

**IN THE UNITED STATES DISTRICT COURT
FOR THE EASTERN DISTRICT OF TEXAS
MARSHALL DIVISION**

ST. LUKE TECHNOLOGIES, LLC,

Plaintiff,

v.

GOOGLE, INC. AND ALPHABET, INC.

Defendants

Civil Action No. _____

JURY TRIAL DEMANDED

COMPLAINT FOR PATENT INFRINGEMENT

Plaintiff St. Luke Technologies, LLC (“St. Luke” or “Plaintiff”), by and through its attorneys, brings this action and makes the following allegations of patent infringement relating to U.S. Patent Nos. 7,181,017 (“the ‘017 patent”); 8,316,237 (“the ‘237 patent”); 7,805,377 (“the ‘377 patent”); 7,587,368 (“the ‘368 patent”); 8,498,941 (“the ‘941 patent”); 8,380,630 (“the ‘630 patent”); and 8,600,895 (“the ‘895 patent”) (collectively, the “patents-in-suit”). Defendants Google, Inc. and Alphabet, Inc. (collectively, “Google” or “Defendant”) infringe the patents-in-suit in violation of the patent laws of the United States of America, 35 U.S.C. § 1 *et seq.*

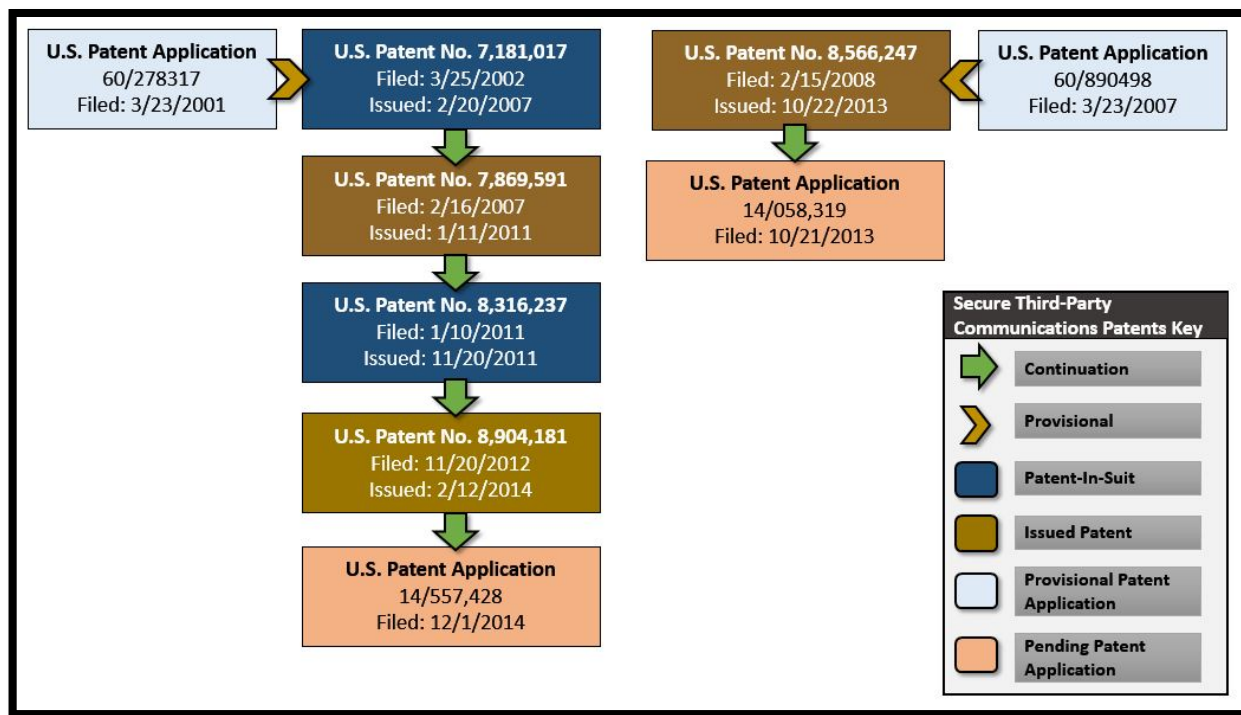
INTRODUCTION

1. In an effort to expand its product base and profit from the sale of infringing cloud computing encryption technologies and information record infrastructure technologies, Google has unlawfully and without permission copied the technologies and inventions of Dr. Robert H. Nagel, David P. Felsher, and Steven M. Hoffberg.

2. Dr. Nagel, Mr. Felsher, and Mr. Hoffberg are the co-inventors of the ‘017 patent, the ‘237 patent, and U.S. Patent Nos. 7,869,591 (“the ‘591 patent”), 8,904,181 (“the ‘181 patent”), and 8,566,247 (“the ‘247 patent”) (collectively, the “Secure Third-Party Communications Patents” or “STPC patents”). The STPC patents have been cited in over 550

United States patents and patent applications as prior art before the United States Patent and Trademark Office.¹ The STPC patents disclose systems and methods for secure communications over a computer network where a third party (intermediary) performs a requisite function with respect to the transaction without requiring the intermediary to be trusted with respect to the private information or cryptographic keys for communicated information. The inventions taught in the STPC patents employ secure cryptographic schemes, which drastically reduce the risk of unauthorized disclosure of encrypted data.

3. The below diagram shows St. Luke’s STPC patents, pending STPC patent applications, and the STPC patents Google infringes.²



4. Over a decade after Dr. Nagel and his co-inventors conceived of the inventions disclosed in the STPC patents, Eran Feigenbaum, Google Applications director of security,

¹ Google has cited the STPC patents as relevant prior art in Patents assigned to Google. See U.S. Patent Nos. 8,978,093 and 9,071,440.

² St. Luke’s STPC patents are in two patent families claiming priority to U.S. Patent Applications 60/278,317 and 60/890,498.

described systems such as Dr. Nagel, Mr. Felsher, and Mr. Hoffberg's secure third party communications system as "part of the core DNA of [Google's] products."

It's a bit of a different model than trying to protect the end point devices, the laptop, the desktop. ***We focus on protecting the data, and have very limited data at the endpoint. We are an internet company — born and raised on the internet — and we have built security as part of the core DNA of our products.*** I find it a model that is really scalable to millions and millions of users, both from a technology and operations perspective.

An Interview with Google Apps Director of Security, Eran Feigenbaum, COBRY BLOG, available at: <http://www.cobry.co.uk/an-interview-with-google-enterprise-director-of-security-eran-feigenbaum/> (emphasis added).

5. Google has described data encryption systems such as the inventions disclosed in the STPC patents as a "primary design consideration for all of Google's infrastructure."

According to a Google white paper:

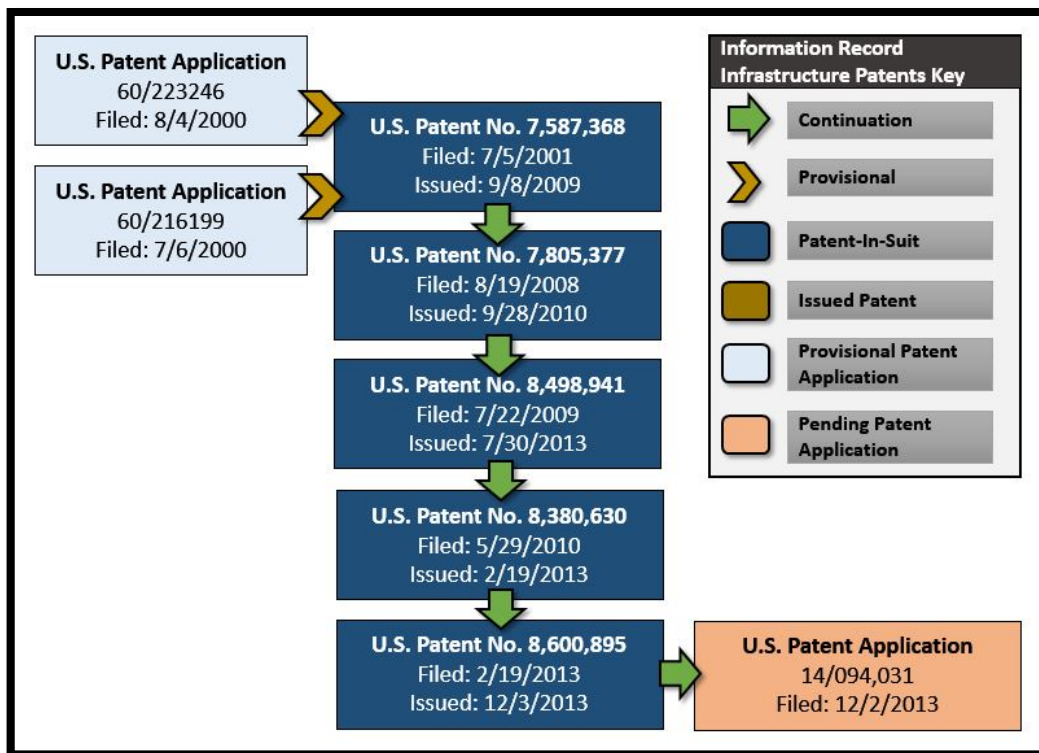
As a cloud pioneer, Google fully understands the security implications of the cloud model. Our cloud services are designed to deliver better security than many traditional on-premises solutions. We make security a priority to protect our own operations, but because Google runs on the same infrastructure that we make available to our customers, your organization can directly benefit from these protections. ***That's why we focus on security, and protection of data is among our primary design criteria.*** Security drives our organizational structure, training priorities and hiring processes. It shapes our data centers and the technology they house. It's central to our everyday operations and disaster planning, including how we address threats. ***It's prioritized in the way we handle customer data. And it's the cornerstone of our account controls, our compliance audits and the certifications we offer our customers.***

Google Cloud Platform Security Whitepaper, GOOGLE SECURITY WHITEPAPER (last updated May 26, 2015), available at: <https://cloud.google.com/security/whitepaper> (emphasis added).

6. Mr. Felsher is the inventor of the '941 patent, the '377 patent, the '368 patent, the '630 patent, and the '895 patent (collectively, the "Information Record Infrastructure Patents" or "IRI patents"). The IRI patents have been cited by over 970 United States patents and patent applications as prior art before the United States Patent and Trademark Office.

7. The IRI patents disclose systems and methods for distributing and granting access to data where data is stored in multiple external computer databases. The IRI patents address the difficult problem of authorizing access to protected information records where authorization will depend on the access privileges of the user.

8. The below diagram shows the IRI patent family tree, a pending IRI patent application, and the IRI patents Google infringes.



THE INVENTORS’ LANDMARK SECURE COMMUNICATION SYSTEMS

9. Mathematician Dr. Robert Nagel, the named inventor of two patents-in-suit, pioneered development of large-scale computer-based data distribution systems. In the 1970s Dr. Nagel developed some of the first computer systems for distributing encrypted data over computer networks. Dr. Nagel is the named inventor of twenty-three United States Patents. Dr. Nagel’s patents have been cited thousands of times by various companies, including Google. Later in life, Dr. Nagel founded two publicly traded companies, and served as a representative to the United Nations.

10. In 1975, Dr. Nagel developed a system harnessing burgeoning microprocessor power to broadcast stock prices and related data over coaxial cable and telephone networks. Dr.

Nagel's patented system was the foundation of Reuters's high-speed transmission technologies for distributing real-time market information.

Computer power behind the new information system is provided by a Digital Equipment Corp. PDP-8E with 32K memory and a multiprocessor system consisting of one PDP-11/35 with 64K memory and 2 PDP-11/50s, each with 96K memory.

The system was developed by Robert H. Nagel of IDR. Another patent for the high-speed transmission technique is expected to be issued shortly.

Reuters Gets News System Patent, COMPUTERWORLD at 36, April 23, 1975 (describing Dr. Nagel's development of one of the first terminals for displaying real-time stock market data).³

11. The data distribution system developed by Dr. Nagel in the mid-1970s was commercialized by Reuters and allowed the rapid transmission of market and news information over coaxial cable and telephone lines.⁴

³ See U.S. Patent Nos. 3,875,329, which issued on April 1, 1975. Dr. Nagel's work at IDR, Inc. (a subsidiary of then Reuters Group PLC) led to the development of U.S. Patent Nos. 3,889,054; 4,042,958; 4,064,494; 4,120,003, 4,135,213; and 4,148,066. These patents have been cited in over 830 patent applications and issued patents of companies including Cisco Technology, Inc., Sony Corporation, Intel Corporation, etc.

⁴ *Reuters Technical Development Chronology 1975-1979*. THE BARON, July 13, 2015). <http://thebaron.info/archives/technology/reuters-technical-development-chronology-1975-1979>.


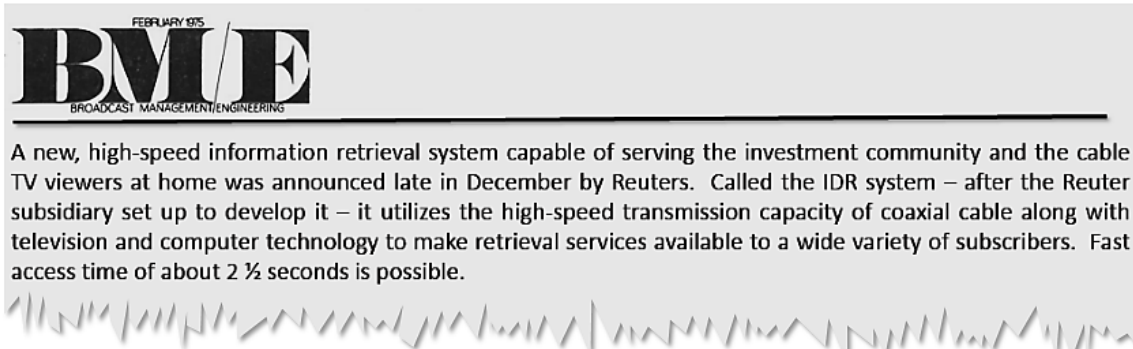


IMAGE OF THE DEC PDP-11/50 SYSTEM, COLUMBIA UNIVERSITY COMPUTING HISTORY ARCHIVE (circa 1976), <http://www.columbia.edu/cu/computinghistory/> (showing an installed PDP-11/50 device that was a component in Dr. Nagel's data distribution system).

12. Reuters sold thousands of information systems modeled on Dr. Nagel's patented inventions.⁵ Hundreds of companies including IBM, Intel, and Xerox cite Dr. Nagel's groundbreaking inventions described in his patents as relevant prior art in their own patents.⁶

⁵ *Reuters Technical Development Chronology 1975-1979*. THE BARON, July 13, 2015). <http://thebaron.info/archives/technology/reuters-technical-development-chronology-1975-1979> (More than 10,000 units are eventually produced. It revolutionizes the Monitor product financials and field staffing and provides valuable cash flow for IDR.”).

⁶ PROCEEDINGS OF THE DIGITAL EQUIPMENT USERS SOCIETY, DIGITAL EQUIPMENT CORPORATION PROCEEDINGS Vol. 3 Issue 1 at 1 (1977) (“Reuters has developed a network to assist stock and commodity brokers and foreign exchange dealers by giving them the latest prices and rate of exchange via terminals in this book.”); ANNUAL REVIEW OF INFORMATION SCIENCE AND TECHNOLOGY, AMERICAN SOCIETY OF INFORMATION SCIENCE, AMERICAN DOCUMENTATION INSTITUTE Vol. 12 at 223 (1977) (“Reuters provides the user with a 1.2 Kbps leased connection to the nearest network processor or multiplexor. The Monitor user configuration is a Digital Equipment Corporation PDP 8 with up to three display units.”); REUTERS BLENDS CATV & COMPUTER SKILLS IN NEWS RETRIEVAL SYSTEM, DATA PROCESSING DIGEST at 12 (1975) (“Reuters has introduced in New York a high-speed information retrieval system for the investment community. The system was developed by Information Dissemination and Retrieval, Inc. (IDR), a Reuters subsidiary, and uses the high-speed transmission capacity of coaxial cable with television and computer technology.”).



Reuters Announces Retrieval System For Cable TV Subscribers, BROADCAST MANAGEMENT/ENGINEERING MAGAZINE at 9, February 1975.

13. In the 1990s, Dr. Nagel was the Chief Technology Officer of eSecure Docs, Inc., Founder of Digits Corporation, and Executive Vice President and Chief Technology Officer of InfoSafe Systems, Inc.⁷ Publications including Fortune Magazine and ComputerWorld described Dr. Nagel as a “noted computer scientist” for his groundbreaking work⁸—work that led to the inventions disclosed in the patents-in-suit.

The technology Nagel designed at InfoSafe Systems, Inc., won the Seybold Award for Excellence as the “most innovative product of the year.” His work in high technology received major press coverage in such publications as Fortune, Forbes, and Business Week. He testified before Congress on the capabilities of a system he designed for NASDAQ.

Aliye Pekin Celik, OUR COMMON HUMANITY IN THE INFORMATION AGE: PRINCIPLES AND VALUES FOR DEVELOPMENT at 191 (2007).

14. Following his development of groundbreaking electronic data distribution systems for Reuters, Dr. Nagel used his insights to develop the secure communications technologies that are used today by Google and many of the world’s largest corporations without attribution or compensation.

⁷ In addition to his work in private industry, Dr. Nagel served as a consultant to the Defense Advanced Research Projects Agency (“DARPA”), responsible for the development of emerging technologies used by the U.S. Department of Defense. Dr. Nagel was a designer of the Navy’s Tactical Air Navigation System (“TACAN”) and assisted in the development of the nuclear reactor that powers the Navy’s Seawolf class of nuclear submarines. Dr. Nagel was also the developer of the Hot Well Liquid Level Control system that is a part of the control system of the nuclear power plant aboard the Seawolf, Defender and other submarines.

⁸ See Rick Tetzeli, et al., *Fortune Checks Out 25 Cool Companies For Products, Ideas, And Investments*, FORTUNE MAGAZINE (July 11, 1994).

15. Dr. Nagel foresaw the need for enabling secure communications between two parties wherein an intermediary performs a requisite function with respect to the transaction without requiring the intermediary to be trusted with respect to the private information or cryptographic keys for communicated information.

16. Dr. Nagel's interest in developing secure systems for the provision of highly secure data was driven in part by his experience being totally blind.⁹ Dr. Nagel recognized that the growing adoption of the Internet and increased computational power presented unique challenges to the security of medical records. Dr. Nagel also had the insight that the challenges presented in controlling access to secure medical records could be applied outside the context of medical records, with wide applicability to the security of data on networks where an intermediary could have access to secure information.


17. The rise of cloud computing (the delivery of on-demand computing resources over a distributed network), has made Dr. Nagel and his co-inventors' insights uniquely valuable. Medical records, financial information, email messages, and other forms of electronic data are now placed on remote servers and accessed via a network by a diverse variety of users, under a diverse variety of circumstances.

18. The inventions disclosed in the STPC patents address shortcomings in systems available at the time of the patents' conception—for example, the need for users in particular contexts, to access and/or modify data stored at or by an intermediary without allowing the intermediary to access an unencrypted version of the data.

19. Prior art systems such as the "Micali Fair Encryption scheme do[] not . . . allow communications of a secret in which only one party gains access to the content, and in which the

⁹ Dr. Nagel served as a representative to the United Nations Committee that authored the International Convention on the Protection of the Rights of Dignity of Persons with Disabilities. See Jan Jekielek, *Human Rights Panel Explores Implementation of Rights and Global Well-Being*, Epoch Times, December 3, 2010, <http://www.cccun.net/cccun-12-2-10-eventepochtim.pdf> ("Nagel, who is blind himself. He expounded on the remarkable accomplishment that is the Convention on the Rights of Persons with Disabilities, the 21st century's first U.N. human rights convention.").

third party or parties and one principal operate only on encrypted or secret information.” ‘237 patent, col. 2:40-44.



‘237 Patent Fig. 1.

20. Dr. Nagel worked with Steven Hoffberg and David P. Felsher to develop the systems and methods disclosed in the STPC patents. The inventions taught in these patents relate to the secure transmission of data—for example, wherein an intermediary performs a requisite function with respect to a secure data transmission without requiring the intermediary to be trusted with the private, secure contents of the transmission and/or without requiring the intermediary to have access to the cryptographic keys required to access the protected information. The STPC patented systems and methods employ secure cryptographic schemes, which reduce the risks and liability of unauthorized disclosure of private information as it travels across a network.

21. Mr. Hoffberg holds a Master of Science degree from the Massachusetts Institute of Technology and an advanced degree in electrical engineering from Rensselaer Polytechnic

Institute. Mr. Hoffberg is a named inventor on sixty-seven patents in the fields of telematics, wireless ad hoc networking, image and audio signal processing, and cryptography. Mr. Hoffberg also spent three years in the University of Connecticut Medical School Medical Doctorate Program.

22. Mr. Felsher is an appellate attorney, health care activist, and inventor. After graduating from MIT with a Bachelor of Science Degree in Chemistry, Mr. Felsher went on to earn an MBA from the Wharton School of Business of the University of Pennsylvania and a J.D. from Fordham Law School.¹⁰ Mr. Felsher has served as counsel to the Association of American Physicians and Surgeons, Inc.

23. The STPC patents have been cited in over 550 United States patents and published patent applications as prior art before the United States Patent and Trademark Office.¹¹ Companies whose patents cite the Secure Third-Party Communication Patents include:

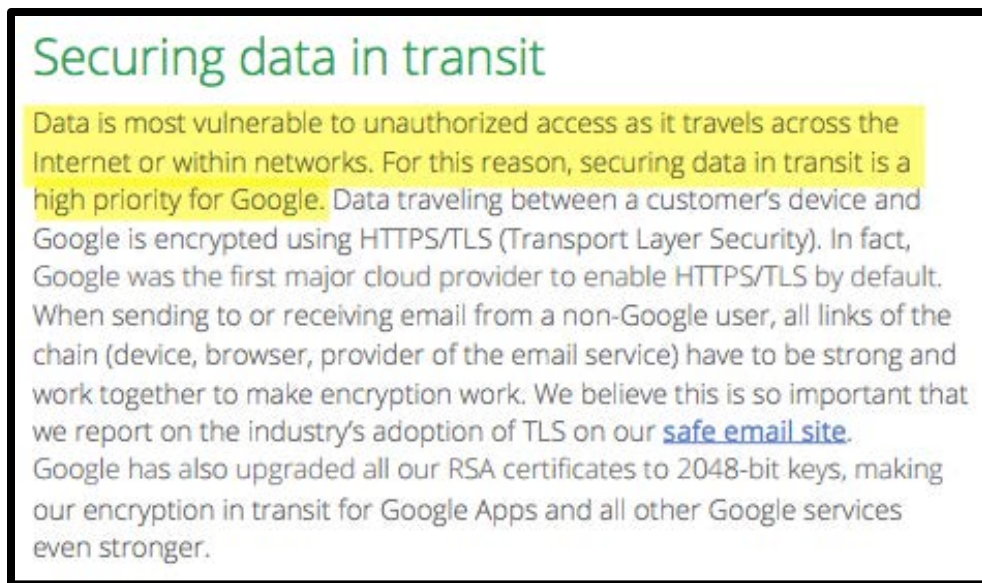
- Microsoft Corporation
- Nokia Corporation
- Apple, Inc.
- International Business Machines Corporation
- Massachusetts Institute of Technology
- NCR Corporation
- NetApp, Inc.
- Adobe Systems Incorporated
- American Express Travel Related Services Company, Inc.
- AT&T Intellectual Property LLP
- Canon Kabushiki Kaisha
- Hytrust, Inc.
- Cisco Technology, Inc.
- Intuit, Inc.
- Cloudera, Inc.
- Novell, Inc.
- **Google, Inc.**
- Teradata US, Inc.
- Mitsubishi Electric Corporation
- Texas Instruments, Inc.
- UnitedHealth Group Incorporated
- Fujitsu Limited

¹⁰ During his legal career, Mr. Felsher has been counsel of record on seventeen briefs to the United States Supreme Court.

