis added). And “the e-mail message could be circulated until a predetermined number of buyers have agreed to purchase, at which point a reply is forward to the vendor and the sale transactions are completed at a group, discounted rate.” *Id.* In this scenario, the group consists of the purchasers and seller, whereas the rules govern when, how, and to whom the messages are sent to, satisfying the claimed “*defining...*” and “*applying the group rules to sent submessages...*”

271. Accordingly, the Rohall/Schiavone combination renders this element obvious.

C. Ground 2: The Combination of Rohall, Schiavone, and Gomes renders obvious claims 9 and 27

272. As demonstrated above, the Rohall/Schiavone combination renders obvious each element of claims 1-8, 10-16, 18, 20-25, and 27-28. For the reasons discussed below, to the extent one asserts certain claims elements are not rendered obvious by the Rohall/Schiavone combination, those elements would have been obvious based on Gomes.

1. The Combination of Rohall, Schiavone, and Gomes

273. A POSITA would have been motivated to combine Rohall's messaging system server with Schiavone's rule-based processing of messages for the reasons demonstrated above. See Section X.B. A POSITA would have additionally been motivated to combine the de-duplication teachings in Gomes with the Rohall/Schiavone combination.

274. Gomes is analogous art as the 911 Patent (and Rohall/Schiavone) because it is directed at electronic mail messaging systems. Gomes, 2:4-16, Abstract. Gomes is reasonably pertinent to a problem described in the 911 Patent: automating the process of filtering e-mails, *e.g.*, to remove duplicates, which the 911 Patent and Gomes explain is a problem of current email systems. Gomes, 7:2-11, 1:49-61. Ex.1001, 1:36-45.

275. A POSITA would have been motivated to combine Gomes with the Rohall/Schiavone combination because Gomes’ teachings of de-duplication would improve the Rohall/Schiavone combination by requiring less memory/storage space, inline with Rohall’s goal of “reduce[ing] memory requirements” and performing “memory management” and Schiavone’s concerns of “significant and increasing costs associated with ... [data] storage.” Rohall, [0064], [0043]. Schiavone, [0002].

276. A POSITA would have been motivated to combine Gomes with the Rohall/Schiavone combination because Gomes’ teachings of applying filters or rules to identify “compliant data” (Gomes, Abstract, 9:21-27) in a database would address Schiavone’s concern that presently “few standards, rules, or laws apply to govern how the e-mail communications medium is used,” and that “even with respect to laws, regulations and best practices ... there are no satisfactory techniques for ensuring compliance of e-mail messages with applicable standards and/or laws and/or otherwise determining appropriateness of a message for an intended recipient,” and therefore, Gomes’ teachings would improve the filtering capabilities of the Rohall/Schiavone combination. Schiavone, [0006]. Moreover, a POSITA would have been further motivated to combine Gomes with the Rohall/Schiavone combination because Gomes’ teachings of data retention policies (Gomes, 9:21-42) would ensure that a messaging system according to the teachings of the combination

meets the legal and statutory requirements of data retention in organizations/businesses.

277. Additionally, a POSITA would have been motivated to combine Gomes with the Rohall/Schiavone combination because Gomes' methods apply to data from different data sources (Gomes, 13:16-21) would improve the Rohall/Schiavone combination by allowing a messaging system to have greater applicability and versatility in handling different data types.

278. Moreover, a POSITA would have had a reasonable expectation of success in implementing Gomes's teachings in the Rohall/Schiavone combination because Gomes describes an embodiment using the same type of email file system as Rohall, "Lotus Notes files." Gomes, 6:27-31. Rohall, [0060]. Additionally, a POSITA would have had a reasonable expectation of success in making the combination because it would be a simple substitution well within the skill of a POSITA. For example, modifying the Rohall/Schiavone combination based on Gomes would be nothing more than adding an extra processing module/engine in the messaging system server that incorporates Gomes's teachings of a deduplication engine and further, adding such a module/engine would not change the principle of operation of the messaging system because it is still functional on Lotus Notes files.