¹¹ The 550 forward citations to the Secure Third-Party Communication Patents do not include patent applications that were abandoned prior to publication in the face of the Secure Third-Party Communication Patents.

- Hewlett-Packard Development Company, L.P.
- Verizon Patent and Licensing, Inc.
- Visa U.S.A. Inc.
- Western Digital Technologies, Inc.
- Xerox Corporation
- Yahoo!, Inc.
- Koninklijke Philips Electronics, N.V.
- Zynga Inc.
- Square, Inc.
- Sprint Communications Company L.P.
- Sony Corporation
- Siemens Aktiengesellschaft
- Sharp Laboratories of America, Inc.
- Sap AG
- EMC Corporation
- Samsung Electronics Co., Ltd.
- Ricoh Co., Ltd.
- Red Hat, Inc.
- Panasonic Corporation
- Broadcom Corporation
- Oracle International Corporation

24. The inventions taught in the STPC patents relate to the encryption of data passed through an intermediary and have been recognized by Google as important and valuable.



How Google Protects Your Data, GOOGLE FOR WORK SECURITY AND COMPLIANCE WHITEPAPER at 9 (highlighting added).

25. The adoption of secure encryption technologies is central to Google's infrastructure and its approach to customer data and it encourages and empowers its users to adopt additional security measures.

Google builds security into its structure, technology, operations and approach to customer data. Our robust security infrastructure and systems become the default for each and every Google Apps customer. But beyond these levels, users are actively empowered to enhance and customize their individual security settings to meet their business needs through dashboards and account security wizards.

How Google Protects Your Data, GOOGLE FOR WORK SECURITY AND COMPLIANCE WHITEPAPER at 16 (2015).

26. The IRI patents have been cited by over 970 United States patents and patent applications as prior art before the United States Patent and Trademark Office.¹² Companies whose patents cite the IRI patents include:

- Bank of America Corporation
- Siemens Medical Solutions Health Services Corporation
- AthenaHealth, Inc.
- Robert Bosch GmbH
- Thompson Reuters (Healthcare), Inc.
- Northrop Grumman Information Technology, Inc.
- McKesson Corporation
- Lockheed Martin Corporation
- Sandisk Technologies, Inc.
- Intel Corporation
- Greenway Medical Technologies, Inc.
- Medtronic, Inc.
- Sybase, Inc.
- General Electric Company
- Epic Systems Corporation
- Allscripts Software, LLC
- Ebay, Inc.
- 3Com Corporation
- Oracle International Corporation
- Intuit Inc.
- Gemalto N.V.
- Adobe Systems Incorporated
- Koninklijke Philips Electronics N.V.
- Electronic Data Systems Corporation
- American Express Travel Related Services Company, Inc.
- **Google, Inc.**
- Apple, Inc.
- McAfee, Inc.
- Hewlett-Packard Development Company L.P.
- EMC Corporation
- Blackboard, Inc.
- AT&T Intellectual Property LLP
- Cerner Innovation, Inc.
- Cisco Technology, Inc.
- Citrix System, Inc.

¹² The 970 forward citations to the IRI Patents and their related patent applications do not include patent applications that were abandoned prior to publication in the face of the IRI Patents.

- International Business Machines Corporation

THE PARTIES

27. Tyler, Texas-based St. Luke is committed to advancing the current state of innovation in the field of data encryption technologies for secure communications over a distributed network. In addition to the ongoing efforts of Messrs. Felsher and Hoffberg, St. Luke employs a resident of Tyler, Texas as a Technology Analyst. St. Luke is a Texas limited liability company with its principal place of business at 719 West Front Street, Suite 247, Tyler, Texas 75710.



28. St. Luke is a small, Texas-based company. St. Luke depends on patent protection to effectively license its innovative technologies and build its business. Like Defendant Google, St. Luke relies on its intellectual property. For example, Eric Schmidt, Google's CEO, explained that:

From a Google perspective, intellectual property rights are fundamental to how we operate because we operate based on a set of proprietary things, which we view ourselves as our own intellectual property, so the company wouldn't exist without basic intellectual property rights. In fact, we've supported a whole bunch of initiatives around that; especially outside the United States, where the rights are not so strong; and the trick is to find legal and business mechanisms that allow people to be properly compensated for their intellectual property while still encouraging use.

Eric Schmidt, CEO of Google Discusses Intellectual Property, FOREIGN POLICY MAGAZINE'S YOUTUBE CHANNEL (February 6, 2007), available at: <https://www.youtube.com/watch?v=Zi3Q40EPUjk#t=38>.

29. On information and belief, Google has acquired patents for technology that it did not invent¹³ and has asserted those patents and others in federal courts.¹⁴

30. On information and belief, Defendant Google, Inc. is a Delaware corporation, with its headquarters at 1600 Amphitheatre Parkway, Mountain View, CA 94043. On information and belief, Google can be served through its registered agent, Corporation Service Company, 211 E. 7th Street, Suite 620, Austin, Texas, 78701-3218.

31. On information and belief Defendant Alphabet, Inc. is a Delaware corporation with its headquarters at 1600 Amphitheatre Parkway, Mountain View, CA 94043. On information and belief, Alphabet, Inc. can be served through its registered agent, Corporation Service Company, 2710 Gateway Oaks Drive Ste. 150n, Sacramento, California, 95833.

32. According to Google's website, infringing products are offered for sale and sold throughout the United States and Canada, including in this District, through various channels. Google offers its infringing products through its distribution channel, which includes numerous distribution points in Texas. Further, Google advertises its infringing products throughout the Eastern District of Texas.

33. In addition, on information and belief, Google:

- Has offices in Dallas, Texas and multiple offices in Austin, Texas, where it employs engineers responsible for the creation, design, and implementation of the accused products;¹⁵

¹³ See, e.g., Quentin Hardy, *Google Buys Motorola for Patent Parts*, FORBES.COM TECHNOLOGY BLOG (August 15, 2011), <http://www.forbes.com/sites/quentinhardy/2011/08/15/google-buys-motorola-for-patent-parts/>; Amir Efrati, *Google Buys IBM Patents*, WALL STREET JOURNAL TECHNOLOGY SECTION, July 29, 2011, available at: <http://www.wsj.com/articles/SB10001424053111904800304576475663046346104> (Google buys more than 1,000 IBM patents); *Patent Transaction Trends for 2012*, RELECURA PATENT TRANSACTION ANALYSIS at 12, 37 (2013), available at: https://relecura.com/reports/Patent_transaction_trends_for_2012.pdf (Google buys 567 patents from Hitachi, 214 patents from KLI Consulting, 156 patents from Mossaid Tech, Inc., 118 patents from Modu Ltd., 100 patents from Computer Associates Think, Inc., 93 patents from Magnolia Broadband Inc., and 41 patents from Unisys Corp.).

¹⁴ *Google, Inc. v. BT Americas, Inc. et al.*, Case No. 13-cv-00254 (C.D. Cal. February 13, 2013) (Google asserted 4 patents, two of which were purchased from IBM and one was purchased from Fujitsu.).

- Has a major data center in Pryor, Oklahoma that plays a central role in the method and manner of operation of the accused products;¹⁶
- Utilizes the State of Texas and its resources for the development and testing of Google products, including but not limited to its autonomous self-driving cars and fiber internet services.¹⁷

JURISDICTION AND VENUE

34. This action arises under the patent laws of the United States, Title 35 of the United States Code. Accordingly, this Court has exclusive subject matter jurisdiction over this action under 28 U.S.C. §§ 1331 and 1338(a).

35. Upon information and belief, this Court has personal jurisdiction over Defendant Google in this action because Google has committed acts within the Eastern District of Texas giving rise to this action and has established minimum contacts with this forum such that the exercise of jurisdiction over Google would not offend traditional notions of fair play and substantial justice. Defendant Google, directly and through subsidiaries or intermediaries (including distributors, retailers, and others), has committed and continues to commit acts of infringement in this District by, among other things, using, offering to sell, and selling products

¹⁵ Jan Buchholz, *Google this: Tech giant takes huge portion of Green Water Treatment redevelopment project*, AUSTIN BUSINESS JOURNAL, May 13, 2015, available at: <http://www.bizjournals.com/austin/blog/real-estate/2015/03/google-this-tech-giant-takes-huge-portion-of-green.html> (Google leasing 200,000 square feet of office space in downtown Austin alone).

¹⁶ Memorandum Opinion and Order Denying Google, Inc.'s Motion to Transfer, Dkt. No. 138 at 7, n. 4, in *Smartflash LLC et al v. Google, Inc. et al*, Case No. 14-cv-435-JRG-KNM (E.D. Tex. April 6, 2015) ("Google's Declaration points out that "No Google data centers (including the secure servers that store documents related to this litigation) are located in Texas." Dubey Decl. ¶ 33. Google does not declare that any Google data centers are located in California either. A brief Internet search confirms that there are none in California, but there are several in the Southeastern United States, including one in Pryor, Oklahoma. *Data Center Locations*, <http://www.google.com/about/datacenters/inside/locations/index.html> (last visited Apr. 6, 2015). The data center in Pryor, Oklahoma is closer to this District than any other data center is to the Northern District of California.")

¹⁷ Graham Rapier, *Google Self-Driving Cars Are Headed for Texas*, BUSINESS INSIDER, July 7, 2015, available at: <http://www.businessinsider.com/google-self-driving-cars-are-headed-for-texas-2015-7>; Conner Forrest, *Google Fiber hits Austin, Texas: Here's what it means for the future*, TECH REPUBLIC, December 4, 2014, available at: <http://www.techrepublic.com/article/google-fiber-hits-austin-texas-heres-what-it-means-for-the-future/>.

and/or services that infringe the asserted patents. Moreover, Google is registered to do business in the state of Texas, and has appointed Corporation Service Company, 211 E. 7th Street, Suite 620, Austin, Texas, 78701-3218, as its agent for service of process. This Court also has personal jurisdiction over Google because it has multiple offices in Texas.

36. Venue is proper in this district under 28 U.S.C. §§ 1391(b)-(d) and 1400(b). Defendant Google is registered to do business in Texas, and upon information and belief, has transacted business in the Eastern District of Texas and has committed acts of direct and indirect infringement in the Eastern District of Texas. In addition, Google has multiple places of business in Texas.

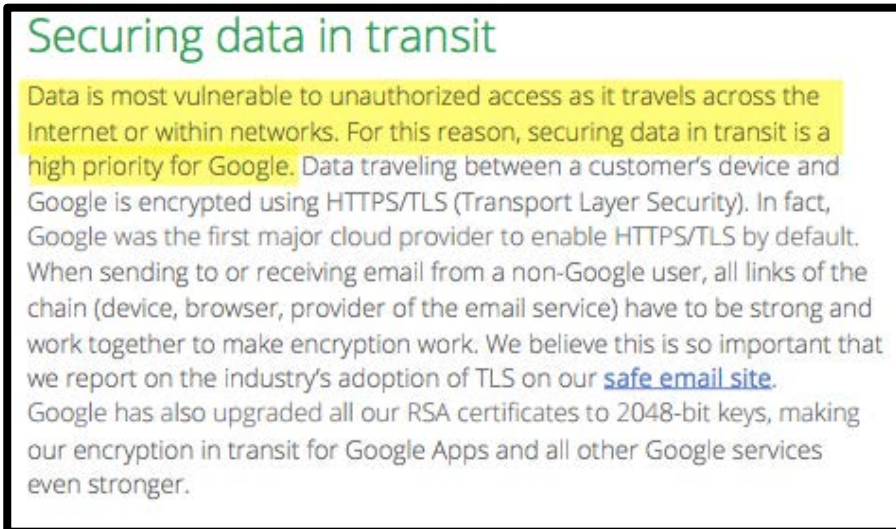
TECHNOLOGY BACKGROUND

37. Advances in computational power and the explosive growth of the Internet have led to the development of secure encryption systems and information record management systems that enable secure communications between two or more computers on a network where the data that is sent and/or processed by an intermediary without access to the plaintext data.

- *The STPC patents* teach specific computer based encryption systems, including systems that use composite key asymmetric cryptographic algorithms to avoid substantially revealing plaintext data during intermediate processing.
- *The IRI patents* teach specific computer based systems and methods, including systems for electronically structuring and controlling access to protected data in a plurality of external databases.

A. Secure Third Party Communications Patents

38. Google prizes systems that provide secure third party communications through an intermediary.



How Google Protects Your Data, GOOGLE FOR WORK SECURITY AND COMPLIANCE WHITEPAPER at 9 (2015).

39. Google's competitors such as Microsoft, Apple, and Oracle have confirmed the importance and value of encryption systems that protect data in the Cloud. Brendon Lynch, Chief Privacy Officer at Microsoft described the importance that Microsoft places on secure encryption in the cloud:

We share the same concerns as our customers do around government surveillance. We know that customers will not use technology that they do not trust that is what people should know about our [Microsoft's] approach to this . . . we're implementing strong encryption right throughout our services to ensure that governments can only access data by lawful means.

Brendon Lynch, *Microsoft Privacy and Compliance in the Cloud*, TRUSTWORTHY COMPUTING - VIDEO TRANSCRIPT (January 9, 2015), <https://www.youtube.com/watch?v=q5rwwQBTJxo>.

40. Tim Cook, Apple's Chief Executive Officer, has repeatedly stated that the use of encryption technologies is central to Apple's business.

Tim Cook: We've also communicated and demonstrated our commitment to respecting and protecting users' privacy with strong encryption and strict policies that govern how our data is handled.

APPLE Q4 2014 EARNING CALL TRANSCRIPT (October 20, 2014), <http://seekingalpha.com/article/2576865-apples-aapl-ceo-tim-cook-on-q4-2014-results-earnings-call-transcript>.

41. Vipin Samar, Vice President of database security product development at Oracle stated in a 2014 press release that, "As regulations worldwide increasingly call for more data to

be encrypted, organizations need a centralized solution to securely manage all the encryption keys and credential files in their data centers.” The press release continued by pointing out the importance of secure encryption in the cloud.

and backup mechanisms. As organizations increasingly encrypt data at rest and on the network, securely managing all the encryption keys and credential files in the data center has become a major challenge.

At the same time, organizations also need to comply with stringent regulatory requirements for managing keys and certificates. Many global regulations and industry standards call for audits demonstrating that keys are routinely rotated, properly destroyed, and accessed solely by authorized entities.

Oracle Customers Secure Critical Encryption Keys with Oracle Key Vault, ORACLE PRESS RELEASE (August 7, 2014).

42. Although secure third party encryption systems that protect access to data at an intermediary are offered by major corporations today, at the time the inventions disclosed in the STPC patents were conceived, no such systems existed.

43. The claims in the STPC patents describe a solution that is unquestionably rooted in computer technology to overcome a problem specific to and characteristic of complex computer networks. Professor of Computer Science at Columbia University, Steven M. Bellovin¹⁸ described in a 1996 academic article, contemporaneous to the development of the patents-in-suit (and cited on the face of the STPC patents) that the development of modern cryptography was a reaction to the rise of the Internet as a mass medium and concerns unique to the exchange of information over the Internet.

¹⁸ At the time, Professor Bellovin was a fellow at AT&T research labs.

In early 1994, CERT announced¹ that widespread password monitoring was occurring on the Internet. In 1995, Joncheray published a paper explaining how an eavesdropper could hijack a TCP connection [Jon95]. In mid-1998, there is still very little use of cryptography. Finally, though, there is some reason for optimism.

A number of factors have combined to change people's behavior. First, of course, there is the rise of the Internet as a mass medium, and along with it the rise of Internet commerce. Consider the following quote from a popular Web site:

Steven M. Bellovin, *Cryptography and the Internet*, AT&T LABS-RESEARCH PAPER (Aug. 1998).

44. Although encryption, in some form, has been an objective of individuals (and governments) for many years, the STPC patents are directed at solving problems that are unique to the realm of computers and specifically network cloud computing. "As we know, public cloud uses virtualization heavily as they share resources between many customers. As a result, this creates security vulnerabilities, both from access levels as well as from exploits in the virtualization software."¹⁹

45. The specific technologies disclosed and claimed in the STPC patents are discussed in detail below. However, the history of cryptography provides context for the inventions disclosed in the STPC patents and confirms that the patented inventions are limited to specific computer systems and methods addressing issues specific to modern computer networks.

46. ***Pre-Mechanical Encryption.*** The origin of cryptography has been around since the reign of Pharaohs; however, the problems that "pre-silicon" societies faced were markedly different than those the patents-in-suit are directed at solving. The unique solutions taught by the patents-in-suit reflect that difference. In 1900 BC, Egyptian scribes developed a rudimentary form of cryptography that allowed the passing of messages written on papyrus. The key to unlocking the meaning of non-standard hieroglyphs (the encrypted message or cipher) was located in an inscription on the same document. Thus, a recipient of a message could decipher the meaning of the encoded message using the key transmitted with the message. This early

¹⁹ Mohd Ujaley, *Cloud Adoption Requires Thorough Risk Assessment: Dell*, EXPRESS COMPUTER (May 26, 2015) (emphasis added) (the quote comes from an interview with Dell General Manager Murli Mohan).

form of encryption was susceptible to frequency analysis, a method utilizing the frequency that certain letters or symbols would be used.²⁰



Alexander Stanoyevitch, INTRODUCTION TO CRYPTOGRAPHY WITH MATHEMATICAL FOUNDATIONS AND COMPUTER IMPLEMENTATIONS PRESS (2002).

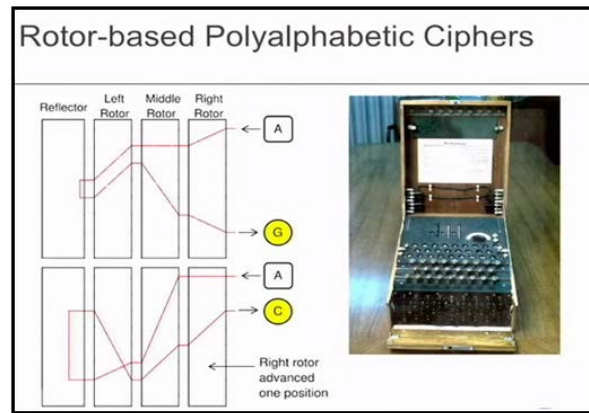
47. Over the following four millennia, the advance of cryptography was limited. In the mid-1400s, Leon Battista Alberti invented an encryption system using a mechanical device with sliding disks that allowed for various methods of substitution.²¹ This is the base concept of a polyalphabetic cipher, which is an encryption method that switches through several substitution ciphers throughout encryption. Polyalphabetic substitution by rotating the discs to change the encryption logic limited the use of frequency analysis to crack the cipher. However, polyalphabetic substitution was susceptible to plain text attacks that would try various permutations of the code.

48. ***Encryption in the Mechanical Age.*** In the 1920s, electro-mechanical devices were developed that used electrical signals to perform rudimentary calculations that would encrypt messages. The Enigma machine developed by the German government at the end of

²⁰ NIGEL SMART, CRYPTOGRAPHY: AN INTRODUCTION 3RD EDITION 40 (2004) ([U]nderlying statistics of the language could be used to break the cipher. For example it was easy to determine which ciphertext letter corresponded to the plaintext letter *E*.”).

²¹ DAVID KAHN, THE CODE BREAKERS: THE STORY OF SECRET WRITING 125 (1967) (David Kahn calls Alberti "the father of western cryptography" based on his development of a device that had two copper disks that fit together. “Each one of them had the alphabet inscribed on it. After every few words, the disks were rotated to change the encryption logic, thereby limiting the use of frequency analysis to crack the cipher.”).

World War I used mechanical devices to encrypt and decrypt messages. Germany's Enigma device used a set of codes that, when programed into a device, would generate an encrypted message. Ciphers generated by the Enigma could thus be decrypted if one had both possession of an Enigma device and the "crib" or the symmetric key that was used to program the device.²² Alan Turing (among others) wanted a technique to break Enigma that did not rely on the key, which could (and frequently did) change.²³ Turing developed several ways of using Bayesian inference coupled with "the Bombe," an electromechanical device that could detect the setting for the Enigma.



Steve Weis, THEORY AND PRACTICE OF CRYPTOGRAPHY 9:23 (November 2007) (image of the Enigma machine).

49. ***The Development of Public Key Encryption.*** Prior to 1976 (roughly three decades before the patents-in-suit issued), the only method of encryption was use of a symmetric key. Egyptian Ciphers, Polyalphabetic Encryption, and the Enigma Machine relied on a sender

²² DAVID KAHN, SEIZING THE ENIGMA: THE RACE TO BREAK THE GERMAN U-BOAT CODES, 1939-1943 (1991) (In 1941 the British were able to decrypt ciphers generated by the enigma machine by discovering that portions of weather reports (Short Weather Codes) transmitted by German Warships were the symmetric key. However, in the fall of 1941 the German cryptographers stopped using short Weather Codes as symmetric keys. Subsequently, Germany out of abundance of caution changed the configuration of the enigma machines.).

²³ DAVID LEAVITT, THE MAN WHO KNEW TOO MUCH: ALAN TURING AND THE INVENTION OF THE COMPUTER (2006) (Turing settled on a known plaintext attack, using what was known at the time as a "crib." A crib was a piece of plaintext that was suspected to lie in the given piece of cipher text. The methodology of this technique was to form a given piece of cipher text and a suspected piece of corresponding plaintext to first deduce a so-called "menu." A menu is simply a graph, which represents the various relationships between cipher text and plaintext letters. Then the menu was used to program an electrical device called a Bombe.).

and receiver sharing the same key (a symmetric key). The advent of computer networks and the increasing computational power of computers spurred the invention of a cryptographic system specifically tailored toward encrypting and decrypting electronic messages communicated using a computer.

50. In a 1976 paper, cited on the face of the STPC patents, Whitfield Diffie and Martin Hellman proposed the notion of *public-key* (frequently, and more generally, called *asymmetric key*) cryptography in which two different but mathematically related keys are used—a *public* key and a *private* key. Systems that utilize *public key* encryption were developed specifically to address problems unique to computer networking. Public key encryption at the time of the invention of the STPC patent technologies was not a long-held view, nor a technology that simply amounted to taking something and “doing it on a computer.” The introduction to Diffie and Hellman’s paper makes clear that public key systems were specific to computer networking.

This paper deals with new problems which arise in the application of cryptography to computer communication systems with large numbers of users. Foremost among these is the key distribution problem. We

Diffie, et al., in *Multiuser Cryptographic Techniques*, AFIPS--CONFERENCE PROCEEDINGS, Vol. 45 at 109 (1976).

51. A public key system contains two keys (numbers) so that calculation of one key (the 'private key') is computationally infeasible from the other (the 'public key'), even though they are necessarily related. Instead, both keys are generated secretly, as an interrelated pair. Public key encryption offered a novel mechanism for allowing two parties to share data over a network.

52. The development of Diffie and Hellman’s first public key system was directly motivated by the need to protect stored or transmitted data on a modern computer network.

In a computer network with a large number of users, cryptography is often essential for protecting stored or transmitted data. While this application closely resembles the age old use of cryptography to protect military and diplomatic communications, there are several important differences which require new protocols and new types of cryptosystems. This paper addresses the multiuser aspect of computer networks and presents ways to preserve privacy of communication despite the large number of user connections which are possible.

Id.