2. Dependent Claim 9

a. Search Rules

[9.0] The messaging server of claim 1, wherein the rules described by the compliance policy include search rules describing which population structure entities can view which sent submessages when searching the sent submessages stored by the messaging system.

279. Although the Rohall/Schiavone combination does not expressly disclose this claim element, it would have been obvious in view of the combination of Rohall and Schiavone in further view of Gomes. And, as explained above, it would have been obvious for a POSITA to modify the Rohall/Schiavone combination to additionally include the teachings of Gomes. *See* X.C.1.

280. Gomes describes an example of a legal repository rule including “defin[ing]” “the criteria for a **search of files, messages**, records, or other data ... based on the names of the parties involved” which would “**allow authorized users access** to the legal repository (or any other “targeted data” database),” *e.g.*, “authorized users may be allowed read-only access to the evidence data and read and write access to the other information.” Ex.1016, 11:40-47, 12:46-55; *see also* Ex.1016, 13:1-6 (disclosing an embodiment in which “only read access may be granted to the authorized users”). Gomes explains that the legal repository “may include files or other records related to a legal matter.” Indeed, Gomes identifies that its methods apply to different file types, including “email files” such as

“Exchange Database (“EDB”) files” and “Lotus Notes files.” Ex.1016, 6:21-31. Therefore, Gomes discloses “a rule including the criteria for a search” (“*a search rule[]*”) “defining” (“*describing*”) “authorized users” (“which population structure entities can view which sent submessages”) “*when*” “accessing” (“*searching*”) the legal repository storing data such as “email messages” (e.g., “*the sent submessages*”).

281. Moreover, at the time of the claimed invention, limiting searches of an electronic database to certain groups of people or population structure entities was a well-known technique to a POSITA. For example, US Patent Pub. No. 2004/0260698 to MacMillan describes:

From block 30, the process advances to block 32, where the e-mail server 14 receives the e-mail sent from the wireless communication device 16 and determines whether the user is authorized to access the private database 12. The e-mail server **14** may determine whether the user is authorized to access the private database 12 by analyzing the e-mail to identify the e-mail address of the wireless communication device 16, then **comparing the identified e-mail address with a list of e-mail addresses of users authorized to access the private database 12. According to one embodiment, the list of e-mail addresses of users authorized to access the private database 12 may reside at the e-mail server 14. According to another embodiment, the list of e-mail addresses of users authorized to access the private database 12 may reside external to the e-mail server 14.**

The identification module 22 may identify the e-mail address of the wireless communication device 16 by analyzing the “from” address of the e-mail received from the wireless communication device 16. The authorization module 24 may compare the identified e-mail address with the list of e-mail addresses of users authorized to access the private database 12. If the list of e-mail addresses includes the e-mail address of the wireless communication device 16, the e-mail server 14 deems the user to be authorized to access the private database 12.

Ex.1017, [0013]-[0014]. See also claims 4-7 in MacMillan describing access limited to authorized e-mail addresses:

4. The method of claim 1, wherein sending the query includes authorizing access to the private database.
5. The method of claim 4, wherein authorizing access to the private database includes identifying an e-mail address of the wireless communication device.
6. The method of claim 5, wherein authorizing access to the private database includes comparing the e-mail address of the wireless communication device with a list of authorized e-mail addresses.
7. The method of claim 6, wherein authorizing access to the private database includes confirming that the list of authorized e-mail addresses includes the e-mail address of the wireless communication device.

MacMillan, Cls. 4-7.

282. Therefore, the combination of Rohall and Schiavone in further view of Gomes renders obvious this claim.

3. Independent Claim 27

283. The combination of Rohall and Schiavone teaches or suggests all of the elements of claim 27 for the same reasons demonstrated above with respect to claims 1, 10, and 13 with the below-mentioned distinctions noted.