53. The Diffie-Hellman public key system illustrates the limitations present in systems for encrypting and decrypting information over a computer network contemporaneous to the STPC patents. The Diffie-Hellman system lacked the ability to enable the exchange of data between two parties through an intermediary where the intermediary would not have the ability to substantially decrypt the data. A 2005 paper (cited on the face of the STPC patents) described the limitations of the Diffie-Hellman system when conducting secure third party communications. The paper also described a problem that the STPC patents solve as one that had only recently been addressed:

It was only recently that the problem has been formally addressed in the three-party model, where the server is considered to be a trusted third party (TTP). This is the same scenario used in the popular 3-party Kerberos authentication system. The main advantage of these systems is that users are only required to remember a single password, the one they share with a trusted server, while still being able to establish secure sessions with many users. ***The main drawback is the need of the trusted server during the establishment of these secure sessions.***

Michel Abdalla and David Pointcheval, *Interactive Diffie-Hellman Assumptions With Applications To Password-Based Authentication*, in *PROCEEDINGS OF THE 9TH INTERNATIONAL CONFERENCE ON FINANCIAL CRYPTOGRAPHY AND DATA SECURITY* (2005) (emphasis added).

54. Another early encryption system developed for communications over a computer network is a method of public-key encryption developed by Ron Rivest, Adi Shamir, and Leonard M. Adleman, now generally referred to as “RSA.” RSA is based on the use of two extremely large prime numbers which fulfill the criteria for a “trap-door, one-way permutation.” Such a permutation function enables the sender to encrypt the message using a non-secret

encryption key, but does not permit an eavesdropper to decrypt the message through cryptanalytic techniques within an acceptable period of time. This is because, for a composite number composed of the product of two very large prime numbers, the computational time necessary to factor this composite number is unacceptably long. A brute force attack requires a sequence of putative keys to be tested to determine which, if any, is appropriate. A brute force attack requires a very large number of iterations. The number of iterations increases exponentially with the key bit size, while the normal decryption generally suffers only an arithmetic-type increase in computational complexity.

55. Like the Diffie-Hellman system, RSA was developed specifically to address problems with sending and receiving encrypted information over a computer network. The original RSA patent (cited on the face of the STPC and IRI patents) describes the use of public key encryption as directed toward a computer network.

With the development of computer technology, the transfer of information in digital form has rapidly increased. There are many applications, including electronic mail systems, bank systems and data processing systems, where the transferred information must pass over communications channels which may be monitored by electronic eavesdroppers.

U.S. Patent No. 4,405,829, col. 1:14-20.

56. Academic articles from creators of the RSA system make clear that the use of public key encryption is specific to problems unique to computer networks.

[W]e present a sketch of how a computer system might be modified to solve the problem of performing operations on encrypted data securely. . . All sensitive data in main memory, in the data bank files, in the ordinary register set, and on the communications channel will be encrypted. During operation, a load/store instruction between main memory and the secure register set will automatically cause the appropriate decryption/encryption operations to be performed.

Ronald L. Rivest, Leonard Adleman, and Michael L. Dertouzos, *On Data Banks and Privacy Homomorphisms*, in ON DATA BANKS AND PRIVACY HOMOMORPHISMS 169 (1978).

57. The RSA system illustrates limitations in encryption technologies that preceded the STPC patents. RSA provided a mechanism for exchanging data between two parties but did not disclose the use of an untrusted intermediary when data was exchanged between two parties. A 1998 article contemporaneous to the development of the STPC patents (and cited on the face

of the STPC patents) describes this as a limitation in the RSA system and other systems known at the time.

We point out that classic techniques of secret sharing [14] are inadequate in this scenario. Secret sharing requires one to reconstruct the secret at a single location before it can be used, hence introducing a single point of failure. The technique described above of sharing the secret key such that it can be used without reconstruction at a single location is known as *Threshold Cryptography*. See [9] for a succinct survey of these ideas and nontrivial problems associated with them.

An important question left out of the above discussion is key generation. Who generates the RSA modulus N and the shares d_1, d_2, d_3 ? Previously the answer

D. Boneh, J. Horwitz, *Generating A Product Of Three Primes With An Unknown Factorization*, in PROC. OF THE THIRD ALGORITHMIC NUMBER THEORY SYMPOSIUM 237 (1998).

58. Silvio Micali's patents (U.S. Pat. Nos. 6,026,163 and 5,315,658; cited on the face of the STPC patents) describe a split key, or so-called "fair" cryptosystem, designed to allow a secret key to be distributed to a plurality of trusted entities, such that the encrypted message is protected unless the key portions are divulged by all of the trusted entities. Thus, a secret key may be recovered through cooperation of a plurality of parties. The Micali system provides that the decryption key is split between a number (n) of trusted entities, meeting the following functional criteria: (1) The private key can be reconstructed given knowledge of all n of the pieces held by the plurality of trusted entities; (2) The private key cannot be guessed at all if one only knows less than all ($<n-1$) of the special pieces; and (3) For $i-1, \dots, n$, the i^{th} special piece can be individually verified to be correct.

59. The Micali system does not allow communication of a secret in which only one party gains access to the content, and in which the third party or parties and one principal operate only on encrypted or secret information.

B. The Value Of The Inventions Disclosed In The STPC Patents

60. Executives at leading technology companies have described the value of specific encryption techniques as critical, lasting, and prominent. Chris Cicotte, a Cloud Architect at EMC, stated strong encryption technologies specific for networked computers "are a vital component of a strong security posture for any size organization, and it should be a standard

offering within the cloud The threat landscape has already begun to evolve, and from an overall security perspective, we need to take a proactive approach by layering in technologies like encryption at every layer."²⁴ The development of secure communications systems and methods, such as the inventions taught in the STPC patents, was motivated by the unique problems created by the internet where secured data is often transmitted through untrusted intermediaries.

Achieving secure communications in networks has been one of the most important problems in information technology. . . . If there is a private and authenticated channel between two parties, then secure communication between them is guaranteed. However, in most cases, many parties are only indirectly connected, as elements of an incomplete network of private and authenticated channels. ***In other words they need to use intermediate or internal nodes.***

Yvo Desmedt and Yongee Wang, *Perfectly Secure Message Transmission Revisited* at 502, *Advances in Cryptology EUROCRYPT Vol. 2332* (2002) (emphasis added).

61. Companies such as Defendant Google, Oracle Corporation, International Business Machines Corporation, and Hewlett-Packard Company confirm the importance of providing strong encryption systems that address the unique threats posed by moving data to the cloud.

Once data is moved to the cloud, ***it becomes vulnerable to a number of new threats*** ranging from stolen administrator credentials to new hacking techniques. In addition, new legislation, such as the USA PATRIOT Act, is making it possible for competitors and governments to access data from cloud providers without the consent of the data owner. Many cloud providers thought they could achieve data sovereignty through locating cloud services in different jurisdictions, but this theory has been shaken by the subpoena classification ruling handed down recently in the U.S. federal court.

HP Atalla Cloud Encryption: Securing Data in the Cloud, HP TECHNICAL WHITE PAPER 2 (2014) (emphasis added).

²⁴ Jude Chao, *Cloud Computing Demands Cloud Data Encryption*, ENTERPRISE NETWORKING PLANET WEBSITE, May 13, 2014, <http://www.enterprisenetworkingplanet.com/netsecur/cloud-computing-demands-cloud-data-encryption.html>.

The need to secure data is driven by an expanding privacy and regulatory environment coupled with an increasingly dangerous world of hackers, insider threats, organized crime, and other groups intent on stealing valuable data. ***The security picture is complicated even more by the rapid expansion of access to sensitive data via the Internet.*** an unprecedented understanding of technology, increasing economic competition, and the push to achieve greater efficiencies through consolidation and cloud computing.

Oracle Database 12C Security and Compliance, ORACLE WHITE PAPER 2 (February 2015) (emphasis added).

With rare exceptions, one of the most important assets for any company is its data. Your data may take the form of financial information, proprietary sales information, marketing information, healthcare information, intellectual property (IP), and more. Losing your data could negatively affect operations and potentially shut down your organization. . . . Cloud-aware applications create unique security challenges in that both Infrastructure as a Service (IaaS) providers and Platform as a Service (PaaS) providers make use of a shared-risk model.

Robi Sen, *Develop Secure Cloud-Aware Applications*, IBM DEVELOPER WORKS 2-3 (May 20, 2015).

Business requirements, industry regulations, and government mandates increasingly dictate that your organization must secure electronic communications. Whether it is financial data, medical records, or proprietary corporate information, you simply must secure the delivery of sensitive content to its destination.

Google Message Encryption, GOOGLE APPLICATION SECURITY PAPER 1 (2008).

62. Numerous academics have concluded the advent of cloud computing has created challenges that are unique to cloud computing and these challenges require specific encryption technologies that were previously unnecessary.

The growing demand for cloud computing stems from the need to securely store, manage, share and analyze immense amounts of complex data in many areas, including health care, national security and alternative energy. And although several companies have launched commercially available cloud systems, two areas still need significant improvements. [Dr. Bhavani] Thuraisingham said: the security mechanisms needed to protect sensitive data as well as the capability to process huge amounts of both geospatial data and what's known as semantic Web data.

Investment in Cloud Computing Research Pays Off, UT Dallas Computer Scientists Make Advances in Key Aspects of Growing Field, UNIVERSITY OF TEXAS AT DALLAS NEWS CENTER (April 19, 2011).²⁵

²⁵ See also Kevin Hamlen et al., *Security Issues For Cloud Computing* at 39, INTERNATIONAL JOURNAL OF INFORMATION SECURITY AND PRIVACY Vol. 4(2) (April-June 2010) (“Because of the critical nature of the applications, it is important that clouds be secure. The major security challenge with clouds is that the owner of the data may not have control of where the data is placed.”); Ryan Layfield, Murat Kantarcioglu, and Bhavani Thuraisingham, *Enforcing Honesty in Assured Information Sharing within a Distributed System*, IFIP WG 11.3 CONFERENCE ON

Security is the most important challenge for cloud technology, as CSP's [Cloud Service Providers] have to protect the consumer's data from theft and ensure the consumer is not exploited. Consumers may be exploited from denial of service (DoS) attacks . . . ***They must also protect the data through the use of advanced encryption algorithms*** and ensure that their data centers are physically secure using advanced biometrics and many other authentication methods.

Sean Carlin & Kevin Curran, *Cloud Computing Technologies*, in INTERNATIONAL JOURNAL OF CLOUD COMPUTING AND SERVICES SCIENCE (IJ-CLOSER) Vol.1, No.2 at 59 (June 2012) (emphasis added).

The growth of computer networks and the opening that their interconnection brings, especially through Internet, mean that a great amount of information is traveling through network and ***crossing numerous intermediate systems. This results in the increase of the number of possible attacks and illegal operations.*** . . . They should guarantee the identity of the communicating parties . . . the protection against unauthorized writing and, in some cases, unauthorized reading of transferred data. These services of authentication, nonrepudiation, integrity and confidentiality, respectively, can be provided using cryptosystems.

Natasha Prohic, *Public Key Infrastructures - PGP vs. X.509* at 1, in INFOTECH SEMINAR ADVANCED COMMUNICATION SERVICES (ACS) (2005) (emphasis added).

63. On information and belief, contemporaneous to, and following conception of the inventions disclosed in the STPC patents, academics, and businesses headquartered in Texas actively entered the field of secure encrypted communications. Computer researchers at the University of Texas at Austin founded the Security Research Group. The University of Texas at Dallas founded the Data Security and Privacy Lab, a center for research on security issues raised by dissemination of data over computer networks.

64. Texas based companies incorporated secure communications technologies into numerous products and many of these same companies cite STPC patents in their own patents. Texas based businesses that developed products incorporating secure communications technologies included: HP Enterprise Services, LLC of Plano, Texas; Texas Instruments, Inc. of Dallas, Texas; Rocksteady Technologies, LLC of Austin, Texas; Dell, Inc. of Round Rock,

DATABASE AND APPLICATIONS SECURITY (2007) ("The growing number of distributed information systems such as the internet has created a need for security in data sharing."); Safwan M. Khan and Kevin W. Hamlen, *AnonymousCloud: A Data Ownership Privacy Provider Framework in Cloud Computing* at 170, in PROCEEDINGS OF THE 11TH IEEE INTERNATIONAL CONFERENCE ON TRUST, SECURITY AND PRIVACY IN COMPUTING AND COMMUNICATIONS (June 2012) ("Revolutionary advances in hardware, middleware, and virtual machines over the past few years have elevated cloud computing to a thriving industry . . . A significant barrier to the adoption of cloud services is customer fear of privacy loss in the cloud.").

Texas; AT&T Intellectual Property whose inventors were based in various locations in Texas; Gazzang, Inc. of Austin Texas; Net.Orange, Inc. of Dallas, Texas; and Futurewei Technologies, Inc. of Plano, Texas. The STPC patents are cited by at least 50 patents that were either initially assigned to or are currently assigned to entities headquartered in Texas.

1. U.S. Patent No. 8,316,237

65. U.S. Patent No. 8,316,237 (the “237 patent”) entitled, System and Method for Secure Three-Party Communications, was filed on January 10, 2011 and claims priority to March 23, 2001. St. Luke is the owner by assignment of the ‘237 patent. A true and correct copy of the ‘237 patent is attached hereto as Exhibit A. The ‘237 patent claims specific methods and systems for securely transcribing protected electronic information transmitted over at least one computer network from a first encrypted form to a second, different encrypted form substantially without intermediate decryption of the protected electronic information.

66. The ‘237 patent has been cited by over 100 issued United States patents as relevant prior art. Specifically, patents issued to the following companies have cited the ‘237 patent as relevant prior art.

- Electronics and Telecommunications Research Institute (ETRI)
- NEC Corporation
- Disney Enterprises, Inc.
- WMS Gaming, Inc.
- Verizon Patent and Licensing, Inc.
- Microsoft Corporation.
- NetApp, Inc.
- NCR Corporation
- EMC Corporation
- AT&T Intellectual Property, L.P.
- Sony Corporation
- SAP AG
- Blackberry Limited
- Adobe Systems Incorporated
- Nippon Telegraph and Telephone Corporation
- Novell, Inc.
- Spring Communications L.P.
- Hytrust, Inc.
- International Business Machines Corporation
- **Google, Inc.**
- Kabushiki Kaisha Toshiba
- Panasonic Intellectual Property Management Co., Ltd.
- Zynga, Inc.

- Certicom Corp.
- Wincor Nixdorf International GmbH
- Oracle International Corporation
- Futurewei Technologies, Inc.
- Dell Products, L.P.
- Intuit, Inc.

67. The '237 patent claims a technical solution to a problem unique to computer networks – securely transmitting encrypted electronic information via an intermediary device, wherein the electronic information is cryptographically secure not only from outside attackers, but also from the intermediary.

68. At the time of the inventions claimed in the '237 patent, securely processing, transmitting, and accessing protected electronic data in a massively distributed computing environment presented new and unique issues over the state of the art. As explained in the '237 patent: “Often, the nature of these communications protocols places the third party (or group of third parties) in a position of trust, meaning that the third party or parties, without access to additional information, can gain access to private communications or otherwise undermine transactional security or privacy.” '237 patent, col. 2:13-17.

Generating and protecting encryption keys while maintaining data availability has traditionally been a major barrier to implementing encryption, especially on an enterprise scale. Key management is complex and challenging, and often fails because issuance, storage, and renewing are difficult. ***Worse yet, lost keys can make important data permanently unrecoverable.***

Sustainable Compliance for the Payment Card Industry Data Security Standard, ORACLE WHITE PAPER 23 (July 2015) (emphasis added).

69. Although the systems and methods taught in the '237 patent have been adopted by leading businesses today, at the time of invention, the technologies taught in the '237 patent claims were innovative and novel. “Typical public key encryption technologies, however, presume that a pair of communications partners seek to communicate directly between each other, without the optional or mandatory participation of a third party, and, in fact, are designed specifically to exclude third party monitoring.” '237 patent, col. 2:56-61. Indeed, companies such as Google competitor Oracle recognized that, until recently, security for distributed systems was not a primary concern.

- Security was not a major issue, even for Oracle
 - Standard passwords (scott/tiger, system/manager, ...)
 - Oracle standard users were installed and left open (though not at SAP!)
 - There are some recommendations, but not much more.
 - From Oracle9i, the issue of security was increasingly addressed by Oracle (DBCA: locking of default accounts, ..., 10.2: CONNECT roles)

Andreas Becker, *High Security for SAP Data with Oracle Database Vault and Transparent Data Encryption*, ORACLE PRESENTATION 6 (2010).

70. Further, the '237 patent claims improve upon the functioning of a computer system by allowing encrypted electronic data to be securely transmitted through an intermediary without the intermediary gaining substantial access to the unencrypted information. This improves the security of the computer system and allows it to be more efficient.²⁶ “Third parties, however, may offer valuable services to the participants in a communication, but existing protocols for involvement of more than two parties are either inefficient or insecure.” ‘237 patent, col. 2:61-64. Studies have confirmed that the inventions disclosed in the ‘237 patent improve the security of systems.

Key management is a big concern with encryption, because the effectiveness of the solution ultimately depends on protecting the key. If the key is exposed, the data being protected with the key is, essentially, exposed. Wherever the key is stored, it must be protected, and it should be changed on occasion. For example, if an administrator with access to a key leaves an organization, the key should be changed.

Tanya Baccam, *Transparent Data Encryption: New Technologies and Best Practices for Database Encryption*, SANS WHITE PAPER 3 (April 2010) (emphasis added).

²⁶ See Kevin Hamlen et al., *Security Issues For Cloud Computing* at 39, INTERNATIONAL JOURNAL OF INFORMATION SECURITY AND PRIVACY VOL. 4(2) (April-June 2010) (“The major security challenge with clouds is that the owner of the data may not have control of where the data is placed. . . . Therefore, we need to safeguard the data in the midst of untrusted processes.”); Elena Ferrari and Bhavani Thuraisingham, *Security and Privacy for Web Databases and Services* at 17, PROCEEDINGS OF THE EDBT CONFERENCE (March 2003) (“very little work has been devoted to security”); Elisa Bertino et al., *Selective and Authentic Third-Party Distribution of XML Documents* at 1263, IEEE TRANSACTIONS ON KNOWLEDGE AND DATA ENGINEERING, Vol. 16 No. 10 (October 2004) (“The most intuitive solution is that of requiring Publishers to be trusted with regard to the considered security properties. However, this solution could not always be feasible in the Web environment since large Web-based systems cannot be easily verified to be secure.”).

71. The '237 patent claims are not directed to a “method of organizing human activity,” “fundamental economic practice long prevalent in our system of commerce,” or “a building block of the modern economy.” Instead, they are limited to a concretely circumscribed set of methods and systems for transcribing electronic information that is transmitted over a computer network via an intermediary.

72. The '237 patent claims are not directed at the broad concept/idea of “encrypting” or “decrypting” information. Instead, they are limited to a concretely circumscribed set of methods and systems for transcribing electronic information that is transmitted over a computer network via an intermediary. These methods and systems are technologies unique to the Internet age.

73. The inventive concepts claimed in the '237 patent are technological, not “entrepreneurial.” For example, transcribing protected electronic information between a first (e.g., server) encrypted form and a second (e.g., network) encrypted form without a substantial intermediate representation of the information in decrypted form is a specific, concrete solution to the technological problem of transferring encrypted information via an intermediary without providing the intermediary substantial access to the information.

74. Researchers have identified the problems the '237 patent is directed at solving arise from new security challenges relating to cloud computing.

Data Security: Data security was the most important concern that had hindered the acceptance of the cloud computing initially. Storing and processing the data, running software, using CPU and virtual Machines on others' infrastructure were some serious concerns for the users initially. Data breach, data integrity and data loss are major security issues that posed threats to organization's data and software. Moreover, the multi-tenancy model and pooled computing resources over cloud have introduced new security challenges requiring new techniques to tackle with [4] [5] [6].

Deepak Panth, Dhananjay Mehta and Rituparna Shelgaonkar, *A Survey on Security Mechanisms of Leading Cloud Service Providers*, in INTERNATIONAL JOURNAL OF COMPUTER APPLICATIONS 98(1) at 34 (July 2014) (emphasis added).²⁷

75. The '237 patent claims are directed toward a solution rooted in computer technology and use technology unique to computers and computer networking to overcome a problem specifically arising in the realm of secure distributed computing. For example, claims of the '237 patent require transcribing protected electronic information using one or more intermediary computing devices specially configured to yield a desired result—a result that overrides the routine and conventional sequence of events in electronic communications, even encrypted electronic communications.

76. The '237 patent is directed to specific problems in the field of cryptography. In the “Background” section of the patent, the '237 patent explains that encryption systems use “keys,” similar to passwords, to control how plaintext is encrypted and decrypted. '237 patent, col. 2:65–3:13. An encryption system thereby encrypts and decrypts information differently depending upon the key input. *Id.* Two common cryptanalytic attacks, linear and differential cryptanalysis, analyze large amounts of cipher text (encrypted information) and different

²⁷ See also Vaibhav Khadilkar, Murat Kantarcioglu, and Bhavani Thuraisingham, *Secure Data Processing in a Hybrid Cloud* at 1-2, Computing Research Repository (CoRR) abs/1105.1982 (2011) (“The emergence of cloud computing has created a paradigm shift by allowing parallel processing of massive amounts of data. . . . [H]ow do users protect themselves from cloud service providers who may be able to access their data? This issue is related to data security and is relevant for users since their data is placed at the provider’s site.”).

possible keys in order to eventually converge on the correct key and break the encryption. *Id.* at col. 3:1–3:13. Both attacks exploit the fact that some encryption systems use static keys to create the cipher text. *Id.* In other words, using the same key repeatedly gives an attacker more information to work with. The inventions of the '237 patent introduce several novel techniques to overcome these weaknesses and allow encrypted information to be securely transferred through an intermediary.

77. The preemptive effect of the claims of the '237 patent are concretely circumscribed by specific limitations. For example, claim 1 of the '237 patent requires:

A transcription device, comprising:

an automated communication port configured to receive a first message representing an encrypted communication associated with a first set of asymmetric keys, to receive a transcription key, and to transmit a second message representing the encrypted communication associated with a second set of asymmetric keys, the first and second sets of encryption keys being distinct;

a memory; and

an automated processor, configured to communicate through the automated communication port and with the memory, to receive the first message, receive the transcription key, automatically transcribe the first message into the second message, and to transmit the second message, wherein the automated processor does not store as a part of the transcription any decrypted representation of the encrypted communication, and the transcription key is employed without revealing any secret cryptographic information usable for decrypting the first message or the second message.

78. The '237 patent does not attempt to preempt every application of the idea of encrypting electronic information transmitted over a computer network, or even the idea of encrypting electronic information transmitted over a computer network via an intermediary.

79. The '237 patent does not preempt the field of secure third-party communications systems, or prevent use of alternative secure third-party communications systems. For example, the '237 patent includes inventive elements—embodied in specific claim limitations—that concretely circumscribe the patented invention and greatly limit its breadth. These inventive

elements are not necessary or obvious tools for achieving secure third-party communications, and they ensure that the claims do not preempt other techniques for secure communications.