Claim 27	Analogous Claim Limitations	Analogous Sections of Ground 1
[27.0]	[1.0], [10.0], [13.0], [27.0]	§X.B.1.a
[27.1]	[1.1], [10.1], [13.2] <i>and see below</i>	§X.B.1.b
[27.2]	[1.2], [10.2], [13.3] <i>and see below</i>	§X.B.1.c
[27.3]	[1.3], [10.3], [13.4]	§X.B.1.d
[27.4]	<i>see below</i>	N/A
[27.5]	[1.4], [10.4], [13.1]	§X.B.1.e
[27.6]	[1.5], [10.5]	§X.B.1.f
[27.7]	[1.6], [10.6], [13.5]	§X.B.1.g
[27.8]	[1.7], [10.7], [13.6]	§X.B.1.h
[27.9]	[1.8], [10.8], [13.7]	§X.B.1.i
[27.10]	[1.9], [10.9], [13.8]	§X.B.1.j

284. Moreover, other than replacing “*message*” in [1.1] and [1.2] with “*email*” in [27.1] and [27.2], the language in claim 27 is substantively the same as claims 1 and 10. Unlike claims 1 and 10, claim 27 includes element [27.4] which recites “*each of the sent submessages is stored exactly once within the email database*” and that “*the first sent submessage is not duplicated within the email database,*” and this element is taught by the combination of Rohall and Schiavone, as further modified by the teachings of Gomes, as explained below.

a. Sent submessages stored exactly once within the email database

[27.4] each of the sent submessages is stored exactly once within the email database, and upon receipt of a reply to a first email having a first email container containing a first sent submessage, the reply is added to the message container as an additional sent submessage but the first sent submessage is not duplicated within the email database; and

285. Gomes discloses this claim element.

286. First, while the 911 Patent explains the problems of “current e-mail systems” in having multiple copies of the same message, the functionality claimed in element [27.4] is not described in the specification—it only appears in the claims. For instance, the term “duplicate” does not appear anywhere in the specification. However, the 911 Patent does explain that “[w]hen a user writes a message, it is

formed into the first e-mail, but may then go on to be included in many other e-mails during its lifetime. This results in many copies of the same, user-authored, message in different, unrelated, mail ‘snapshots.’ Storing multiple copies of the same messages is inefficient and undesirable.” Ex.1001, 1:39-45. Thus the 911 Patent thus explains that storing multiple copies of the same message is inefficient, for example, from the perspective of memory/storage requirements.

287. Moreover, at the time of filing of the 911 Patent, as explained above, *see* VIII.D (citing Ex.1009, 4-5 and 8; Ex.1018, 4:54-58 and Fig. 14; and Ex.1019, 3), techniques for removing duplicate records of data (*e.g.*, emails stored in the Lotus Notes format and other types of documents) was known in the art.

288. Gomes discloses “email files can be ... de-duplicated, including de-duplication of attachments, message bodies and metadata ... **such that only one copy of the content is in the database** regardless of the location of the file,” (Gomes, 2:4-16, Abstract) and it would have been obvious for a POSITA that Rohall’s messaging system server, as modified by Schiavone’s teachings of rule-based processing of electronic mail messages would be further modified to include Gomes’ teachings, and thus Rohall’s modified messaging system server would additionally include a de-duplication engine taught in Gomes resulting in “*each of the sent submessages is stored exactly once within the email database*” and that “*the*

first sent submessage is not duplicated within the email database,” for the reasons explained in Section X.C.1.

289. Gomes recognizes the benefits of eliminating duplicates in environments processing enterprise-wide data. “By having fewer files, less storage space is needed and quicker searches can be performed.” Gomes, 3:40-42. And further, by “automating the operation, the de-duplication operation can occur real-time without any human intervention.” Gomes, 13:32-34.

290. Gomes describes the details of how it performs de-duplication:

[T]he de-duplication operation, as performed by de-duplication engine 144, can comprise separating each of the original data into a content portion and a metadata portion at step 422. For a file, the metadata portion may include an Internet Protocol address or other address of a computer from which the file is obtained. Additional data within the metadata portion may include the file path and filename, date the file was last saved on that particular computer, a pointer to the content portion within that computer, size of the file, and the like.

The method can further comprise hashing at least a portion of the content and metadata portions of each original data to form hashed values using de-duplication engine 144 at step 424. One hashed value corresponds to the content portion, and another hashed value corresponds to the metadata portion.

Gomes, 7:11-26.

The method can further comprise comparing the hashed values for the content portions using de-duplication engine 144 (step 426). For previously collected data within

collective database 162, the de-duplication engine 144 compares the hashed values for those previously collected data files' content portions are stored within database 162. Before storing the content portion of a current electronic file within database 162, the hashed value for the content portion of the current file is compared to hashed values of the content portions of other previously stored files within database 162. If a match is not found, the method can comprise storing a copy of the content portion in collective database 162 at step 502 in FIG. 5. If a match between hashed values is found, the content portion is already within database 162. In one embodiment, rather than storing multiple copies of the same content portion, only one copy of the content is stored within collective database 162.

Gomes, 7:61-8:10.

291. Additionally, because de-duplication techniques were known in the art (see Section VIII.D), Gomes recognizes that de-duplication can be done in multiple ways.

After reading this specification, skilled artisans realize that the foregoing are simply examples of the de-duplication embodiments and de-duplication can be implemented in a variety of different manners. In one embodiment, the hash of the content or metadata portion may be performed on only a portion of the content or metadata, respectively. For example, the hash of the metadata portion may be performed on the creation date, author, type of document, originating department, or nearly any other attribute of the metadata, by itself, or in combination with other attributes. Similarly, any portion of the content portion may be used for creating a hash value using the hashing algorithm. In addition, any number of hashes may be performed (e.g., hashing at least two hashed values to generate another hashed value). Still

further, the hash can be formed from combination(s) of portions or all of the content and metadata. Still further, a different hash algorithm may be used. After reading this specification, skilled artisans will appreciate that listing every possible embodiment would be nearly impossible and that many other methodologies may be used to replace or in conjunction with those described herein without departing from the scope of the invention.

Gomes, 7:39-60.

292. Accordingly, the combination of Rohall, Schiavone, and Gomes renders obvious element [27.4].

XI. OBJECTIVE INDICIA OF NON-OBVIOUSNESS

293. I am not aware of any objective indicia of non-obviousness that would change my opinion that the challenged claims of the 911 Patent are obvious in view of the prior art discussed above in this declaration. It is my opinion that any potential objective indicia of non-obviousness would not overcome the clear disclosures in the prior art showing that the subject matter of these claims would have been obvious to a POSITA as of May 12, 2004. However, I reserve the right to supplement my opinions Patent Owner alleges or offers any purported evidence of objective indicia of non-obviousness.

XII. CONCLUSION

294. I declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true, and that these statements were made with knowledge that willful false statements and

the like so made are punishable by fine or imprisonment, or both, under section 1001 of Title 18 of the United States Code.

A handwritten signature in black ink, appearing to read 'J. R. Black, Jr.', with a stylized, cursive script.

Dr. John R. Black, Jr.

Dated: 1/17/2025

CLAIM LISTING

Claim 1	
[1.0]	A computerized messaging server in an electronic messaging system, comprising:
[1.1]	a messaging module adapted to control a message database storing messages sent among users of the messaging system,
[1.2]	at least one of the sent messages stored in the message database comprising a message container containing relational references pointing to a plurality of sent submessages stored externally to the message container in the message database,
[1.3]	at least one of the sent submessages being one of a reply to and a forward of other sent submessages of the sent message; and
[1.4]	a governance module adapted to control a governance policy database storing a compliance policy describing rules applicable to the sent submessages,
[1.5]	the governance module further adapted to determine rules described by the compliance policy that are applicable to ones of the plurality of sent submessages,
[1.6]	wherein different rules of the compliance policy are applicable to different ones of the plurality of sent submessages;
[1.7]	wherein the messaging system is utilized by a set of entities in a population structure and wherein a rule of the compliance policy includes one or more rule elements describing whether the rule is applicable to a sent submessage, and
[1.8]	wherein a rule of the compliance policy includes one or more rule elements describing whether the rule is applicable to a sent submessage,
[1.9]	at least one of the rule elements selected from the group consisting of: a source element describing a population structure entity that is a source of the sent submessage; a target element describing a population structure entity that is a recipient of the sent submessage; and a rank describing a ranking of the rule relative to other rules in the compliance policy.
Claim 2	
[2.0]	The messaging server of claim 1, wherein the compliance policy describes a population structure of end-users of the messaging system, a set of jobcodes with which the sent submessages can be classified,