80. For example, the '237 patent describes numerous techniques for secure third-party communications that inform the invention's development but do not, standing alone, fall within the scope of its claims:

- Key Escrow. U.S. Pat. No. 6,009,177 to Sudia, relates to a cryptographic system and method with a key escrow feature that uses a method for verifiably splitting users' private encryption keys into components and for sending those components to trusted agents chosen by the particular users.
- Partitioning of Information Storage Systems. U.S. Patent No. 5,956,400 to Chaum, relates to partitioned information storage systems with controlled retrieval.
- Use of a Trusted Intermediary. U.S. Patent No. 6,161,181 to Haynes, describing secure electronic transactions using a trusted Intermediary; U.S. Patent No. 6,145,079 to Misty, describing secure electronic transactions using a trusted intermediary to perform electronic services.
- Split Key Storage. U.S. Patent No. 6,118,874 to Okamoto, teaching encrypted data using split storage key and system.
- Use of a Cryptographic File Labeling System. U.S. Pat. No. 5,953,419 to Lohstroh, disclosing cryptographic file labeling system for supporting secured access by multiple users.
- Computer Security Devices. U.S. Pat. No. 5,982,520 to Weiser, disclosing a personal storage device for receipt, storage, and transfer of digital information to other electronic devices; *see also* U.S. Pat. No. 5,991,519 to Benhammou; U.S. Pat. No. 5,999,629 to Heer; and U.S. Pat. No. 6,034,618 to Tatebayashi.
- Computer Network Firewalls And Agents. U.S. Pat. No. 6,061,798 to Coley, disclosed the use of an assigned proxy agent to verify the authority of an incoming request to access a network element indicated in the request. Once verified, the proxy agent completes the connection to the protected network element on behalf of the source of the incoming request; *see also* U.S. Pat. No. 6,023,762 to Dean, disclosing a data access and retrieval system which comprises a plurality of user data sources each storing electronic data signals describing data specific to a user, or enabling services selected by a user; an agent device which is configurable to select individual ones of the user data sources and

present selections of user data and service data to a set of callers who may interrogate the agent device remotely over a communications network; and U.S. Pat. No. 6,029,150 to Kravitz, disclosing a system and method of payment in an electronic payment system wherein a plurality of customers have accounts with an agent. Further, the patent lists thirty-three other patented systems involving Computer Network Firewalls that are not, standing alone, preempted by the inventions claimed in the patents-in-suit.

- Virtual Private Networks. As described in: U.S. Pat. No. 6,079,020 to Liu and U.S. Pat. No. 6,081,900 and twenty other patented systems involving virtual private networks that are not, standing alone, preempted by the inventions claimed in the patents-in-suit.
- Biometric Authentication. U.S. Pat. No. 5,193,855 to Shamos, disclosing the use of biometrics such as fingerprints to facilitate secure communications and identification of users. Further, the '237 lists 238 patented systems that use biometric authentication that are not, standing alone, preempted by the inventions claimed in the patents-in-suit.

81. Although “[e]ncryption, in general, represents a basic building block of human ingenuity that has been used for hundreds, if not thousands, of years,”²⁸ the ‘237 patent does not claim, or attempt to preempt, “some process that involves the encryption of data for some purpose” (or similar abstraction).

82. The ‘237 patent does not claim, or attempt to preempt, the performance of an abstract business practice on the Internet or using a conventional computer.

83. The claimed subject matter of the ‘237 patent is not a pre-existing but undiscovered algorithm.

84. The ‘237 patent claims systems and methods that “could not conceivably be performed in the human mind or pencil and paper.”²⁹

²⁸ *Paone v. Broadcom Corp.*, Case No. 15 CIV. 0596 BMC GRB, 2015 WL 4988279, at *7 (E.D.N.Y. Aug. 19, 2015) (citing *Fid. Nat'l Info. Servs., Inc.*, Petitioner, CBM2014-00021, 2015 WL 1967328, at *8 (Apr. 29, 2015) (both upholding the patent eligibility of patents directed toward encryption).

²⁹ *TQP Dev., LLC v. Intuit Inc.*, Case No. 2:12-CV-180-WCB, 2014 WL 651935, at *4 (E.D. Tex. Feb. 19, 2014) (finding claims directed to encryption to be patent eligible); *Paone v. Broadcom Corp.*, Case No. 15 CIV. 0596 BMC GRB, 2015 WL 4988279, at *7 (E.D.N.Y. Aug. 19, 2015); see also *Prism Technologies, LLC v. T-Mobile USA, Inc.*, 12-cv-124, Dkt. No. 428 at 7 (D. Neb. Sept. 22, 2015) (Finding on cross motions for summary judgment that patents directed at delivering resources over an untrusted network were patent eligible. “The problems addressed by Prism’s claims are ones that ‘arose uniquely in the context of the Internet.’”).

85. The '237 patent claims require the use of a computer system.

86. The claims in the '237 patent require the modifying of data that has concrete and valuable effects in the field of secure third-party communications. By allowing an intermediary to receive secure information but not gain access to the unencrypted form of the information, the '237 patent improves the security of computer systems. Prior art systems that the '237 patent remedies enabled unauthorized "access to private communications or otherwise undermine[d] transactional security or privacy." Companies have described the use of encryption in the cloud as important to improve the security and functioning of systems.

For many organizations, keeping data private and secure has also become a compliance requirement. Standards including Health Insurance Portability and Accountability Act of 1996 (HIPAA), Sarbanes-Oxley (SOX), Payment Card Industry Data Security Standard (PCI DSS), the Gramm-Leach-Bliley Act, and EU Data Protection Directives all **require that organizations protect their data at rest and provide defenses against threats.**

HP Atalla Cloud Encryption: Securing Data in the Cloud, HP TECHNICAL WHITE PAPER 2 (2014) (emphasis added).

87. The '237 patent claims systems and methods not merely for transferring secure information over a computer network, but for making the computer network itself more secure.³⁰

88. The claimed invention in the '237 claims is rooted in computer technology and overcomes problems specifically arising in the realm of computer networks.

89. The systems and methods claimed in the '237 patent were not a longstanding or fundamental economic practice at the time of the patented inventions. Nor were they fundamental principles in ubiquitous use on the Internet or computers in general. As just one

³⁰ Limitations in the prior art that the '237 patent was directed to solving included: computer systems where a "third party plays a requisite role in the transaction but which need not be trusted with access to the information or the cryptographic key" (*Id.*, col. 2:5-7); "[p]asswords may be written near access terminals (*Id.* col. 1:50-51);" "[s]ecurity tokens can be stolen or misplaced" (*Id.*, col. 1:51-52); "users may share supposedly secret information" (*Id.*, col. 1:52); and "unauthorized uses of the system" (*Id.*, col. 11:28). The '237 patent "allows the entity that transmits the information to be assured that the transmission will be secure, even with respect to a trusted third party, while ensuring that the intended recipient must cooperate with the intended third party." '237 patent, col. 8:48-52.

example, at the time the inventions disclosed in the '237 patent were conceived, the use of asymmetric encryption keys was described by Oracle as "relatively new."³¹

A Public Key Infrastructure (PKI) consists of protocols, services, and standards supporting applications of public key cryptography. ***Because the technology is still relatively new***, the term PKI is somewhat loosely defined.

Introduction to the SSL Technology, ORACLE DOCUMENTATION (February 1, 2001), http://docs.oracle.com/cd/E53645_01/tuxedo/docs12cr2/security/publickey.html (emphasis added).

90. The asserted claims do not involve a method of doing business that happens to be implemented on a computer; instead, the '237 patent teaches changing data in a way that will affect the communication system itself, by making it more secure. The security challenges that the '237 patent is directed at overcoming were new and unique to distributed networks, as confirmed in a recent paper from Accenture Services Pvt. Ltd.: "The unprecedented growth of cloud computing has created new security challenges. The problem is ever more complex as there is a transition from traditional computing to a service-based computing."³²

91. The '237 patent claims are not directed at a mathematical relationship or formula. The '237 patent claims concrete, specific computer systems and methods for cryptographically protecting and managing access to secure data in multi-party communications.

92. '237 patent claims transform data from one form into another that will be recognizable by the intended recipient but secure against decryption by unintended recipients.


93. IBM in its computer reference guides ("redbooks") refers to encryption as "transform[ing] data that is unprotected."

³¹ See also *BackupEDGE Encryption Whitepaper*, MICROLITE CORPORATION at 2 (2003) (describing the technology of asymmetric keys as "new"); Roger Clarke, MESSAGE TRANSMISSION SECURITY (May 1998), <http://www.rogerclarke.com/II/CryptoSecy.html> ("Public key cryptography is relatively new and technically complex.").

³² Deepak Panth, Dhananjay Mehta and Rituparna Shelgaonkar, *A Survey on Security Mechanisms of Leading Cloud Service Providers*, in INTERNATIONAL JOURNAL OF COMPUTER APPLICATIONS 98(1) at 34 (July 2014).


Encryption concepts and terminology

Encryption transforms data that is unprotected, or *plain text*, into encrypted data, or *ciphertext*, by using a *key*. Without knowledge of the encryption key, the ciphertext cannot be converted back to plain text.



Bertrand Dufrasne and Robert Tondini, IBM DS8870 DISK ENCRYPTION 6th Edition at 4 (2015) (from a reference guide published by IBM).

94. One or more claims of the '237 patent require a specific configuration of electronic devices, a network configuration, and the use of encryption systems to secure communications from access by an intermediary. These are meaningful limitations that tie the claimed methods and systems to specific machines. For example, the below diagram from the '237 patent illustrates a specific configuration of hardware disclosed in the patent.



'237 patent, Fig. 1.

2. U.S. Patent No. 7,181,017

95. U.S. Patent No. 7,181,017 (the "'017 patent") entitled, System and Method for Secure Three-Party Communications, was filed on March 25, 2002, and claims priority to March 23, 2001. St. Luke is the owner by assignment of the '017 patent. A true and correct copy of the

'017 patent is attached hereto as Exhibit B. The '017 patent claims specific methods and systems for secure third-party communications—for example, a system and method for communicating information between a first party and a second party that includes identifying desired information; negotiating, through an intermediary, a cryptographic comprehension function for obscuring at least a portion of the information communicated between the first party and the second party; communicating the encrypted information to the second party, and decrypting the encrypted information using the negotiated cryptographic comprehension function. Moreover, in the patented systems and methods, the intermediary does not itself possess sufficient information to decrypt the encrypted information, thus allowing use of an “untrusted” intermediary.

96. The '017 patent has been cited by over 350 issued United States patents as relevant prior art. Specifically, patents issued to the following companies have cited the '017 patent.

- Electronics and Telecommunications Research Institute (ETRI)
- Sharp Laboratories of America, Inc.
- International Business Machines Corporation
- Microsoft Corporation
- Sony Corporation
- France Telecom
- Siemens Medical Solutions USA, Inc.
- Canon Kabushiki Kaisha
- Nikon Corporation
- Apple, Inc.
- Fujitsu Limited
- Hewlett-Packard Development Company, L.P.
- SAP AG
- Guardian Data Storage, LLC
- Teradata US, Inc.
- AT&T Intellectual Property I, L.P.
- Panasonic Corporation
- Sharp Laboratories of America, Inc.
- Ricoh Company, Ltd.
- Nokia Corporation
- Boss Logic, LLC
- Juniper Networks, Inc.
- American Express Travel Related Services Company, Inc.
- Kyocera Mita Corporation
- Oracle International Corporation
- Medox Exchange, Inc.
- Nortel Networks Limited

- Hitachi-Omron Terminal Solutions, Corporation
- Medapps, Inc.
- Samsung Electronics Co., Ltd.
- NEC Corporation
- Visa International Service Corporation
- Cisco Technology, Inc.
- Yahoo!, Inc.
- Flexera Software LLC
- CompuGroup Medical AG
- Datcard Systems, Inc.
- Futurewei Technologies, Inc.
- Telecom Italia S.P.A.
- General Electric Company
- Fuji Xerox Co., Ltd.
- Massachusetts Institute Of Technology
- NetApp, Inc.
- Koninklijke Philips N.V.
- Computer Associates Think, Inc.
- Huawei Technologies Co., Ltd.
- Texas Instruments, Inc.
- Nippon Telegraph and Telephone Corporation
- Research in Motion Limited.
- Net.Orange, Inc.
- Nokia Siemens Networks Oy
- Honeywell Int., Inc.

97. The claims in the '017 patent are directed at a technical solution to a problem unique to computer networks – securely transmitting encrypted electronic information via an intermediary device, wherein the electronic information is cryptographically secure not only from outside attackers, but also from the intermediary.

98. At the time of the inventions claimed in the '017 patent, securely processing, transmitting, and accessing protected electronic data in a massively distributed computing environment presented new and unique issues over the state of the art. As explained in the '017 patent: “Often, the nature of these communications protocols places the third party (or group of third parties) in a position of trust, meaning that the third party or parties, without access to additional information, can gain access to private communications or otherwise undermine transactional security or privacy.” '017 patent, col. 1:54-61.

Generating and protecting encryption keys while maintaining data availability has traditionally been a major barrier to implementing encryption, especially on an enterprise scale. Key management is complex and challenging, and often fails because issuance, storage, and renewing are difficult. ***Worse yet, lost keys can make important data permanently unrecoverable.***

Sustainable Compliance for the Payment Card Industry Data Security Standard, ORACLE WHITE PAPER 23 (July 2015) (emphasis added).

99. Although the systems and methods taught in the '017 patent have been adopted by leading businesses today, at the time of invention, the technologies taught in the '017 patent claims were innovative and novel. "Typical public key encryption technologies, however, presume that a pair of communications partners seek to communicate directly between each other, without the optional or mandatory participation of a third party, and, in fact, are designed specifically to exclude third party monitoring." '017 patent, col. 4:40-45. As described in an article contemporaneous to the '017 patent, the rise of cloud computing and distributed networks gave rise to a need to use key encryption to resolve security issues.

stored or communicated. As information becomes increasingly mobile, moving rapidly from application to application and system to system, this feature becomes more and more desirable. Public-key schemes are scalable: their operation is well-suited to environments with lots of users. The advent of large-scale open networks like the Internet necessitates this property.

Simon Blake-Wilson, *Information Security, Mathematics and Public-Key Cryptography*, in *Designs, Codes and Cryptography* Vol. 19 at 81 (2000).

100. Further, the '017 patent claims improve upon the functioning of a computer system by allowing encrypted electronic data to be securely transmitted through an intermediary without the intermediary gaining access to the unencrypted information. This improves the security of the computer system and allows it to be more efficient.³³ "Third parties, however,

³³ See Kevin Hamlen et al., *Security Issues For Cloud Computing* at 39, INTERNATIONAL JOURNAL OF INFORMATION SECURITY AND PRIVACY VOL. 4(2) (April-June 2010) ("The major security challenge with clouds is that the owner of the data may not have control of where the data is placed. . . . Therefore, we need to safeguard the data in the midst of untrusted processes."); Elena Ferrari and Bhavani Thuraisingham, *Security and Privacy for Web Databases and Services* at 17, PROCEEDINGS OF THE EDBT CONFERENCE (March 2003) ("very little work has been devoted to security"); Elisa Bertino et al.; *Selective and Authentic Third-Party Distribution of XML Documents* at 1263, IEEE TRANSACTIONS ON KNOWLEDGE AND DATA ENGINEERING, Vol. 16 No. 10 (October 2004) ("The most intuitive solution is that of requiring Publishers to be trusted with regard to the considered security properties. However,

may offer valuable services to the participants in a communication, but existing protocols for involvement of more than two parties are either inefficient or insecure.” ‘017 patent, col. 4:45-48. Studies have confirmed that the inventions disclosed in the ‘017 patent improve the security of systems.

Key management is a big concern with encryption, because the effectiveness of the solution ultimately depends on protecting the key. If the key is exposed, the data being protected with the key is, essentially, exposed. Wherever the key is stored, it must be protected, and it should be changed on occasion. For example, if an administrator with access to a key leaves an organization, the key should be changed.

Tanya Baccam, *Transparent Data Encryption: New Technologies and Best Practices for Database Encryption*, SANS WHITE PAPER 3 (April 2010) (emphasis added).

101. The ‘017 patent claims are not directed to a “method of organizing human activity,” “fundamental economic practice long prevalent in our system of commerce,” or “a building block of the modern economy.” Instead, they are limited to a concretely circumscribed set of methods and systems for transcribing electronic information that is transmitted over a computer network via an intermediary.

102. The ‘017 patent claims are not directed at the broad concept/idea of “encrypting” or “decrypting” information. Instead, they are limited to a concretely circumscribed set of methods and systems for transcribing electronic information that is transmitted over a computer network via an intermediary. This type of method and system is unique to the Internet age.

103. The inventive concepts claimed in the ‘017 patent are technological, not “entrepreneurial.” For example, transcribing protected electronic information between a first (e.g., server) encrypted form and a second (e.g., network) encrypted form without a substantial intermediate representation of the information in decrypted form is a specific, concrete solution to the technological problem of transferring encrypted information via an intermediary without providing the intermediary substantial access to the information.

this solution could not always be feasible in the Web environment since large Web-based systems cannot be easily verified to be secure.”)

104. Companies (such as Google competitor) Oracle recognized that until recently security for distributed systems was not a primary concern.

- Security was not a major issue, even for Oracle
 - Standard passwords (scott/tiger, system/manager, ...)
 - Oracle standard users were installed and left open (though not at SAP!)
 - There are some recommendations, but not much more.
 - From Oracle9i, the issue of security was increasingly addressed by Oracle (DBCA: locking of default accounts, ..., 10.2: CONNECT roles)

Andreas Becker, *High Security for SAP Data with Oracle Database Vault and Transparent Data Encryption*, ORACLE PRESENTATION at 6 (2010).

105. Researchers have identified the problems the '017 patent is directed at solving arise from new security challenges relating to cloud computing.

Data Security: Data security was the most important concern that had hindered the acceptance of the cloud computing initially. Storing and processing the data, running software, using CPU and virtual Machines on others' infrastructure were some serious concerns for the users initially. Data breach, data integrity and data loss are major security issues that posed threats to organization's data and software. Moreover, the multi-tenancy model and pooled computing resources over cloud have introduced new security challenges requiring new techniques to tackle with [4] [5] [6].

Deepak Panth, Dhananjay Mehta and Rituparna Shelgaonkar, *A Survey on Security Mechanisms of Leading Cloud Service Providers*, in INTERNATIONAL JOURNAL OF COMPUTER APPLICATIONS 98(1) at 34 (July 2014) (emphasis added).³⁴

106. The '017 patent claims are directed toward a solution rooted in computer technology and use technology unique to computers and computer networking to overcome a problem specifically arising in the realm of secure distributed computing. For example, claims of the '017 patent require cryptographically manipulating protected electronic information using

³⁴ See also Vaibhav Khadilkar, Murat Kantarcioglu, and Bhavani Thuraisingham. *Secure Data Processing in a Hybrid Cloud* at 1-2, Computing Research Repository (CoRR) abs/1105.1982 (2011) (“The emergence of cloud computing has created a paradigm shift by allowing parallel processing of massive amounts of data. . . . [H]ow do users protect themselves from cloud service providers who may be able to access their data? This issue is related to data security and is relevant for users since their data is placed at the provider’s site.”).

one or more intermediary computing devices specially configured to yield a desired result—a result that overrides the routine and conventional sequence of events in electronic communications, even encrypted electronic communications.

107. The '017 patent is directed to specific problems in the field of cryptography. In the “Background” section of the patent, the '017 patent explains that encryption systems use “keys,” similar to passwords, to control how plaintext is encrypted and decrypted. '017 patent, col. 4:39–4:64. An encryption system thereby encrypts and decrypts information differently depending upon the key input. *Id.* Two common cryptanalytic attacks, linear and differential cryptanalysis, analyze large amounts of cipher text (encrypted information) and different possible keys in order to eventually converge on the correct key and break the encryption. *Id.* Both attacks exploit the fact that some encryption systems use static keys to create the cipher text. *Id.* In other words, using the same key repeatedly gives an attacker more information to work with. The inventions of the '017 patent introduce several novel techniques to overcome these weaknesses, particularly where encrypted information is held by an intermediary.

108. The preemptive effect of the '017 patent is concretely circumscribed by specific limitations. For example, claim 1 of the '017 patent requires:

A method for processing information, comprising the steps of:

receiving information to be processed:

defining a cryptographic comprehension function for the information, adapted for making at least a portion of the information incomprehensible;

receiving asymmetric cryptographic key information, comprising at least asymmetric encryption key information and asymmetric decryption key information;

negotiating a new cryptographic comprehension function between two parties to a communication using an intermediary;

processing the information to invert the cryptographic comprehension function and impose the new cryptographic comprehension function in an integral process, in dependence on at least the asymmetric cryptographic key information, without providing the intermediary with sufficient asymmetric cryptographic key information to decrypt the processed information; and

outputting processed information,

wherein the ability of the asymmetric decryption key information to decrypt the processed information changes dynamically.

109. The '017 patent does not attempt to preempt every application of the idea of encrypting electronic information transmitted over a computer network, or even the idea of encrypting electronic information transmitted over a computer network via an intermediary.

110. The '017 patent does not preempt the field of secure third-party communications systems, or prevent use of alternative secure third-party communications systems. For example, the '017 patent includes inventive elements—embodied in specific claim limitations—that concretely circumscribe the patented invention and limit its breadth. These inventive elements are not necessary or obvious tools for achieving secure third-party communications, and they ensure that the claims do not preempt other techniques for secure communications.

111. For example, the '017 patent describes numerous techniques for secure third-party communications that inform the invention's development but do not, standing alone, fall within the scope of its claims:

- Key Escrow. U.S. Pat. No. 6,009,177 to Sudia, relates to a cryptographic system and method with a key escrow feature that uses a method for verifiably splitting users' private encryption keys into components and for sending those components to trusted agents chosen by the particular users.
- Partitioning of Information Storage Systems. U.S. Patent No. 5,956,400 to Chaum, relates to partitioned information storage systems with controlled retrieval.
- Use of a Trusted Intermediary. U.S. Patent No. 6,161,181 to Haynes, describing secure electronic transactions using a trusted Intermediary; U.S. Patent No. 6,145,079 to Misty, describing secure electronic transactions using a trusted intermediary to perform electronic services.
- Split Key Storage. U.S. Patent No. 6,118,874 to Okamoto, teaching encrypted data using split storage key and system.
- Use of a Cryptographic File Labeling System. U.S. Pat. No. 5,953,419 to Lohstroh, disclosing cryptographic file labeling system for supporting secured access by multiple users.

- Computer Security Devices. U.S. Pat. No. 5,982,520 to Weiser, disclosing a personal storage device for receipt, storage, and transfer of digital information to other electronic devices; *see also* U.S. Pat. No. 5,991,519 to Benhammou; U.S. Pat. No. 5,999,629 to Heer; and U.S. Pat. No. 6,034,618 to Tatebayashi.
- Computer Network Firewalls and Agents. U.S. Pat. No. 6,061,798 to Coley, disclosed the use of an assigned proxy agent to verify the authority of an incoming request to access a network element indicated in the request. Once verified, the proxy agent completes the connection to the protected network element on behalf of the source of the incoming request; *see also* U.S. Pat. No. 6,023,762 to Dean, disclosing a data access and retrieval system which comprises a plurality of user data sources each storing electronic data signals describing data specific to a user, or enabling services selected by a user; an agent device which is configurable to select individual ones of the user data sources and present selections of user data and service data to a set of callers who may interrogate the agent device remotely over a communications network; and U.S. Pat. No. 6,029,150 to Kravitz, disclosing a system and method of payment in an electronic payment system wherein a plurality of customers have accounts with an agent. Further, the patent lists thirty-three other patented systems involving Computer Network Firewalls that are not, standing alone, preempted by the inventions claimed in the patents-in-suit.
- Virtual Private Networks. As described in: U.S. Pat. No. 6,079,020 to Liu and U.S. Pat. No. 6,081,900 and twenty other patented systems involving virtual private networks that are not, standing alone, preempted by the inventions claimed in the patents-in-suit.
- Biometric Authentication. U.S. Pat. No. 5,193,855 to Shamos, disclosing the use of biometrics such as fingerprints to facilitate secure communications and identification of users. Further, the '017 lists numerous patented systems that use biometric authentication that are not, standing alone, preempted by the inventions claimed in the patents-in-suit.