	and a set of rules applicable to the population structure, jobcodes, and sent submessages.
Claim 3	
[3.0]	The messaging server of claim 2, wherein a jobcode describes a set of jobcode rules applicable to all sent submessages classified with the jobcode.
Claim 4	
[4.0]	The messaging server of claim 2, wherein the messaging module is further adapted to receive an indication of a jobcode and of a sent submessage from an end-user, and is further adapted to associate the jobcode with the sent submessage responsive to receipt of the indication of the jobcode from the end-user.
Claim 5	
[5.0]	The messaging server of claim 1, wherein the rules described by the compliance policy include one or more types of rules from the set consisting of: communication rules describing which population structure entities can communicate with other population structure entities via the messaging system; and readership rules describing which population structure entities can read messages from other population structure entities via the messaging system.
Claim 6	
[6.0]	The messaging server of claim 1, wherein the rules described by the compliance policy include one or more types of rules from the set consisting of: validity rules describing a time period for which a sent submessage or entity in the population structure is useable; and retention rules describing a time period after which a sent submessage or entity in the population structure should be deleted from the messaging system.
Claim 7	
[7.0]	The messaging server of claim 1, wherein the governance policy database further stores at least one prior compliance policy active at a time prior to a time at which the compliance policy was active, and wherein the sent submessages of the stored messages are associated with a compliance policy that was active at a time that content of the sent submessages was created.
Claim 8	
[8.0]	The messaging server of claim 7, wherein at least one of the messages stored in the message database has a first sent submessage associated

	with a first compliance policy and a second sent submessage associated with a second compliance policy different from the first compliance policy, and wherein the governance module applies the rules described by the first compliance policy to the first sent submessage and applies the rules described by the second compliance policy to the second sent submessage.
Claim 9	
[9.0]	The messaging server of claim 1, wherein the rules described by the compliance policy include search rules describing which population structure entities can view which sent submessages when searching the sent submessages stored by the messaging system.
Claim 10	
[10.0]	A computer program product having a non-transitory computer-readable medium having embodied thereon program code for use in an electronic messaging system, the program code comprising:
[10.1]	a messaging module adapted to control a message database storing messages sent among users of the messaging system to at least one other end user of the messaging system,
[10.2]	at least one of the sent messages comprising a message container containing relational references pointing to a plurality of sent submessages stored in the message database externally to the message container,
[10.3]	at least one of the sent submessages being one of a reply to and a forward of other sent submessages of the sent message; and
[10.4]	a governance module adapted to control a governance policy database storing a compliance policy describing rules applicable to the sent submessages,
[10.5]	the governance module further adapted to determine rules described by the compliance policy that are applicable to ones of the plurality of sent submessages
[10.6]	wherein different rules of the compliance policy are applicable to different ones of the plurality of sent submessages
[10.7]	wherein the messaging system is utilized by a set of entities in a population structure and wherein a rule of the compliance policy includes one or more rule elements describing whether the rule is applicable to a sent submessage,

[10.8]	wherein a rule of the compliance policy includes one or more rule elements describing whether the rule is applicable to a sent submessage,
[10.9]	at least one of the elements selected from the group consisting of: a source element describing a population structure entity that is a source of the sent submessage; a target element describing a population structure entity that is a recipient of the sent submessage; and a rank describing a ranking of the rule relative to other rules in the compliance policy.
Claim 11	
[11.0]	The computer program product of claim 10, wherein the rules described by the compliance policy include one or more types of rules from the set consisting of: communication rules describing which population structure entities can communicate with other population structure entities via the messaging system; readership rules describing which population structure entities can read messages from other population structure entities via the messaging system; action rules describing actions a population structure entity can perform on a message component in the messaging system; and search rules describing which population structure entities can view which message components when searching the messaging components stored by the messaging system.
Claim 12	
[12.0]	The computer program product of claim 10, wherein the rules described by the compliance policy include one or more types of rules from the set consisting of: validity rules describing a time period for which a sent submessage or entity in the population structure is useable; and retention rules describing whether a sent submessage or entity in the population structure should be deleted from the messaging system.
Claim 13	
[13.0]	A computer-implemented method of managing messages in an electronic messaging system, comprising:
[13.1]	defining a compliance policy describing rules applicable to sent submessages of sent messages in the electronic messaging system;
[13.2]	receiving a message sent by one end user of the messaging system to at least one other end user of the messaging system,
[13.3]	the sent message comprising a message container containing relational references pointing to a plurality of sent submessages stored in a message database externally to the message container,