112. Although “[e]ncryption, in general, represents a basic building block of human ingenuity that has been used for hundreds, if not thousands, of years,”³⁵ the claims in the ‘017 patent do not claim, or attempt to preempt, “some process that involves the encryption of data for some purpose” (or similar abstraction).

³⁵ *Paone v. Broadcom Corp.*, Case No. 15 CIV. 0596 BMC GRB, 2015 WL 4988279, at *7 (E.D.N.Y. Aug. 19, 2015) (*citing Fid. Nat'l Info. Servs., Inc.*, Petitioner, CBM2014-00021, 2015 WL 1967328, at *8 (Apr. 29, 2015) (both upholding the patent eligibility of patents directed toward encryption).

113. The '017 patent does not claim, or attempt to preempt, the performance of an abstract business practice on the Internet or using a conventional computer

114. The claimed subject matter of the '017 patent is not a pre-existing but undiscovered algorithm.

115. The '017 patent claims systems and methods that “could not conceivably be performed in the human mind or pencil and paper.”³⁶

116. The claims in the '017 patent require the modifying of data that has concrete and valuable effects in the field of secure third-party communications. By allowing an intermediary to receive secure information but not gain access to the unencrypted form of the information, the '017 patent improves the security of computer systems. Prior art systems that the '017 patent remedies enabled unauthorized “access to private communications or otherwise undermine[d] transactional security or privacy.” Companies have described the use of encryption in the cloud as important to improve the security and functioning of systems.

For many organizations, keeping data private and secure has also become a compliance requirement. Standards including Health Insurance Portability and Accountability Act of 1996 (HIPAA), Sarbanes-Oxley (SOX), Payment Card Industry Data Security Standard (PCI DSS), the Gramm-Leach-Bliley Act, and EU Data Protection Directives all ***require that organizations protect their data at rest and provide defenses against threats.***

HP Atalla Cloud Encryption: Securing Data in The Cloud, HP TECHNICAL WHITE PAPER 2 (2014) (emphasis added).

117. The '017 patent claims systems and methods not merely for transferring secure information over a computer network, but for making the computer network itself more secure.

118. The claimed invention in the '017 claims is rooted in computer technology and overcame problems specifically arising in the realm of computer networks.

³⁶ *TQP Dev., LLC v. Intuit Inc.*, Case No. 2:12-CV-180-WCB, 2014 WL 651935, at *4 (E.D. Tex. Feb. 19, 2014) ((finding claims directed to encryption to be patent eligible); *see also Paone v. Broadcom Corp.*, Case No. 15 CIV. 0596 BMC GRB, 2015 WL 4988279, at *7 (E.D.N.Y. Aug. 19, 2015); *see also Prism Technologies, LLC v. T-Mobile USA, Inc.*, 12-cv-124, Dkt. No. 428 at 7 (D. Neb. Sept. 22, 2015) (Finding on cross motions for summary judgment that patents directed at delivering resources over an untrusted network were patent eligible. “The problems addressed by Prism’s claims are ones that ‘arose uniquely in the context of the Internet.’”).

119. The systems and methods claimed in the '017 patent were not a longstanding or fundamental economic practice at the time of the patented inventions. Nor were they fundamental principles in ubiquitous use on the Internet or computers in general. As just one example, at the time the inventions disclosed in the '017 patent were conceived, the use of asymmetric encryption keys was described by Oracle as “relatively new.”³⁷

A Public Key Infrastructure (PKI) consists of protocols, services, and standards supporting applications of public key cryptography. ***Because the technology is still relatively new***, the term PKI is somewhat loosely defined.

Introduction to the SSL Technology, ORACLE DOCUMENTATION (February 1, 2001), http://docs.oracle.com/cd/E53645_01/tuxedo/docs12cr2/security/publickey.html (emphasis added).

120. The asserted claims do not involve a method of doing business implemented on a computer; instead, it involves a method for changing data in a way that will affect the communication system itself, by making it more secure. The security challenges that the '017 patent is directed at were new and unique to distributed networks as confirmed in a recent paper from Accenture Services Pvt. Ltd. “The unprecedented growth of cloud computing has created new security challenges. The problem is ever more complex as there is a transition from traditional computing to a service-based computing.”³⁸

121. The '017 patent claims are not directed to a mathematical relationship or formula. The '017 patent claims concrete, specific computer systems and methods for cryptographically protecting and managing access to secure data in multi-party communications.

122. The '017 patent claims cover a systems and methods that transform data from one form into another that will be recognizable by the intended recipient but secure against

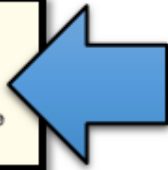
³⁷ See also BACKUPEDGE ENCRYPTION WHITEPAPER, MICROLITE CORPORATION at 2 (2003) (describing the technology of asymmetric keys as “new”); Roger Clarke, MESSAGE TRANSMISSION SECURITY (May 1998), <http://www.rogerclarke.com/II/CryptoSecy.html> (“Public key cryptography is relatively new and technically complex.”).

³⁸ Deepak Panth, Dhananjay Mehta and Rituparna Shelgaonkar, *A Survey on Security Mechanisms of Leading Cloud Service Providers*, in INTERNATIONAL JOURNAL OF COMPUTER APPLICATIONS 98(1) at 34 (July 2014).

decryption by unintended recipients. IBM, in its reference guides (“redbooks”), refers to encryption as “transform[ing] data that is unprotected.”


Encryption concepts and terminology

Encryption transforms data that is unprotected, or *plain text*, into encrypted data, or *ciphertext*, by using a *key*. Without knowledge of the encryption key, the ciphertext cannot be converted back to plain text.



Bertrand Dufrasne and Robert Tondini, IBM DS8870 DISK ENCRYPTION 6th Edition at 4 (2015) (from a reference guide published by IBM).

123. One or more claims of the ‘017 patent require a specific configuration of electronic devices, a network configuration, and the use of encryption systems to secure communications from access by an intermediary. These are meaningful limitations that tie the claimed methods and systems to specific machines. For example, the below diagram from the ‘017 patent illustrates a specific configuration of hardware disclosed in the patent.



‘017 patent, Fig. 2.

C. Information Record Infrastructure Patents

124. The IRI patents disclose specific computer based systems and methods for electronically structuring and controlling access to protected data in a plurality of external databases.

125. Over fifteen years ago, Mr. Felsher conceived of the inventions disclosed in the IRI patents, based on his experiences with the limitations in existing systems for controlling access to electronic medical records and protected electronic data.

126. During Mr. Felsher's work in the field of electronic medical records, he witnessed first-hand the drawbacks to existing computer systems and methods for controlling access to protected data. Existing systems failed to efficiently transmit unstructured protected information. '368 patent, col. 3:5-10. Other problems included the inability to secure the protection of data, integrate content management functions, and create a trust infrastructure wherein an independent third party represents and serves as an agent for the content owner. *Id.* at col. 3:4-54:16. The result was an inability to effectively manage access to protective data. The IRI patents disclosed systems and methods that overcome these drawbacks. The inventions disclosed in the IRI patents improved upon the then-available technology, enabled efficient access control of unstructured data, reduced costs, and ultimately resulted in a more secure system.

127. Google values systems that provide secure systems and methods for controlling access to protected data such as the system disclosed in the IRI patents and makes the implementation of such systems a priority.

As a cloud pioneer, Google fully understands the security implications of the cloud model. Our cloud services are designed to deliver better security than many traditional on-premises solutions. We make security a priority to protect our own operations, but because Google runs on the same infrastructure that we make available to our customers, your organization can directly benefit from these protections. ***That's why we focus on security, and protection of data is among our primary design criteria. Security drives our organizational structure, training priorities and hiring processes.*** It shapes our data centers and the technology they house. It's central to our everyday operations and disaster planning, including how we address threats. It's prioritized in the way we handle customer data. And it's the cornerstone of our account controls, our compliance audits and the certifications we offer our customers.

Google Cloud Platform Security Whitepaper, GOOGLE SECURITY WHITEPAPER, <https://cloud.google.com/security/whitepaper?hl=en> (emphasis added).

128. Google's competitors, such as Hewlett-Packard Company and Microsoft Corporation, have confirmed the importance and value of systems and methods that manage access to protected data.

Today, the need for data protection and security goes well beyond the realm of access privileges and firewalls. Organizations of all sizes, in public and private sectors, must not only protect information from unauthorized access and intrusion but also manage how documents, presentations, spreadsheets, and e-mails are handled in the normal course of daily business

HP Information Rights Management Solutions Ensuring Life Cycle Protection Of Digital Information in Microsoft Environments, HP WHITE PAPER (2005).³⁹

Such cloud adoption within the healthcare industry is gaining momentum because the economic, clinician productivity and care team collaboration advantages of the cloud are undeniable. However, as was the case for UCHealth, there's ***one fundamental concern that continues to weigh heavily on the minds of providers: Is patient data safe, secure and private in the cloud.***

University of Colorado Health Adopts Microsoft Office 365 for its data privacy and security commitment, MICROSOFT ON THE ISSUES BLOG (December 18, 2013), <http://blogs.microsoft.com/on-the-issues/2013/12/18/university-of-colorado-health-adopts-microsoft-office-365-for-its-data-privacy-and-security-commitment/> (emphasis added).

129. Academics have confirmed the value of secure information access management systems such as the inventions disclosed in the IRI patents.

With the proliferation of the Internet, the speed and ease of digital data exchange has increased, together with the number of potential parties that can exchange data. This has also meant that digital data security is no longer confined to the computer that holds the original data, or even behind corporate firewalls. Furthermore, data security no longer applies only to the access to data, but also to what the user can do with the data

Alapan Arnab and Andrew Hutchinson, *Digital Rights Management - An Overview of Current Challenges and Solutions*, in PROCEEDINGS OF INFORMATION SECURITY SOUTH AFRICA CONFERENCE (2004) (emphasis added).⁴⁰

³⁹ See also Albert Biketi, *HP Gets Serious About End-To-End Data Protection*, HP SECURITY BLOG (February 19, 2015) (Mr. Biketi, vice president and general manager of data security and encryption at Hewlett-Packard stated "***What our customers need is a data-centric solution that protects sensitive information*** from the moment it's created throughout its entire lifecycle. That means protecting data wherever it moves – from emails to databases and attachments . . . in the cloud, in use, at rest, and in motion.") (emphasis added).

⁴⁰ See also Murat Kantarcioglu, Wei Jiang, and Bradley Malin, *A Privacy-Preserving Framework for Integrating Person-Specific Databases* at 299, PRIVACY IN STATISTICAL DATABASES LNCS 5262 (2008) (Describing the difficulty in managing medical records stored in multiple electronic databases "in the healthcare realm, patients are mobile and their data can be collected by multiple locations, such as when a patient visits one hospital for primary care and a second hospital to participate in a clinical trial.").

130. Although major corporations offer systems for providing secure access to protected data today, at the time the inventions disclosed in the IRI patents were conceived, systems had significant limitations that were addressed by the inventions disclosed in the IRI patents.

While “awareness of risks and of possible technical solutions is increasing,” the authors would appear to be describing a rather precarious environment, at least in the short run. The picture does not improve when one focuses on the details of some of the technical fixes. Barrows and Clayton deem “tight” prospective access restrictions—a “need to know” mandatory access control model—as largely incompatible with the dynamic health care environment.

Reid Cushman, *Serious Technology Assessment for Health Care Information Technology*, JOURNAL OF THE AMERICAN MEDICAL INFORMATICS ASSOCIATION 4(4) (1997).⁴¹

131. The claims in the IRI patents describe solutions that are rooted in computer technology to overcome problems specific to and characteristic of complex computer networks where protected data is stored. For example, academics identified distributed information systems as leading to new problems regarding information rights management that the IRI patents solve.

The development and wider use of wireless networks and mobile devices has led to novel pervasive computing environments *which pose new problems for software rights management* and enforcement on resource-constrained and occasionally connected devices. . . . The latter opens new channels for super-distribution and sharing of software applications that do not impose a cost on the user.

Ivana Dusparic, Dominik Dahlem, and Jim Dowling, *Flexible Application Rights Management in a Pervasive Environment*, in IEEE INTERNATIONAL CONFERENCE ON E-TECHNOLOGY, E-COMMERCE AND E-SERVICE, pages 680–685 (2005) (emphasis added).⁴²

⁴¹ This reference is cited on the face of the IRI patents as an exemplar illustrating limitations in systems existing at the time the inventions disclosed in the IRI patents were conceived; *see also* Alapan Arnab and Andrew Hutchinson, *Digital Rights Management - An Overview of Current Challenges and Solutions*, in PROCEEDINGS OF INFORMATION SECURITY SOUTH AFRICA CONFERENCE (2004) (emphasis added) (“none of these products provide for all the needs of an enterprise, and furthermore these products do not offer all the benefits that DRM potentially offers to an enterprise).

⁴² *See also* Aaron Franks, Stephen LaRoy, Miek Wood, and Mike Worth. *Idrm: An Analysis Of Digital Rights Management For The Itunes Music Store*, TECHNICAL REPORT, UNIVERSITY OF BRITISH COLUMBIA (2005) (“The need for secure digital rights management (DRM) is more urgent today than ever before. With the rapid increase in broadband availability, Internet file sharing has become a threat to content providers’ bottom line.”); Mike Godwin, *What Every Citizen Should Know About DRM, A.K.A. ‘Digital Rights Management,’* PUBLIC KNOWLEDGE (2004) (“As circumvention tools evolve, and as new technologies pose new infringement

Then there is the cloud. Cloud, cloud, cloud, it's on every webcast, in every article. The cloud has many advantages. Why wouldn't you want to outsource all your costs of network management, storage, system administration? The cloud makes perfect sense but has one massive concern... security.

Simon Thorpe, *Security in the Enterprise 2.0 World: Conflicts of Collaboration*, ORACLE OFFICIAL BLOG, September 27, 2010, <https://blogs.oracle.com/irm/>.

132. Although secure and effective information rights management, in some form, has been an objective of corporations and researchers for many years ('368 patent, col. 6:61-7:3), the IRI patents are directed at solving problems that are unique to the realm of computers and specifically network cloud computing.

133. The systems and methods disclosed in the IRI patents have particular application to two primary fields: electronic medical records and electronic rights management. Shortcomings in available technology at the time the inventions disclosed in the IRI patents were conceived, led to the development of the IRI patents.

134. A brief overview of the state of the prior art in these two areas provides context to understanding the truly inventive nature of the IRI patents. The specific systems and methods disclosed and claimed in the IRI patents are discussed in detail later in this Complaint.

135. Background on the state of the art at the time of the inventions disclosed in the IRI patents confirms that the patented inventions are limited to specific computer systems and methods and address issues specific to accessing protected data using modern computer networks.

136. ***Information Rights Management.*** The inventions disclosed in the IRI patents have particular application to the management of rights in digital works, to allow a content owner to exploit the value of the works while assuring control over the use and dissemination.

problems, the locking of industrial sectors into a particular "standard" scheme, mediated and supervised by government, actually slows the ability of the content sector to respond to new problems.); HP DIGITAL RIGHTS MANAGEMENT (DRM) FOR NETWORK AND SERVICE PROVIDERS (NSPs), HP SOLUTION BRIEF (2003) ("DRM [Digital Rights Management] is an emerging technology with fragmented addressable markets, solution capabilities and standards."); Arun Kulkarni, Harikrisha Gunturu, and Srikanth Datla, *Association-Based Image Retrieval* at 183, WSEAS TRANS. SIG. PROC. Vol.4(4) (April 2008) ("With advances in computer technology and the World Wide Web there has been an explosion in the amount and complexity of multimedia data that are generated, stored, transmitted, analyzed, and accessed.").

The IRI patents address problems specific to and arising from distribution and protected works on the internet.

137. At the time the inventions disclosed in the IRI patents were conceived, the growth of the internet created unique problems relating to managing rights to protected works.

There's too much data being collected in so many ways, and a lot of it in ways that you don't feel you had a role in the specific transaction," he [Craig Mundie] said. "Now that you're just being observed, whether it's for commercial purposes or other activities, *we have to move to a new model.*" . . . Under the model imagined by Mundie [a] central authority would distribute encryption keys to applications, allowing them to access protected data in the ways approved by the data's owners.

Tom Simonite, *Microsoft Thinks DRM Can Solve the Privacy Problem*, MIT TECHNOLOGY REVIEW, October 10, 2013 (emphasis added) (Craig Mundie is Senior Advisor to the CEO at Microsoft and its former Chief Research and Strategy Officer).⁴³

138. In the late 1990s and early 2000s, information rights management systems had significant limitations. Prior art systems did not create a trust infrastructure, wherein an independent third party represents and serves as agent for the content owner, implementing a set of restrictive rules for use of the content, and interacting and servicing customers.

139. Rudimentary rights management systems such as Microsoft's PlayForSure and RealNetwork's Rhapsody were still years from being released. Even when these systems were released in 2004 they had significant limitations. Both systems lacked the ability of a third party to act as an intermediary between a content creator and a user. The state of the art at the time the inventions disclosed in the IRI patents were conceived underscores the inventive nature of the IRI patents.

140. *Electronic Medical Records*. The IRI patents disclose systems and methods for controlling access to protected health information where the information is stored in one or more external databases. Systems for controlling access to medical records, contemporaneous to the

⁴³ See also Martin Abrahams, *Document Theft - IRM as a Last Line of Defense*, ORACLE IRM, THE OFFICIAL BLOG, August 1, 2011, <https://blogs.oracle.com/irm/> ("The relevance of IRM is clear. . . . In a cloudy world, where perimeters are of diminishing relevance, you need to apply controls to the assets themselves.").

IRI patents had significant limitations that the IRI patents address.⁴⁴ These systems included: (1) Anonymizing Records. A method used in contemporaneous systems to the IRI patents is the maintenance of anonymous medical records. However, anonymizing techniques did not provide patients and medical professionals the ability to access patient specific records. (2) Indexing. Systems contemporaneous to the IRI patents indexed medical records with anonymous identification codes.⁴⁵ While these systems preserved privacy, these systems made locating a database record other than by patient identifier, or its accession identifier, difficult. (3) Proxy Systems. Other contemporaneous systems used a proxy server to protect user privacy. However, systems using an Internet proxy resulted in a loss of rights and did not act in a representative capacity for the content owner, and did not integrate content management functions.

141. In addition, access to these early medical records systems was limited to authorized individuals who were on-site, as these systems provided little-to-no connectivity to anyone outside of the organization or to the Internet generally. Because access was restricted to on-site users on a local network using stationary terminals in designated areas, there was very little emphasis placed on data security.

⁴⁴ See Reid Cushman, *Serious Technology Assessment for Health Care Information Technology*, J. AM. MED. INFORM. ASSOC. 4: 259-265 (1997) (This article is cited on the face of the IRI patents and finds “Data protection practices in the typical late twentieth-century organization are not very good, even in putatively “secure” institutions. . . The forthcoming study of health care security by the National Academy of Sciences, to be released in February 1997, is expected to reach a similar conclusion. The widespread deficits in security are hardly a secret; they are common fodder among information systems professionals.”); see also Bhavani Thuraisingham, *Data and Applications Security: Developments and Directions* at 2, PROCEEDINGS IEEE COMPSAC (2002) (Discussing issues with electronic medical records “There are numerous security issues for such systems including secure information sharing and collaboration. Furthermore, data is no longer only in structured databases. . . . Security for such data has not received much attention.”).

⁴⁵ See also Murat Kantarcioglu and Chris Clifton, *Security Issues in Querying Encrypted Data* at 2, TECHNICAL REPORT CSD TR 04-013, Purdue University Computer Sciences Department (2004) (“methods that quantize or “bin” values reveal data distributions. Methods that hide distribution, but preserve order, can also disclose information if used naively”).

142. In sharp contrast to the flexible, modular, and tightly integrated multi-layer security and access control framework disclosed and claimed in the IRI patents, systems such as Epic System Corporation's CareWeb⁴⁶ had significant limitations, including: inability to effectively control access on a record-by-record basis within respective external databases, as claimed in several IRI patents; inability to distinguish between records within an external or backend database, the databases accessed through CareWeb were basically opaque to the "CareWeb" system; and CareWeb's fixed structure was expressly limited to a particular, monolithic front-end architecture for secure implementation.

143. At the time the inventions disclosed in the IRI patents were conceived, the medical community showed little sign of implementing a system for controlling access to medical records that were stored in external databases. Further, computer networks presented new challenges and unique problems that the IRI patents addressed.

As health care moves from paper to electronic data collection, providing easier access and dissemination of health information, the development of guiding privacy, confidentiality, and security principles is necessary to help balance the protection of patients' privacy interests against appropriate information access. . . . It is imperative that all participants in our health care system work actively toward a viable resolution of this information privacy debate.

Suzy Buckovich, Helga Rippen, and Michael Rozen, *Driving Toward Guiding Principles: A Goal for Privacy, Confidentiality, and Security of Health Information*, J. AM. MED. INFORM. ASSOC. 6 (1999).

144. The need for a secure system for providing access to medical records was specifically required in the cloud computing context where medical records were stored in one or more external databases.

The healthcare industry is in a major period of transformation and IT modernization. More than ever, healthcare providers and professionals are faced with the need to be more efficient, reduce costs and collaborate seamlessly as virtual teams to deliver higher quality care for more people at a lower cost point. Healthcare organizations are increasingly looking to cloud technologies to help

⁴⁶ John D. Halamka, Peter Szolovits, David Rind, and Charles Safran, *A WWW Implementation of National Recommendations for Protecting Electronic Health Information*, J. AM. MED. INFORM. ASSOC. 4: 458-464 (1997) (The limitations of the CareWeb system are discussed in depth in the specification of the IRI patents.).

them meet these goals. However, a natural concern with using cloud technology is keeping sensitive health information private and secure.

Hemant Pathak, *Data Privacy and Compliance in the Cloud Is Essential for the Healthcare Industry*, MICROSOFT HEALTH TECHNOLOGY BLOG (December 2013), <http://www.microsoft.com/en-us/health/blogs/data-privacy-and-compliance-in-the-cloud-is-essential-for-the-healthcare-industry/default.aspx>.

145. On information and belief, contemporaneous to, and following conception of the inventions disclosed in the IRI patents, Texas educational institutions, Texas governmental entities, and businesses headquartered in Texas actively entered the field of electronically structuring and controlling access to protected health data stored in a plurality of external databases. In 2006, Texas Gov. Rick Perry called for widespread adoption of health information technology (“HIT”).⁴⁷ Governor Perry signed Senate Bill 45, which created the Health Information Technology Advisory Committee (HITAC) within the Texas Statewide Health Coordinating Council in the Department of State Health Services.⁴⁸ In addition, various universities studied and implemented systems for securely managing access to distributed medical records.⁴⁹

146. Texas based companies incorporated systems and methods for electronically structuring and controlling access to protected data in a plurality of external databases into numerous products. Many of these same companies cite the IRI patents in their own patents. Texas based businesses that developed products and/or technologies incorporating these systems included: HP Enterprise Services, LLC of Plano, Texas; Hospitalists Now, Inc. of Austin, Texas; StandardCall, LLC of Frisco, Texas; Security First Corp whose inventors were based in various locations in Texas; Huawei Technologies Co., Ltd. of Plano, Texas; Omnyx LLC whose

⁴⁷ Gov. Rick Perry, *State-of-the-State Speech* (February 6, 2007), available at: <http://governor.state.tx.us/news/speech/5567/>.