[13.4]	the plurality of sent submessages including a first sent submessage and a second sent submessage, at least one of the sent submessages being one of a reply to and a forward of another sent submessage of the plurality of sent submessages; and
[13.5]	selectively applying rules of the compliance policy to the sent submessages;
[13.6]	wherein the messaging system is utilized by a set of entities in a population structure and
[13.7]	wherein a rule of the compliance policy includes one or more rule elements describing whether the rule is applicable to a sent submessage,
[13.8]	at least one of the elements selected from the group consisting of: a source element describing a population structure entity that is a source of the sent submessage; a target element describing a population structure entity that is a recipient of the sent submessage; and a rank describing a ranking of the rule relative to other rules in the compliance policy.
Claim 15	
[15.0]	The computer-implemented method of claim 13, wherein the compliance policy describes the set of entities in the population structure, the method further comprising: defining a group containing a set of entities, the entities in the group defined by their locations in a hierarchy defined by the hierarchical population structure; and defining a set of rules applicable to the entities in the group.
Claim 18	
[18]	The computer-implemented method of claim 13, wherein the rules described by the compliance policy include one or more types of rules from the set consisting of: communication rules describing which population structure entities can communicate with other population structure entities via the messaging system; readership rules describing which population structure entities can read sent messages from other population structure entities via the messaging system; action rules describing actions a population structure entity can perform on a sent submessage in the messaging system; and search rules describing which population structure entities can view which sent submessages when searching the sent submessages stored by the messaging system.

Claim 20	
[20.0]	The computer-implemented method of claim 13, wherein the rules described by the compliance policy include action rules describing actions an entity can perform on a sent submessage in the messaging system, and wherein the action rules describe a number of times within a given time period that the entity can perform a specified action.
Claim 21	
[21.0]	The computer-implemented method of claim 13, further comprising:
	defining a group of entities in the population structure;
	defining group rules applicable to sent submessages related to the group of entities; and
	applying the group rules to a sent submessage related to the group of entities.
Claim 27	
[27.0]	A computerized messaging server in an electronic messaging system, comprising:
[27.1]	a messaging module adapted to control an email database storing emails sent among users of the messaging system, wherein
[27.2]	at least one of the sent emails comprises an email container containing relational references pointing to a plurality of sent submessages stored in the email database externally to the email container,
[27.3]	at least one of the sent submessages being one of a reply to and a forward of other sent submessages of the message, and
[27.4]	each of the sent submessages is stored exactly once within the email database, and upon receipt of a reply to a first email having a first email container containing a first sent submessage, the reply is added to the message container as an additional sent submessage but the first sent submessage is not duplicated within the email database; and
[27.5]	a governance module adapted to control a governance policy database storing a compliance policy describing rules applicable to the sent submessages,
[27.6]	the governance module further adapted to determine rules described by the compliance policy that are applicable to ones of the plurality of sent submessages,
[27.7]	wherein different rules of the compliance policy are applicable to different ones of the plurality of sent submessages,
[27.8]	wherein the messaging system is utilized by a set of entities in a population structure and

[27.9]	wherein a rule of the compliance policy includes one or more rule elements describing whether the rule is applicable to a sent submessage,
[27.10]	at least one of the elements selected from the group consisting of: a source element describing a population structure entity that is a source of the sent submessage; a target element describing a population structure entity that is a recipient of the sent submessage; and a rank describing a ranking of the rule relative to other rules in the compliance policy.