⁴⁸ Texas Senate Bill 45, Texas 79th Regular Legislative Session (25 TAC §§571.11-571.13); *see also* Texas Executive Order RP-61, *Relating to the Creation, Composition, and Operation of the Governor's Health System Integrity Partnership for the State of Texas* (October 9, 2006) (The Partnership was directed to develop a method for secure exchange of electronic health information.).

⁴⁹ *See* David E. Gerber et al., *Predictors and Intensity of Online Access to Electronic Medical Records Among Patients with Cancer*, J ONCOL PRACT. Vol. 10(5) (Sept. 2014) (studying electronic medical record infrastructure implementations at and Texas hospitals).

inventors included individuals based in Texas; Electronic Data Systems Corporation of Plano, Texas and South Texas Accelerated Research Therapeutics, LLC of San Antonio, Texas.

1. U.S. Patent No. 7,805,377

147. U.S. Patent No. 7,805,377 (the “’377 patent”) entitled, Information Record Infrastructure, System and Method, was filed on August 19, 2008, and claims priority to July 6, 2000. St. Luke is the owner by assignment of the ‘377 patent. A true and correct copy of the ‘377 patent is attached hereto as Exhibit C. The ‘377 patent claims specific methods and systems for securely controlling access to a plurality of digital records by a remote computer.

148. The ‘377 patent has been cited by over 30 United States patents and patent applications as relevant prior art. Specifically, patents issued to the following companies have cited the ‘377 patent as relevant prior art.

- Symantec Corporation
- Siemens Medical Solutions USA, Inc.
- AT&T Corporation
- Hospitalists Now, Inc.
- MasterCard International Incorporated
- J.D. Power and Associates
- Middlegate, Inc.
- Cardiac Pacemakers, Inc.
- Robert Bosch GmbH

149. The ‘377 patent claims a technical solution to a problem unique to computer networks – securely transmitting encrypted digital records and controlling access to digital records requested by a remote computer.

150. At the time of the inventions claimed in the ‘377 patent, electronically structuring and controlling access to protected data in a plurality of external databases presented new and unique issues over the state of the art. As explained in the ‘377 patent: “The present invention therefore seeks to provide a comprehensive set of technologies to address the full scope of issues presented in implementing a secure and versatile information content infrastructure that respects the rights of content owners and users to privileges, such as confidentiality.” ‘377 patent, col. 54:27-33.

151. Although the systems and methods taught in the '377 patent have been adopted by leading businesses today, at the time of invention, the technologies taught in the '377 patent claims were innovative and novel. "Existing systems do not create a trust infrastructure, wherein an independent third party represents and serves as an agent for the content owner, implementing a set of restrictive rules for use of content . . . Thus, existing intermediaries do not act in a representative capacity for the content owner, and do not integrate content management functions." '377 patent, col. 5:8-20.

152. Further, the '377 patent claims improve upon the functioning of a computer system by allowing encrypted electronic data to be securely transmitted through an intermediary. This improves the security of the computer system and allows it to be more efficient. "[B]y consolidating a plurality of institutions [referring to digital records stored in external databases], uniformity, interoperability, cost reductions, and improved security result." '377 patent, col. 69:28-30.

153. The '377 patent claims are not directed to a "method of organizing human activity," "fundamental economic practice long prevalent in our system of commerce," or "a building block of the modern economy." Instead, they are limited to a concretely circumscribed set of methods and systems that provide a conduit for the authorized transmission of digital records, while maintaining the security of the records against unauthorized access.

154. The '377 patent claims are not directed at the broad concept/idea of "managing digital records." Instead, the '377 patent claims are limited to a concretely circumscribed set of methods and systems for authorizing and transmitting secure digital records. These methods and systems are technologies unique to the Internet age.⁵⁰

⁵⁰ See *Trusted Cloud: Microsoft Azure Security, Privacy, and Compliance* at 5, MICROSOFT WHITE PAPER (April 2015) ("Cloud services raise unique privacy challenges for businesses. As companies look to the cloud to save on infrastructure costs and improve their flexibility, they also worry about losing control of where their data is stored, who is accessing it, and how it gets used.").

155. The '377 patent claims are directed toward a solution rooted in computer technology and use technology unique to computers and computer networking to overcome a problem specifically arising in the realm of secure distributed computing. For example, one or more claims of the '377 patent require a database adapted to store information for determining patient-controlled access control criteria, authenticate the requestor and determine sufficiency of the patient-provided access control authorization, and generating an electronic payment authorization.

156. The '377 patent is directed to specific problems in the field of digital record access and transmission.

157. The preemptive effect of the claims of the '377 patent are concretely circumscribed by specific limitations. For example, claim 7 of the '377 patent requires:

A system adapted to control access to a patient medical record hosted by at least one medical record repository comprising a plurality of record portions, each record portion being associated with different patient-controlled access control criteria, said system comprising an automated processor, a database adapted to store information for authenticating requestors, a database adapted to store information for determining patient-controlled access control criteria for respective record portions of a patient medical record, and a computer network interface, said processor being controlled by instructions stored on a computer readable storage medium to:

- (a) receive a request for a medical record from a requestor, said request comprising a medical record identifier, a requestor identifier, requestor authentication information, and patient-provided access control authorization;
- (b) process the request for the medical record, to authenticate the requestor and determine sufficiency of the patient-provided access control authorization to meet the patient-controlled access control criteria for each respective record portion encompassed by the request;
- (c) selectively communicate through the computer network interface to the at least one medical record repository, an identification of each record portion for which access control

criteria are determined to be sufficient for access by the requestor;
and

(d) generating an electronic payment authorization associated with the request, for compensation of at least one of said system and the at least one medical record repository.

158. The '377 patent does not attempt to preempt every application of the idea of controlling access to an electronic medical record over a computer network.

159. The '377 patent does not preempt the field of electronically structuring and controlling access to protected medical records in a plurality of external databases. For example, the '377 patent includes inventive elements—embodied in specific claim limitations—that concretely circumscribe the patented invention and greatly limit its breadth. These inventive elements are not necessary or obvious tools for achieving secure third-party communications, and they ensure that the claims do not preempt other techniques for secure communications.

160. For example, the '377 patent describes numerous techniques for electronically structuring and controlling access to protected data in a plurality of external databases. The techniques inform the invention's development but do not, standing alone, fall within the scope of its claims:

- Rights-Based Access to Database Records. U.S. Pat. No. 5,325,294 to Keene, relates to a system that receives and stores the individual's medical information, after the individual is tested to establish this information and the date on which such information was most recently obtained.
- Security Tokens. U.S. Patent No. 5,978,918 to Scholnick, discloses a back-end process returns a time sensitive token that the “sender” sends to the “receiver.” The “receiver” takes the time sensitive token and uses it to retrieve the private data.⁵¹
- Role-Based Access. U.S. Pat. No. 6,023,765 to Kuhn, relates to a role-based access control in multi-level secure systems.

⁵¹ See also Arindam Khaled et al., *A Token-based Access Control System for RDF Data in the Clouds* at 104, in PROCEEDINGS OF THE 2ND IEEE INTERNATIONAL CONFERENCE ON CLOUD COMPUTING TECHNOLOGY AND SCIENCE (2010) (discussing the use of a “token-based access control system . . . implemented in Hadoop (an open source cloud computing framework)”).

- Secure Networks. U.S. Pat. No. 5,579,393 to Conner, relates to a system and method for secure digital records, comprising a provider system and a payer system.
- Cryptographic Technology. U.S. Pat. No. 5,956,408 to Arnold, relates to an apparatus and method for secure distribution of data. Data, including program and software updates, are encrypted by a public key encryption system using a private key.
- Watermarking. U.S. Pat. No. 5,699,427 to Chow, relates to a method to deter document and intellectual property piracy through individualization, and a system for identifying the authorized receiver of any particular copy of a document.
- Computer System Security. U.S. Pat. No. 5,881,225 to Worth, relates to a security monitor for controlling functional access to a computer system. A security monitor controls security functions for a computer system. A user desiring access to the system inputs a user identification and password combination, and a role the user to assume is selected from among one or more roles defined in the system.
- Computer Security Devices. U.S. Pat. No. 5,982,520 to Weiser, relates to a personal storage device for receipt, storage, and transfer of digital information to other electronic devices has a pocket sized crush resistant casing with a volume of less than about ten cubic centimeters.
- Computer Network Firewall. U.S. Pat. No. 5,944,823 to Jade, relates to a system and method for providing outside access to computer resources through a firewall. A firewall isolates computer and network resources inside the firewall from networks, computers and computer applications outside the firewall.
- Virtual Private Network. U.S. Pat. No. 6,079,020 to Liu, relates to a method and an apparatus for managing a virtual private network operating over a public data network. This public data network has been augmented to include a plurality of virtual private network gateways so that communications across the virtual private network are channeled through the virtual private network gateways.
- Biometric Authentication. U.S. Pat. No. 5,193,855 to Shamos relates to a patient and healthcare provider identification system which includes a database of patient and healthcare provider information including the identity of each patient and provider and some identification criteria (such as fingerprint data).⁵²

⁵² Nary Subramanian, *Biometric Authentication*, in *ENCYCLOPEDIA OF CRYPTOGRAPHY AND SECURITY* (S. Jajodia and H.C.A. van Tilborg 2nd ed. 2011) (“Biometric authentication is a technique for identifying the person accessing a secured asset . . . by comparing their unique biological features . . . [an] issue with biometric authentication is privacy of personal data.”).

161. The '377 patent does not claim, or attempt to preempt, the performance of an abstract business practice on the Internet or using a conventional computer.

162. The '377 patent claims systems and methods that “could not conceivably be performed in the human mind or pencil and paper.”⁵³

163. The '377 patent claims systems and methods not merely for transferring secure information over a computer network, but for making the computer network itself more secure.


164. The claimed inventions in the '377 claims is rooted in computer technology and overcomes problems specifically arising in the realm of computer networks.

165. The systems and methods claimed in the '377 patent were not a longstanding or fundamental economic practice at the time of patented inventions. Nor were they fundamental principles in ubiquitous use on the Internet or computers in general.

166. The asserted claims do not involve a method of doing business that happens to be implemented on a computer; instead, they involve a method for changing digital records in a way that will affect the communication system itself, by making it more secure.

167. One or more claims of the '377 patent require a specific configuration of electronic devices, a network configuration, external databases, a computer network interface, etc.. These are meaningful limitations that tie the claimed methods and systems to specific machines. For example, the below diagram from the '377 patent illustrates a specific configuration of hardware disclosed in the patent.

⁵³ *TQP Dev., LLC v. Intuit Inc.*, Case No. 2:12-CV-180-WCB, 2014 WL 651935, at *4 (E.D. Tex. Feb. 19, 2014) (finding claims directed to encryption to be patent eligible); *see also Paone v. Broadcom Corp.*, Case No. 15 CIV. 0596 BMC GRB, 2015 WL 4988279, at *7 (E.D.N.Y. Aug. 19, 2015); *see also Prism Technologies, LLC v. T-Mobile USA, Inc.*, 12-cv-124, Dkt. No. 428 at 7 (D. Neb. Sept. 22, 2015) (Finding on cross motions for summary judgment that patents directed at delivering resources over an untrusted network were patent eligible. “The problems addressed by Prism’s claims are ones that ‘arose uniquely in the context of the Internet.’”).



'377 patent, Fig. 1.

2. U.S. Patent No. 7,587,368

168. U.S. Patent No. 7,587,368 ("the '368 patent") entitled, Information Record Infrastructure, System and Method, was filed on July 5, 2001, and claims priority to July 6, 2000. St. Luke is the owner by assignment of the '368 patent. A true and correct copy of the '368 patent is attached hereto as Exhibit D. The '368 patent claims specific methods and systems for securely controlling access to a plurality of digital records by a remote computer.

169. The '368 patent has been cited by over 100 United States patents and patent applications as relevant prior art. Specifically, patents issued to the following companies have cited the '368 patent as relevant prior art.

- Microsoft Corporation
- LG Electronics, Inc.
- Canon Kabushiki Kaisha
- Hewlett-Packard Development Company, L.P.
- Voltage Security, Inc.
- Northrop Grumman Systems Corporation
- International Business Machines Corporation
- McAfee, Inc.
- J.D. Power And Associates
- NEC Corporation
- Electronics And Telecommunications Research Institute (ETRI)
- Koninklijke Philips Electronics N.V.

- Huawei Technologies Co., Ltd.
- Ricoh Co., Ltd.
- Massachusetts Institute Of Technology

170. The '368 patent claims a technical solution to a problem unique to computer networks – securely transmitting encrypted digital records and controlling access to digital records requested by a remote computer.

171. At the time of the inventions claimed in the '368 patent, electronically structuring and controlling access to protected data in a plurality of external databases presented new and unique issues over the state of the art. As explained in the '368 patent: “The present invention therefore seeks to provide a comprehensive set of technologies to address the full scope of issues presented in implementing a secure and versatile information content infrastructure that respects the rights of content owners and users to privileges, such as confidentiality.” '368 patent, col. 54:27-33.

172. Although the systems and methods taught in the '368 patent have been adopted by leading businesses today, at the time of invention, the technologies taught in the '368 patent claims were innovative and novel. “Existing systems do not create a trust infrastructure, wherein an independent third party represents and serves as an agent for the content owner, implementing a set of restrictive rules for use of content . . . Thus, existing intermediaries do not act in a representative capacity for the content owner, and do not integrate content management functions.” '368 patent, col. 5:4-16.

173. Further, the '368 patent claims improve upon the functioning of a computer system by allowing encrypted electronic data to be securely transmitted through an intermediary. This improves the security of the computer system and allows it to be more efficient. “[B]y consolidating a plurality of institutions [referring to digital records stored in external databases], uniformity, interoperability, cost reductions, and improved security result.” '368 patent, col. 67:65-67.

174. The '368 patent claims are not directed to a “method of organizing human activity,” “fundamental economic practice long prevalent in our system of commerce,” or “a

building block of the modern economy.” Instead, they are limited to a concretely circumscribed set of methods and systems that provide a conduit for the authorized transmission of digital records, while maintaining the security of the records against unauthorized access.

175. The '368 patent claims are not directed at the broad concept/idea of “managing digital records.” Instead, the '368 patent claims are limited to a concretely circumscribed set of methods and systems for authorizing and transmitting secure digital records. These methods and systems are technologies unique to the Internet age.

176. The '368 patent claims are directed toward a solution rooted in computer technology and use technology unique to computers and computer networking to overcome a problem specifically arising in the realm of secure distributed computing. For example, one or more claims of the '368 patent require encrypting and sending, by the server system, the requested digital record which has been validated, using the public key and the session key to encrypt the digital record - a procedure that overrides the routine and conventional sequence of events in electronic communications, even encrypted electronic communications.

177. The '368 patent is directed to specific problems in the field of digital record access and transmission.

178. The preemptive effect of the claims of the '368 patent are concretely circumscribed by specific limitations. For example, claim 1 of the '368 patent requires:

A method, comprising the steps of:

storing a plurality of digital records and respective access rules for each digital record in a computer memory associated with a server system;

receiving a request for access, from a remote computer, to access a digital record stored in the computer memory;

validating, by the server system, the received request to access the digital record by applying a respective set of access rules for the digital record stored in the computer memory;

retrieving, by the server system, a public key having an associated private key, and associating a logging wrapper having a respective session key with the digital record, after validating the received request, wherein the session key is distinct from the public key and the private key;

encrypting and sending, by the server system, the requested digital record which has been validated, using the public key and the session key to encrypt the digital record;

receiving and decrypting the encrypted digital record, by the remote computer, using the private key, and the session key in conjunction with the logging wrapper;

generating by the logging wrapper, at the remote computer, a logging event; and

recording the logging event in an access log.

179. The '368 patent does not attempt to preempt every application of the idea of controlling access to an encrypted digital record over a computer network.

180. The '368 patent does not preempt the field of electronically structuring and controlling access to protected data in a plurality of external databases. For example, the '368 patent includes inventive elements—embodied in specific claim limitations—that concretely circumscribe the patented invention and greatly limit its breadth. These inventive elements are not necessary or obvious tools for achieving secure third-party communications, and they ensure that the claims do not preempt other techniques for secure communications.

181. For example, the '368 patent describes numerous techniques for electronically structuring and controlling access to protected data in a plurality of external databases. The techniques inform the invention's development but do not, standing alone, fall within the scope of its claims:

- Rights-Based Access to Database Records. U.S. Pat. No. 5,325,294 to Keene, relates to a system that receives and stores the individual's medical information, after the individual is tested to establish this information and the date on which such information was most recently obtained.
- Security Tokens. U.S. Patent No. 5,978,918 to Scholnick, discloses a back-end process returns a time sensitive token that the "sender" sends to the "receiver." The "receiver" takes the time sensitive token and uses it to retrieve the private data.⁵⁴

⁵⁴ See also Arindam Khaled et al., *A Token-based Access Control System for RDF Data in the Clouds* at 104, in PROCEEDINGS OF THE 2ND IEEE INTERNATIONAL CONFERENCE ON CLOUD COMPUTING TECHNOLOGY AND SCIENCE (2010) (discussing the use of a "token-based access control system . . . implemented in Hadoop (an open source cloud computing framework)").

- Role-Based Access. U.S. Pat. No. 6,023,765 to Kuhn, relates to a role-based access control in multi-level secure systems.
- Secure Networks. U.S. Pat. No. 5,579,393 to Conner, relates to a system and method for secure digital records, comprising a provider system and a payer system.
- Cryptographic Technology. U.S. Pat. No. 5,956,408 to Arnold, relates to an apparatus and method for secure distribution of data. Data, including program and software updates, is encrypted by a public key encryption system using the private key of the data sender.
- Watermarking. U.S. Pat. No. 5,699,427 to Chow, relates to a method to deter document and intellectual property piracy through individualization, and a system for identifying the authorized receiver of any particular copy of a document.
- Computer System Security. U.S. Pat. No. 5,881,225 to Worth, relates to a security monitor for controlling functional access to a computer system. A security monitor controls security functions for a computer system. A user desiring access to the system inputs a user identification and password combination, and a role the user to assume is selected from among one or more roles defined in the system.
- Computer Security Devices. U.S. Pat. No. 5,982,520 to Weiser, relates to a personal storage device for receipt, storage, and transfer of digital information to other electronic devices has a pocket sized crush resistant casing with a volume of less than about ten cubic centimeters.
- Computer Network Firewall. U.S. Pat. No. 5,944,823 to Jade, relates to a system and method for providing outside access to computer resources through a firewall. A firewall isolates computer and network resources inside the firewall from networks, computers and computer applications outside the firewall.
- Virtual Private Network. U.S. Pat. No. 6,079,020 to Liu, relates to a method and an apparatus for managing a virtual private network operating over a public data network. This public data network has been augmented to include a plurality of virtual private network gateways so that communications across the virtual private network are channeled through the virtual private network gateways.
- Biometric Authentication. U.S. Pat. No. 5,193,855 to Shamos relates to a patient and healthcare provider identification system which includes a database of patient and healthcare provider information including the identity of each patient and provider and some identification criteria (such as fingerprint data).⁵⁵

⁵⁵ Nary Subramanian, *Biometric Authentication*, in ENCYCLOPEDIA OF CRYPTOGRAPHY AND SECURITY (S. Jajodia and H.C.A. van Tilborg 2nd ed. 2011) (“Biometric authentication is a

182. Although “[e]ncryption, in general, represents a basic building block of human ingenuity that has been used for hundreds, if not thousands, of years,”⁵⁶ the ‘368 patent does not claim, or attempt to preempt, “some process that involves the encryption of data for some purpose” (or similar abstraction).

183. The ‘368 patent does not claim, or attempt to preempt, the performance of an abstract business practice on the Internet or using a conventional computer.

184. The claimed subject matter of the ‘368 patent is not a pre-existing but undiscovered algorithm.

185. The ‘368 patent claims systems and methods that “could not conceivably be performed in the human mind or pencil and paper.”⁵⁷

186. The ‘368 patent claims require the use of a computer system.

187. The ‘368 patent claims systems and methods not merely for transferring secure information over a computer network, but for making the computer network itself more secure.

188. The claimed invention in the ‘368 claims is rooted in computer technology and overcomes problems specifically arising in the realm of computer networks.

189. The systems and methods claimed in the ‘368 patent were not a longstanding or fundamental economic practice at the time of the patented inventions. Nor were they fundamental principles in ubiquitous use on the Internet or computers in general.


technique for identifying the person accessing a secured asset . . . by comparing their unique biological features . . . [an] issue with biometric authentication is privacy of personal data.”).

⁵⁶ *Paone v. Broadcom Corp.*, Case No. 15 Civ. 0596-BMC-GRB, 2015 WL 4988279, at *7 (E.D.N.Y. Aug. 19, 2015) (citing *Fid. Nat'l Info. Servs., Inc.*, Petitioner, CBM2014-00021, 2015 WL 1967328, at *8 (Apr. 29, 2015) (both upholding the patent eligibility of patents directed toward encryption).

⁵⁷ *TQP Dev., LLC v. Intuit Inc.*, Case No. 2:12-CV-180-WCB, 2014 WL 651935, at *4 (E.D. Tex. Feb. 19, 2014) (finding claims directed to encryption to be patent eligible); see also *Paone v. Broadcom Corp.*, Case No. 15 CIV. 0596 BMC GRB, 2015 WL 4988279, at *7 (E.D.N.Y. Aug. 19, 2015); see also *Prism Technologies, LLC v. T-Mobile USA, Inc.*, 12-cv-124, Dkt. No. 428 at 7 (D. Neb. Sept. 22, 2015) (Finding on cross motions for summary judgment that patents directed at delivering resources over an untrusted network were patent eligible. “The problems addressed by Prism’s claims are ones that ‘arose uniquely in the context of the Internet.’”).

190. The asserted claims do not involve a method of doing business that happens to be implemented on a computer; instead, they involve a method for changing digital records in a way that will affect the communication system itself, by making it more secure.

191. One or more claims of the '368 patent require a specific configuration of electronic devices, a network configuration, and the use of encryption systems to secure communications and manage access to secure digital records. These are meaningful limitations that tie the claimed methods and systems to specific machines. For example, the below diagram from the '368 patent illustrates a specific configuration of hardware disclosed in the patent.



'368 patent, Fig. 1.

3. U.S. Patent No. 8,498,941

192. U.S. Patent No. 8,498,941 (the "'941 patent'") entitled, Information Record Infrastructure, System and Method, was filed on July 22, 2009, and claims priority to July 6, 2000. St. Luke is the owner by assignment of the '941 patent. A true and correct copy of the '941 patent is attached hereto as Exhibit E. The '941 patent claims specific methods and systems

for securely controlling access to a plurality of digital records by a remote computer where each record has associated access rules.

193. The '941 patent has been cited by 10 United States patents and patent applications as relevant prior art. Specifically, patents issued to the following companies have cited the '941 patent as relevant prior art.

- Red Hat, Inc.
- Intuit, Inc.
- Microsoft Corporation
- Silver Spring Networks, Inc.
- Royal Canadian Mint
- Extendabrain Corporation

194. The '941 patent claims a technical solution to a problem unique to computer networks – controlling access to a plurality of records provided within a plurality of automated electronic databases, each record having an associated set of access rules.

195. At the time of the inventions claimed in the '941 patent, electronically structuring and controlling access to protected data in a plurality of external databases presented new and unique issues over the state of the art. As explained in the '941 patent: “The present invention therefore seeks to provide a comprehensive set of technologies to address the full scope of issues presented in implementing a secure and versatile information content infrastructure that respects the rights of content owners and users to privileges, such as confidentiality.” '941 patent, col. 53:35-39.

196. Although the systems and methods taught in the '941 patent have been adopted by leading businesses today, at the time of invention, the technologies taught in the '941 patent claims were innovative and novel. “Existing systems do not create a trust infrastructure, wherein an independent third party represents and serves as an agent for the content owner, implementing a set of restrictive rules for use of content . . . Thus, existing intermediaries do not act in a representative capacity for the content owner, and do not integrate content management functions.” '941 patent, col. 5:17-20.

197. Further, the '941 patent claims improve upon the functioning of a computer system by allowing encrypted electronic data to be securely transmitted through an intermediary. This improves the security of the computer system and allows it to be more efficient. “[B]y consolidating a plurality of institutions [referring to digital records stored in external databases], uniformity, interoperability, cost reductions, and improved security result.” ‘941 patent, col. 66:21-23.

198. The ‘941 patent claims are not directed to a “method of organizing human activity,” “fundamental economic practice long prevalent in our system of commerce,” or “a building block of the modern economy.” Instead, they are limited to a concretely circumscribed set of methods and systems that provide a conduit for the authorized transmission of digital records, while maintaining the security of the records against unauthorized access.

199. The ‘941 patent claims are not directed at the broad concept/idea of “managing digital records.” Instead, the ‘941 patent claims are limited to a concretely circumscribed set of methods and systems for authorizing and transmitting secure digital records. These methods and systems are technologies unique to the Internet age.

200. The ‘941 patent claims are directed toward a solution rooted in computer technology and use technology unique to computers and computer networking to overcome a problem specifically arising in the realm of secure distributed computing. For example, one or more claims of the ‘941 patent require the generation of an information polymer - a procedure that overrides the routine and conventional sequence of events in electronic communications, even encrypted electronic communications.

201. The ‘941 patent is directed to specific problems in the field of digital record access and transmission.

202. The preemptive effect of the claims of the ‘941 patent are concretely circumscribed by specific limitations. For example, claim 1 of the ‘941 patent requires:

A method for controlling access to a plurality of records provided within a plurality of automated electronic databases, each record having an associated set of access rules, comprising:

receiving a request from a requestor, the requestor having at least one attribute;

searching the plurality of automated electronic databases to find records in dependence on the request and on connections between respective records;

applying a set of access rules associated with each found record by at least one automated processor, to produce a set of accessible records;

linking the set of accessible records into an information polymer using a server device;

applying at least one compensation rule by at least one automated processor, dependent on the at least one attribute of the requestor;

logging at least the request for access by at least one automated processor; and

communicating the information polymer to the requestor.

203. The '941 patent does not attempt to preempt every application of the idea of controlling access to a digital record over a computer network where the digital records are within a plurality of automated electronic databases.

204. The '941 patent does not preempt the field of electronically structuring and controlling access to protected data in a plurality of external databases. For example, the '941 patent includes inventive elements—embodied in specific claim limitations—that concretely circumscribe the patented invention and greatly limit its breadth. These inventive elements are not necessary or obvious tools for achieving secure third-party communications, and they ensure that the claims do not preempt other techniques for secure communications.

205. For example, the '941 patent describes numerous techniques for electronically structuring and controlling access to protected data in a plurality of external databases. The techniques inform the invention's development but do not, standing alone, fall within the scope of its claims:

- Rights-Based Access to Database Records. U.S. Pat. No. 5,325,294 to Keene, relates to a system that receives and stores the individual's medical information, after the individual is tested to establish this information and the date on which such information was most recently obtained

- Role-Based Access. U.S. Pat. No. 6,023,765 to Kuhn, relates to a role-based access control in multi-level secure systems.
- Security Tokens. U.S. Patent No. 5,978,918 to Scholnick, discloses a back-end process returns a time sensitive token that the “sender” sends to the “receiver.” The “receiver” takes the time sensitive token and uses it to retrieve the private data.⁵⁸
- Secure Networks. U.S. Pat. No. 5,579,393 to Conner, relates to a system and method for secure digital records, comprising a provider system and a payer system.
- Cryptographic Technology. U.S. Pat. No. 5,956,408 to Arnold, relates to an apparatus and method for secure distribution of data. Data, including program and software updates, is encrypted by a public key encryption system using the private key of the data sender.
- Watermarking. U.S. Pat. No. 5,699,427 to Chow, relates to a method to deter document and intellectual property piracy through individualization, and a system for identifying the authorized receiver of any particular copy of a document.
- Computer System Security. U.S. Pat. No. 5,881,225 to Worth, relates to a security monitor for controlling functional access to a computer system. A security monitor controls security functions for a computer system. A user desiring access to the system inputs a user identification and password combination, and a role the user to assume is selected from among one or more roles defined in the system.
- Computer Security Devices. U.S. Pat. No. 5,982,520 to Weiser, relates to a personal storage device for receipt, storage, and transfer of digital information to other electronic devices has a pocket sized crush resistant casing with a volume of less than about ten cubic centimeters.
- Computer Network Firewall. U.S. Pat. No. 5,944,823 to Jade, relates to a system and method for providing outside access to computer resources through a firewall. A firewall isolates computer and network resources inside the firewall from networks, computers and computer applications outside the firewall.
- Virtual Private Network. U.S. Pat. No. 6,079,020 to Liu, relates to a method and an apparatus for managing a virtual private network operating over a public data network.

⁵⁸ See also Arindam Khaled et. al, *A Token-based Access Control System for RDF Data in the Clouds* at 104, in PROCEEDINGS OF THE 2ND IEEE INTERNATIONAL CONFERENCE ON CLOUD COMPUTING TECHNOLOGY AND SCIENCE (2010) (discussing the use of a “token-based access control system . . . implemented in Hadoop (an open source cloud computing framework)”).

This public data network has been augmented to include a plurality of virtual private network gateways so that communications across the virtual private network are channeled through the virtual private network gateways.

- Biometric Authentication. U.S. Pat. No. 5,193,855 to Shamos relates to a patient and healthcare provider identification system which includes a database of patient and healthcare provider information including the identity of each patient and provider and some identification criteria (such as fingerprint data).⁵⁹

206. The '941 patent does not claim, or attempt to preempt, the performance of an abstract business practice on the Internet or using a conventional computer.

207. The claimed subject matter of the '941 patent is not a pre-existing but undiscovered algorithm.

208. The '941 patent claims require the use of a computer system.

209. The '941 patent claims systems and methods not merely for transferring secure information over a computer network, but for making the computer network itself more secure.

210. The claimed invention in the '941 claims is rooted in computer technology and overcomes problems specifically arising in the realm of computer networks.

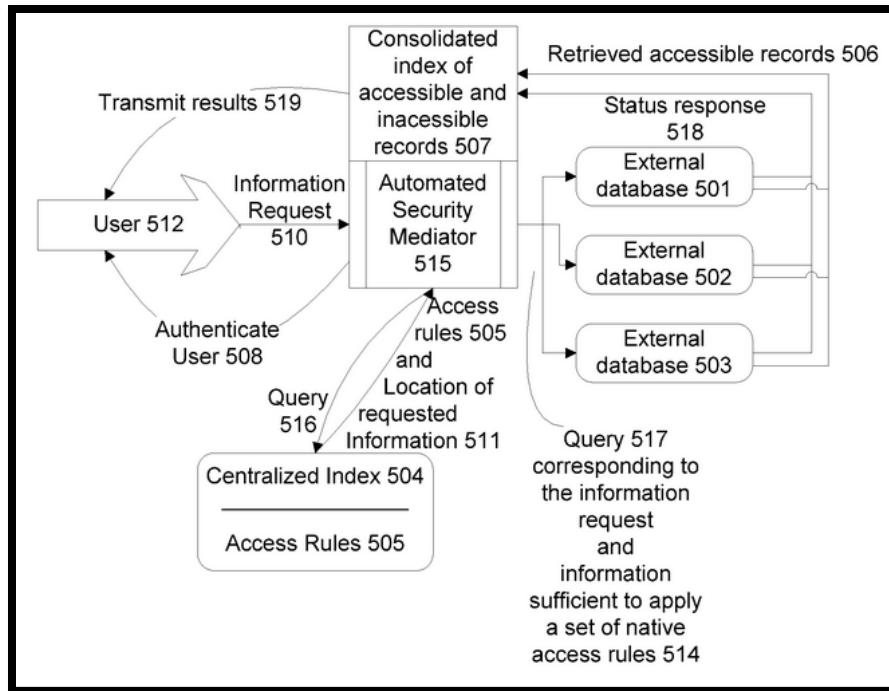
211. The systems and methods claimed in the '941 patent were not a longstanding or fundamental economic practice at the time of patented inventions. Nor were they fundamental principles in ubiquitous use on the Internet or computers in general.

212. The asserted claims do not involve a method of doing business that happens to be implemented on a computer; instead, they involve a method for changing digital records in a way that will affect the communication system itself, by making it more secure.

213. One or more claims of the '941 patent require a specific configuration of electronic devices, a network configuration, and the use of encryption systems to secure communications and manage access to secure digital records. These are meaningful limitations

⁵⁹ Nary Subramanian, *Biometric Authentication*, in *ENCYCLOPEDIA OF CRYPTOGRAPHY AND SECURITY* (S. Jajodia and H.C.A. van Tilborg 2nd ed. 2011) (“Biometric authentication is a technique for identifying the person accessing a secured asset . . . by comparing their unique biological features . . . [an] issue with biometric authentication is privacy of personal data.”).

that tie the claimed methods and systems to specific machines. For example, the below diagram from the '941 patent illustrates a specific configuration of hardware disclosed in the patent.



'941 patent, Fig. 6.

4. U.S. Patent No. 8,380,630

214. U.S. Patent No. 8,380,630 (the “‘630 patent”) entitled, Information Record Infrastructure, System and Method, was filed on May 29, 2010, and claims priority to July 6, 2000. St. Luke is the owner by assignment of the ‘630 patent. A true and correct copy of the ‘630 patent is attached hereto as Exhibit F. The ‘630 patent claims specific methods and systems for securely controlling access to a plurality of digital records by a remote computer, using a security mediator, where each record has associated access rules.

215. The ‘630 patent has been cited by ten United States patents and published patent applications as relevant prior art. Specifically, patents issued to the following companies have cited the ‘630 patent as relevant prior art.

- Informatica Corporation
- Electronics and Telecommunications Research Institute (“ETRI”)
- J.D. Power and Associates
- CA, Inc.

- Microsoft Corporation

216. The '630 patent claims a technical solution to a problem unique to computer networks – controlling access to a plurality of records provided within a plurality of automated electronic databases, each record having an associated set of access rules.

217. At the time of the inventions claimed in the '630 patent, electronically structuring and controlling access to protected data in a plurality of external databases presented new and unique issues over the state of the art. As explained in the '630 patent: “The present invention therefore seeks to provide a comprehensive set of technologies to address the full scope of issues presented in implementing a secure and versatile information content infrastructure that respects the rights of content owners and users to privileges, such as confidentiality.” '630 patent, col. 53:45-49.

218. Although the systems and methods taught in the '630 patent have been adopted by leading businesses today, at the time of invention, the technologies taught in the '630 patent claims were innovative and novel. “Existing systems do not create a trust infrastructure, wherein an independent third party represents and serves as an agent for the content owner, implementing a set of restrictive rules for use of content . . . Thus, existing intermediaries do not act in a representative capacity for the content owner, and do not integrate content management functions.” '630 patent, col. 5:11-23.

219. Further, the '630 patent claims improve upon the functioning of a computer system by allowing encrypted electronic data to be securely transmitted through an intermediary. This improves the security of the computer system and allows it to be more efficient. “[B]y consolidating a plurality of institutions [referring to digital records stored in external databases], uniformity, interoperability, cost reductions, and improved security result.” '630 patent, col. 66:33-35.

220. The '630 patent claims require an automated security mediator (“ASM”).

221. The '630 patent claims require the ASM query the automated centralized index (“ACI”) to locate the record information within a plurality of external databases.

222. The '630 patent claims require that the ASM generate an index of accessible location record information that is available in a plurality of externally databases.

223. The '630 patent claims are not directed to a “method of organizing human activity,” “fundamental economic practice long prevalent in our system of commerce,” or “a building block of the modern economy.” Instead, they are limited to a concretely circumscribed set of methods and systems that provide a conduit for the authorized transmission of digital records, while maintaining the security of the records against unauthorized access.

224. The '630 patent claims are not directed at the broad concept/idea of “managing digital records.” Instead, the '630 patent claims are limited to a concretely circumscribed set of methods and systems for authorizing and transmitting secure digital records. These methods and systems are technologies unique to the Internet age.

225. The '630 patent claims are directed toward a solution rooted in computer technology and use technology unique to computers and computer networking to overcome a problem specifically arising in the realm of secure distributed computing. For example, one or more claims of the '630 patent require an ASM, require the generation of an Automated Centralized Index (“ACI”), require applying the access rules associated with the located requested information (“LRI”), require the ASM query the ACI to locate the record information within the plurality of external databases, and require that the ASM generate an index of LRI accessible in a plurality of external databases - a procedure that overrides the routine and conventional sequence of events in electronic communications.

226. The '630 patent is directed to specific problems in the field of digital record access and transmission.

227. The preemptive effect of the claims of the '630 patent are concretely circumscribed by specific limitations. For example, claim 1 of the '630 patent requires:

A method for security mediation, comprising:

receiving an information request for information stored within a plurality of external databases (“POEDs”) from a user, wherein the information request is received by an automated security mediator

(“ASM”) which is neither an owner nor custodian of the requested information;

authenticating the user;

querying an automated centralized index (“ACI”), maintained by the ASM to locate the requested information within the POEDs, wherein the ACI includes a location and a set of access rules for each entry;

applying the access rules associated with the located requested information (“LRI”);

automatically communicating from the ASM to each of the POEDs storing the LRI: a query corresponding to the information request, and information sufficient to apply a set of native access rules of the respective POEDs storing the LRI to further control access to the LRI;

receiving at least a status response from at least one of the POEDs storing the LRI indicating whether the LRI is accessible or inaccessible;

automatically indexing the accessible and inaccessible LRI; and

at least one of:

retrieving, by the ASM, the accessible LRI from the POEDs storing the LRI and communicating, from the ASM to the user a consolidation of the retrieved accessible LRI; and

communicating, from the ASM to the user a consolidated index of the accessible LRI.

228. The ‘630 patent does not attempt to preempt every application of the idea of controlling access to a digital record over a computer network where the digital records are within a plurality of automated electronic databases.

229. The ‘630 patent does not preempt the field of electronically structuring and controlling access to protected data in a plurality of external databases. For example, the ‘630 patent includes inventive elements—embodied in specific claim limitations—that concretely circumscribe the patented invention and greatly limit its breadth. These inventive elements are not necessary or obvious tools for achieving secure third-party communications, and they ensure that the claims do not preempt other techniques for secure communications.

230. For example, the ‘630 patent describes numerous techniques for electronically structuring and controlling access to protected data in a plurality of external databases. The

techniques inform the invention's development but do not, standing alone, fall within the scope of its claims:

- Rights-Based Access to Database Records. U.S. Pat. No. 5,325,294 to Keene, relates to a system that receives and stores the individual's medical information, after the individual is tested to establish this information and the date on which such information was most recently obtained
- Role-Based Access. U.S. Pat. No. 6,023,765 to Kuhn, relates to a role-based access control in multi-level secure systems.
- Secure Networks. U.S. Pat. No. 5,579,393 to Conner, relates to a system and method for secure digital records, comprising a provider system and a payer system.
- Cryptographic Technology. U.S. Pat. No. 5,956,408 to Arnold, relates to an apparatus and method for secure distribution of data. Data, including program and software updates, is encrypted by a public key encryption system using the private key of the data sender.
- Watermarking. U.S. Pat. No. 5,699,427 to Chow, relates to a method to deter document and intellectual property piracy through individualization, and a system for identifying the authorized receiver of any particular copy of a document.
- Computer System Security. U.S. Pat. No. 5,881,225 to Worth, relates to a security monitor for controlling functional access to a computer system. A security monitor controls security functions for a computer system. A user desiring access to the system inputs a user identification and password combination, and a role the user to assume is selected from among one or more roles defined in the system.
- Computer Security Devices. U.S. Pat. No. 5,982,520 to Weiser, relates to a personal storage device for receipt, storage, and transfer of digital information to other electronic devices has a pocket sized crush resistant casing with a volume of less than about ten cubic centimeters.
- Computer Network Firewall. U.S. Pat. No. 5,944,823 to Jade, relates to a system and method for providing outside access to computer resources through a firewall. A firewall isolates computer and network resources inside the firewall from networks, computers and computer applications outside the firewall.
- Virtual Private Network. U.S. Pat. No. 6,079,020 to Liu, relates to a method and an apparatus for managing a virtual private network operating over a public data network.

This public data network has been augmented to include a plurality of virtual private network gateways so that communications across the virtual private network are channeled through the virtual private network gateways.

- Biometric Authentication. U.S. Pat. No. 5,193,855 to Shamos relates to a patient and healthcare provider identification system which includes a database of patient and healthcare provider information including the identity of each patient and provider and some identification criteria (such as fingerprint data).

231. The '630 patent does not claim, or attempt to preempt, the performance of an abstract business practice on the Internet or using a conventional computer.

232. The claimed subject matter of the '630 patent is not a pre-existing but undiscovered algorithm.

233. The '630 patent claims require the use of a computer system.

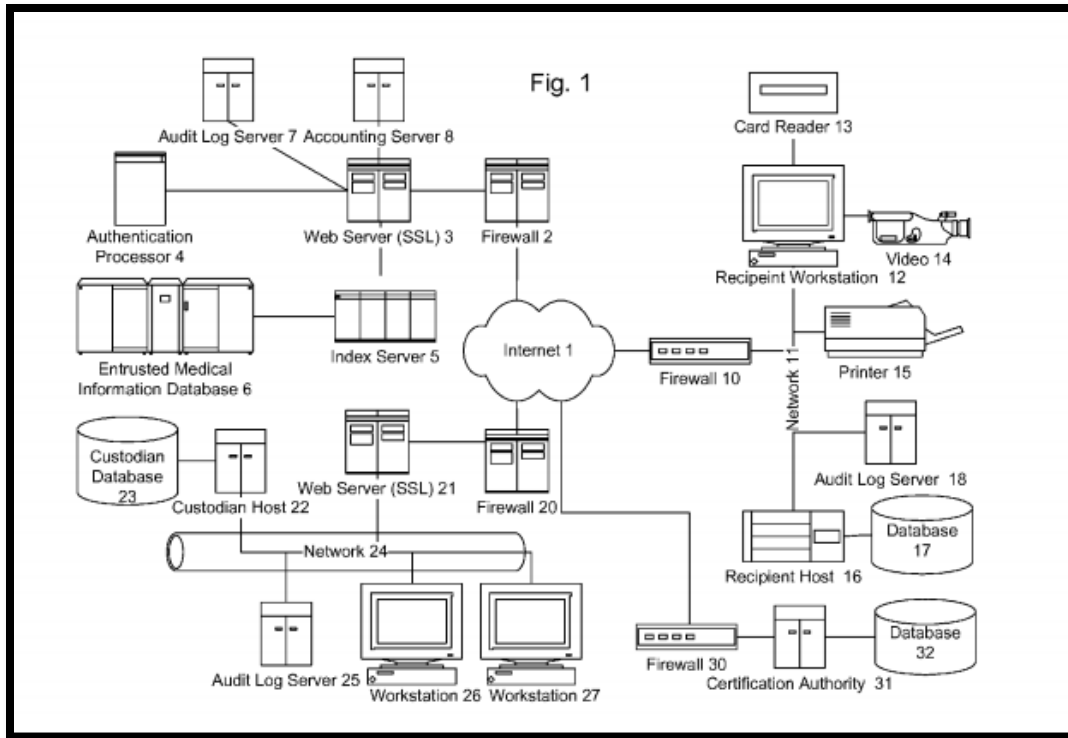
234. The '630 patent claims systems and methods not merely for transferring secure information over a computer network, but for making the computer network itself more secure.

235. The claimed invention in the '630 claims is rooted in computer technology and overcomes problems specifically arising in the realm of computer networks.

236. The systems and methods claimed in the '630 patent were not a longstanding or fundamental economic practice at the time of the patented inventions. Nor were they fundamental principles in ubiquitous use on the Internet or computers in general.

237. The asserted claims do not involve a method of doing business that happens to be implemented on a computer; instead, it involves a method for changing digital records in a way that will affect the communication system itself, by making it more secure.

238. One or more claims of the '630 patent require a specific configuration of electronic devices, a network configuration, and the use of encryption systems to secure communications and manage access to secure digital records. These are meaningful limitations that tie the claimed methods and systems to specific machines. For example, the below diagram from the '630 patent illustrates a specific configuration of hardware disclosed in the patent.



'630 patent, Fig. 1.

5. **U.S. Patent No. 8,600,895**

239. U.S. Patent No. 8,600,895 (the "'895 patent") entitled, Information Record Infrastructure, System and Method, was filed on February 19, 2013, and claims priority to July 6, 2000. St. Luke is the owner by assignment of the '895 patent. A true and correct copy of the '895 patent is attached hereto as Exhibit G. The '895 patent claims specific methods and systems for securely controlling access to a plurality of digital records by a remote computer, using a security mediator, where each record has associated access rules.

240. The '895 patent has been cited by four United States patents and patent applications as relevant prior art.⁶⁰ Specifically, patents issued to the following companies have cited the '895 patent as relevant prior art.

- J.D. Power and Associates

⁶⁰ Although the '895 patent has only been cited 4 times, the patent applications to which the '895 patent claims priority have been cited by hundreds of companies. U.S. Patent Application 12/790,818 was cited in 45 issued patents and published patent applications, U.S. Patent Application was cited in 27 patents and published patent applications, and U.S. Patent Application 09/899,787 was cited in 751 patents and published patent applications.

- Fujitsu Limited
- Extendabrain Corporation

241. The '895 patent claims a technical solution to a problem unique to computer networks – controlling access to a plurality of records provided within a plurality of automated electronic databases, each record having an associated set of access rules.

242. At the time of the inventions claimed in the '895 patent, electronically structuring and controlling access to protected data in a plurality of external databases presented new and unique issues over the state of the art. As explained in the '895 patent: “The present invention therefore seeks to provide a comprehensive set of technologies to address the full scope of issues presented in implementing a secure and versatile information content infrastructure that respects the rights of content owners and users to privileges, such as confidentiality.” '895 patent, col. 53:53-57.

243. Although the systems and methods taught in the '895 patent have been adopted by leading businesses today, at the time of invention, the technologies taught in the '895 patent claims were innovative and novel. “Existing systems do not create a trust infrastructure, wherein an independent third party represents and serves as an agent for the content owner, implementing a set of restrictive rules for use of content . . . Thus, existing intermediaries do not act in a representative capacity for the content owner, and do not integrate content management functions.” '895 patent, col. 5:18-30.

244. Further, the '895 patent claims improve upon the functioning of a computer system by allowing encrypted electronic data to be securely transmitted through an intermediary. This improves the security of the computer system and allows it to be more efficient. “[B]y consolidating a plurality of institutions [referring to digital records stored in external databases], uniformity, interoperability, cost reductions, and improved security result.” '895 patent, col. 66:41-44.

245. The '895 patent claims require controlling access to a plurality of records stored within a plurality of automated external databases.

246. The '895 patent claims require an automated centralized index ("ACI") that includes, for each record, a (1) location identifier (LI), (2) content identifier (CI), and (3) associated set of access rules (ASAR).

247. The '895 patent claims require logically associating the releasable accessible record ("AR") into a linked set of releasable ARs (LAS) and communicating the LAS to the requestor.

248. The '895 patent claims are not directed to a "method of organizing human activity," "fundamental economic practice long prevalent in our system of commerce," or "a building block of the modern economy." Instead, they are limited to a concretely circumscribed set of methods and systems that provide a conduit for the authorized transmission of digital records, while maintaining the security of the records against unauthorized access.

249. The '895 patent claims are not directed at the broad concept/idea of "managing digital records." Instead, the '895 patent claims are limited to a concretely circumscribed set of methods and systems for authorizing and transmitting secure digital records. These methods and systems are technologies unique to the Internet age.

250. The '895 patent claims are directed toward a solution rooted in computer technology and use technology unique to computers and computer networking to overcome a problem specifically arising in the realm of secure distributed computing. For example, one or more claims of the '895 patent require an ACI, require a content identifier ("CI"), require querying ACI to find entries containing CI, require for each accessible record (AR) communicate to the plurality of external databases information sufficient for the external databases to apply native access rules to determine whether the AR is releasable.

251. The '895 patent is directed to specific problems in the field of digital record access and transmission.

252. The preemptive effect of the claims of the '895 patent are concretely circumscribed by specific limitations. For example, claim 16 of the '895 patent requires:

An apparatus for controlling access to a plurality of records stored within a plurality of automated external databases (“AXES”), comprising:

- an automated centralized index (“ACI”), stored in a memory, configured to store an entry for each record consisting of a location identifier (“LI”), an associated set of access rules (“ASAR”), and a content identifier (“CI”);

- an input port configured to receive a request from a requestor for access to one or more records stored in the plurality of AXES, wherein the request specifies a CI with which to query the ACI;

- at least one processor configured to:

 - generate a query based on the specified CI (“SCI”);

 - find entries in the ACI containing the SCI;

 - for each found entry, apply the ASAR corresponding to the LI to determine if the record stored in a respective one of the AXES corresponding to the LI is accessible;

 - generate a communication, for communication to the respective one of the AXES storing an accessible record (“AR”), wherein the communication contains information sufficient for the respective one of the AXES storing the AR to apply a set of native access rules (“NAR”) it maintains to determine if the AR is releasable;

 - form a linked set of releasable ARs by logically associating the releasable ARs; and

 - generate a communication containing the linked set of releasable ARs; and

- at least one communications port configured to communicate:

 - the generated communication to the respective one of the AXES storing the ARs; and

 - the linked set of releasable ARs.

253. The ‘895 patent does not attempt to preempt every application of the idea of controlling access to a digital record over a computer network where the digital records are within a plurality of automated electronic databases.

254. The ‘895 patent does not preempt the field of electronically structuring and controlling access to protected data in a plurality of external databases. For example, the ‘895 patent includes inventive elements—embodied in specific claim limitations—that concretely circumscribe the patented invention and greatly limit its breadth. These inventive elements are not necessary or obvious tools for achieving secure third-party communications, and they ensure that the claims do not preempt other techniques for secure communications.

255. For example, the '895 patent describes numerous techniques for electronically structuring and controlling access to protected data in a plurality of external databases. The techniques inform the invention's development but do not, standing alone, fall within the scope of its claims:

- Rights-Based Access to Database Records. U.S. Pat. No. 5,325,294 to Keene, relates to a system that receives and stores the individual's medical information, after the individual is tested to establish this information and the date on which such information was most recently obtained
- Role-Based Access. U.S. Pat. No. 6,023,765 to Kuhn, relates to a role-based access control in multi-level secure systems.
- Secure Networks. U.S. Pat. No. 5,579,393 to Conner, relates to a system and method for secure digital records, comprising a provider system and a payer system.
- Cryptographic Technology. U.S. Pat. No. 5,956,408 to Arnold, relates to an apparatus and method for secure distribution of data. Data, including program and software updates, is encrypted by a public key encryption system using the private key of the data sender.
- Watermarking. U.S. Pat. No. 5,699,427 to Chow, relates to a method to deter document and intellectual property piracy through individualization, and a system for identifying the authorized receiver of any particular copy of a document.
- Computer System Security. U.S. Pat. No. 5,881,225 to Worth, relates to a security monitor for controlling functional access to a computer system. A security monitor controls security functions for a computer system. A user desiring access to the system inputs a user identification and password combination, and a role the user to assume is selected from among one or more roles defined in the system.
- Computer Security Devices. U.S. Pat. No. 5,982,520 to Weiser, relates to a personal storage device for receipt, storage, and transfer of digital information to other electronic devices has a pocket sized crush resistant casing with a volume of less than about ten cubic centimeters.
- Computer Network Firewall. U.S. Pat. No. 5,944,823 to Jade, relates to a system and method for providing outside access to computer resources through a firewall. A firewall isolates computer and network resources inside the firewall from networks, computers and computer applications outside the firewall.

- Virtual Private Network. U.S. Pat. No. 6,079,020 to Liu, relates to a method and an apparatus for managing a virtual private network operating over a public data network. This public data network has been augmented to include a plurality of virtual private network gateways so that communications across the virtual private network are channeled through the virtual private network gateways.
- Biometric Authentication. U.S. Pat. No. 5,193,855 to Shamos relates to a patient and healthcare provider identification system which includes a database of patient and healthcare provider information including the identity of each patient and provider and some identification criteria (such as fingerprint data).

256. The '895 patent does not claim, or attempt to preempt, the performance of an abstract business practice on the Internet or using a conventional computer.

257. The claimed subject matter of the '895 patent is not a pre-existing but undiscovered algorithm.

258. The '895 patent claims require the use of a computer system.

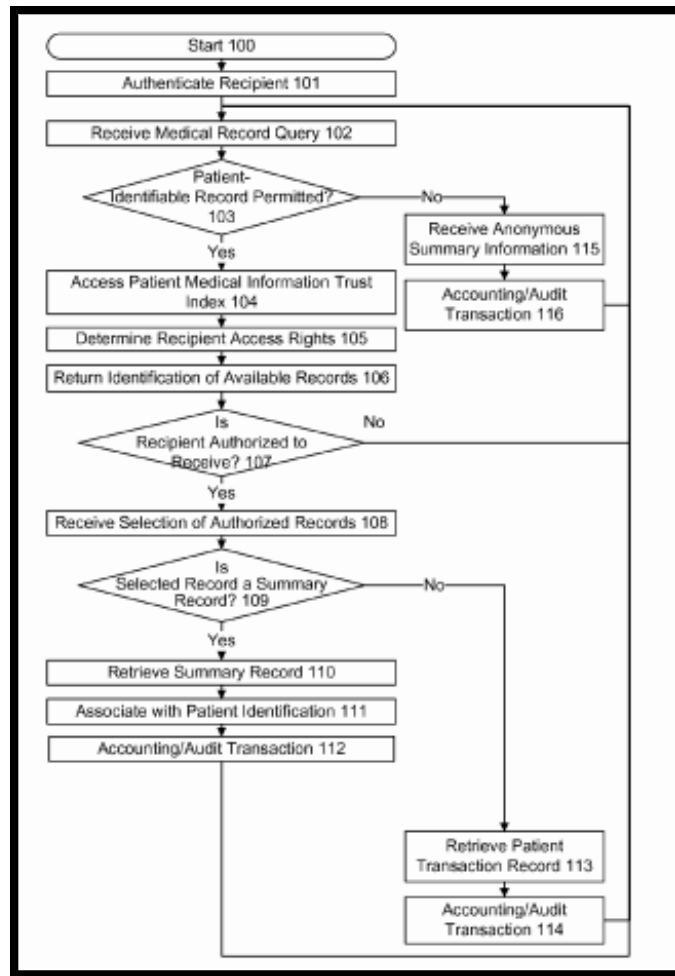
259. The '895 patent claims systems and methods not merely for transferring secure information over a computer network, but for making the computer network itself more secure.

260. The claimed invention in the '895 claims is rooted in computer technology and overcomes problems specifically arising in the realm of computer networks.

261. The systems and methods claimed in the '895 patent were not a longstanding or fundamental economic practice at the time of the patented inventions. Nor were they fundamental principles in ubiquitous use on the Internet or computers in general.

262. The asserted claims do not involve a method of doing business that happens to be implemented on a computer; instead, it involves a method for changing digital records in a way that will affect the communication system itself, by making it more secure.

263. One or more claims of the '895 patent require a specific configuration of electronic devices, a network configuration, and the use of access rules to secure communications and manage access to secure digital records. These are meaningful limitations that tie the claimed methods and systems to specific machines. For example, the below diagram from the '895 patent illustrates a specific configuration of hardware disclosed in the patent.



'895 patent, Fig. 4.

COUNT I
INFRINGEMENT OF U.S. PATENT NO. 8,316,237

264. St. Luke references and incorporates by reference the preceding paragraphs of this Complaint as if fully set forth herein.

265. Google designs, makes, sells, offers to sell, imports, and/or uses in the United States products and/or services for secure three-party communications.

266. Google designs, makes, sells, offers to sell, imports, and/or uses in the United States the Google Play Store platform, including but not limited to Google Play and/or Widevine packaging and license/key management server hardware and/or programming; Google Play web and mobile applications and services (e.g., Google Play Movies & TV applications and services

for Android, iOS, and Chrome); the Google Play website; and Google Play/Widevine secure client hardware and/or programming on end-user devices (e.g., Android smartphones, tablets, and streaming media devices; Chrome OS computers and streaming media devices; Widevine CDM; and/or plugin provisioned computers and mobile devices with Chrome Browser software and services) (collectively, “Google Play”).

267. Google designs, makes, sells, offers to sell, imports, and/or uses in the United States the YouTube platform, including but not limited to YouTube and/or Widevine packaging and license/key management server hardware and/or programming; YouTube web and mobile applications and services (e.g., YouTube applications and services for Android, iOS, and Chrome); the YouTube website; and YouTube/Widevine secure client hardware and/or programming on end-user devices (e.g., Android smartphones, tablets, and streaming media devices; Chrome OS computers and streaming media devices; Widevine CDM; and/or plugin provisioned computers and mobile devices with Chrome Browser software and services) (collectively, “YouTube”).

268. Google designs, makes, sells, offers to sell, imports, and/or uses in the United States the Widevine DRM platform, including but not limited to Widevine packaging, provisioning, and license/key management server hardware and/or programming; Widevine secure client libraries and programming embedded in and/or integrated with web and mobile applications and services (e.g., YouTube and Google Play Movies & TV applications and services for Android, iOS, and Chrome); and Widevine secure client hardware and/or programming on end-user devices (e.g., Android smartphones, tablets, and streaming media devices; Chrome OS computers and streaming media devices; Widevine CDM; and/or plugin provisioned computers and mobile devices with Chrome Browser software and services) (collectively, “Widevine”).

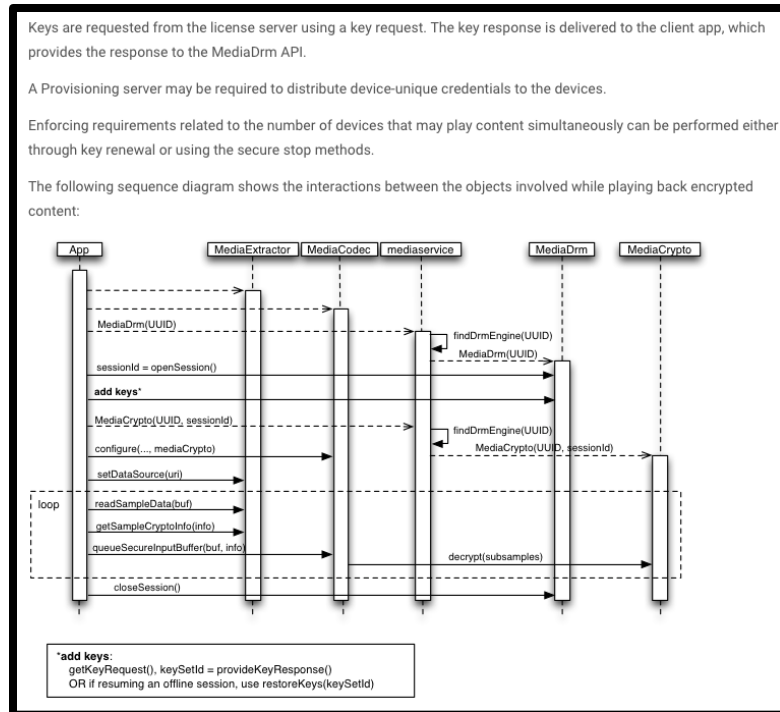
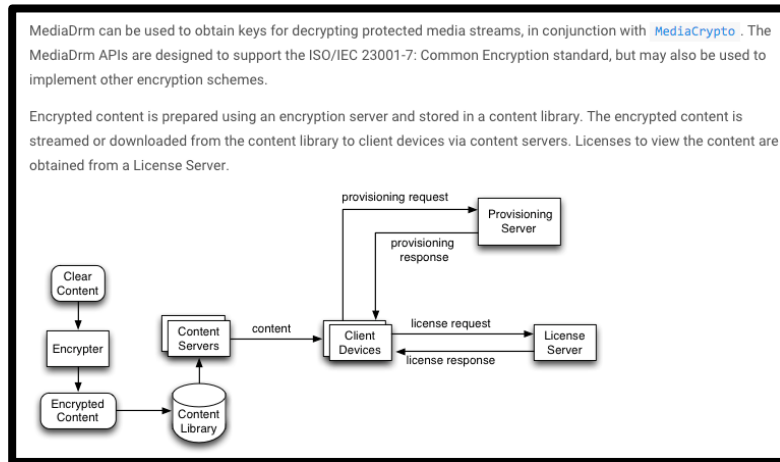
269. Google designs, makes, sells, offers to sell, imports, and/or uses in the United States Google Play, YouTube, and Widevine (collectively, the “Google ‘237 Products”).

270. On information and belief, the Google ‘237 Products comprise at least one transcription device (e.g., a Google Play, YouTube, and/or Widevine physical and/or virtual server appliance configured to transcribe encrypted Google Play and/or YouTube license, content, and/or cryptographic key communications).

271. On information and belief, the at least one transcription device (e.g., the Google Play, YouTube, and/or Widevine physical and/or virtual server appliance) comprises an automated communication port (e.g., an automated wired and/or wireless Internet Protocol communications port).

272. On information and belief, the automated communication port (e.g., the automated wired and/or wireless Internet Protocol communications port in the Google Play, YouTube, and/or Widevine physical and/or virtual server appliance) is configured to receive a first message (e.g., an Internet Protocol message sent to the Google Play, YouTube, and/or Widevine physical and/or virtual server appliance by and/or on behalf of a Google Play and/or YouTube Content Partner (“CP”) or a Certified Widevine Integrator Partner (“CWIP”)) representing an encrypted communication (e.g., an encrypted Google Play and/or YouTube license, content, and/or cryptographic key communication sent to the Google Play, YouTube, and/or Widevine physical and/or virtual server appliance by or on behalf of the CP or CWIP) associated with a first set of asymmetric keys (e.g., a set of RSA and/or Elliptical Curve asymmetric keys associated with the CP or CWIP).

273. On information and belief, the automated communication port (e.g., the automated wired and/or wireless Internet Protocol communications port in the Google Play, YouTube, and/or Widevine physical and/or virtual server appliance) is configured to receive a transcription key (e.g., a composite cryptographic key derived from and/or comprising at least CP- or CWIP-specific asymmetric decryption key information; client-specific asymmetric encryption key information; and session-specific cryptographic key information).



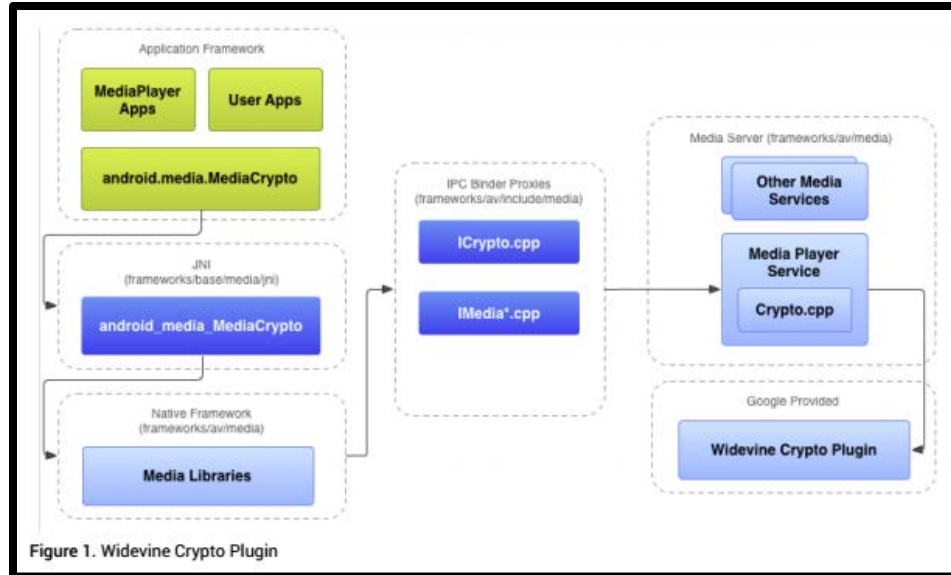
The app first constructs `MediaExtractor` and `MediaCodec` objects. It accesses the DRM-scheme-identifying UUID, typically from metadata in the content, and uses this UUID to construct an instance of a `MediaDrm` object that is able to support the DRM scheme required by the content. Crypto schemes are assigned 16 byte UUIDs. The method `isCryptoSchemeSupported(UUID)` can be used to query if a given scheme is supported on the device.

The app calls `openSession()` to generate a `sessionId` that will uniquely identify the session in subsequent interactions. The app next uses the `MediaDrm` object to obtain a key request message and send it to the license server, then provide the server's response to the `MediaDrm` object.

Once the app has a `sessionId`, it can construct a `MediaCrypto` object from the UUID and `sessionId`. The `MediaCrypto` object is registered with the `MediaCodec` in the `configure(MediaFormat, Surface, MediaCrypto, int)` method to enable the codec to decrypt content.

MediaDRM, ANDROID DEVELOPERS REFERENCE, <http://developer.android.com/reference/android/media/MediaDrm.html> (retrieved Oct. 13, 2015).

274. On information and belief, the automated communication port (e.g., the automated wired and/or wireless Internet Protocol communications port in the Google Play, YouTube, and/or Widevine physical and/or virtual server appliance) is configured to transmit a second message (e.g., an Internet Protocol message sent from the Google Play, YouTube, and/or Widevine physical and/or virtual server appliance to a Widevine-provisioned Google Play and/or YouTube client) representing the encrypted communication (e.g., the encrypted Google Play and/or YouTube license, content, and/or cryptographic key communication sent to the Google Play, YouTube, and/or Widevine physical and/or virtual server appliance by or on behalf of the CP or CWIP) associated with a second set of asymmetric keys (e.g., a set of RSA and/or Elliptical Curve asymmetric keys associated with the Widevine-provisioned Google Play and/or YouTube client).



Devices: DRM, ANDROID SOURCE, <http://source.android.com/devices/drm.html> (Sep. 14, 2014 version).

275. On information and belief, the first and second sets of encryption keys (e.g., the CP- or CWIP-specific set of asymmetric cryptographic keys and the Widevine-provisioned Google Play and/or YouTube client-specific set of asymmetric cryptographic keys, respectively) are distinct.

276. On information and belief, the at least one transcryption device (e.g., the Google Play, YouTube, and/or Widevine physical and/or virtual server appliance) comprises a memory (e.g., non-volatile flash and/or SSD memory).

277. On information and belief, the at least one transcryption device (e.g., the Google Play, YouTube, and/or Widevine physical and/or virtual server appliance) comprises an automated processor (e.g., a Google server automated processor).

278. On information and belief, the automated processor (e.g., the Google server automated processor in the Google Play, YouTube, and/or Widevine physical and/or virtual server appliance) is configured to communicate through the automated communication port (e.g., the automated wired and/or wireless Internet Protocol communications port in the Google Play, YouTube, and/or Widevine physical and/or virtual server appliance) and with the memory (e.g.,

the non-volatile flash and/or SSD memory in the Google Play, YouTube, and/or Widevine physical and/or virtual server appliance).

Google Infrastructure

Android devices, though powerful, rely on cloud-based services for much of their functionality. A large portion of the infrastructure behind these services is hosted by Google itself. The functionality provided by these services ranges from contact and e-mail data used by the phone dialer and Gmail to sophisticated remote management features. As such, these cloud services present an interesting attack surface, albeit not one that is usually reachable by a typical attacker. Many of these services are authenticated by Google's Single Sign On (SSO) system. Such a system lends itself to abuse because credentials stolen from one application could be used to access another application. This section discusses several relevant back-end infrastructure components and how they can be used to remotely compromise an Android device.

Google Infrastructure, ANDROID HACKER'S HANDBOOK at 149 (Wiley 2014).

279. On information and belief, the automated processor (e.g., the Google server automated processor in the Google Play, YouTube, and/or Widevine physical and/or virtual server appliance) is configured to receive the first message (e.g., the Internet Protocol message sent to the Google Play, YouTube, and/or Widevine physical and/or virtual server appliance by and/or on behalf of the CP or CWIP).

280. On information and belief, the automated processor (e.g., the Google server automated processor in the Google Play, YouTube, and/or Widevine physical and/or virtual server appliance) is configured to receive the transcription key (e.g., the composite cryptographic key derived from and/or comprising at least CP- or CWIP-specific asymmetric decryption key information; client-specific asymmetric encryption key information; and session-specific cryptographic key information).

281. On information and belief, the automated processor (e.g., the Google server automated processor in the Google Play, YouTube, and/or Widevine physical and/or virtual server appliance) is configured to automatically transcribe the first message (e.g., the Internet Protocol message sent to the Google Play, YouTube, and/or Widevine physical and/or virtual server appliance by and/or on behalf of the CP or CWIP) into the second message (e.g., the

Internet Protocol message sent from the Google Play, YouTube, and/or Widevine physical and/or virtual server appliance to a Widevine-provisioned Google Play and/or YouTube client).

2. Security Solution Robustness. With respect to the playback of High Definition Feature Films, the Content Protection System shall employ Licensor-approved tamper-resistant technology on hardware and software components (e.g., technology to prevent such hacks as a clock rollback, spoofing, use of common debugging tools, and intercepting unencrypted content in memory buffers). Examples of tamper resistant software techniques include, without limitation:

- a. *Code and data obfuscation*: The executable binary dynamically encrypts and decrypts itself in memory so that the algorithm is not unnecessarily exposed to disassembly or reverse engineering.
- b. *Integrity detection*: Using one-way cryptographic hashes of the executable code segments and/or self-referential integrity dependencies, the trusted software fails to execute and deletes all CSPs if it is altered prior to or during runtime.
- c. *Anti-debugging*: The decryption engine prevents the use of common debugging tools.
- d. *Red herring code*: The security modules use extra software routines that mimic security modules but do not have access to CSPs.

Schedule C: Content Protection Requirements and Obligations, CULVER DIGITAL DISTRIBUTION-GOOGLE CONTENT LICENSE AGREEMENT FOR YOUTUBE, GOOGLE, AND/OR ANDROID-BRANDED VOD SERVICES (HEREINAFTER, "Google/YouTube VOD License") (Mar. 17, 2011), publicly available at <https://wikileaks.org> (last accessed Oct. 4, 2015).

282. On information and belief, the automated processor (e.g., the Google server automated processor in the Google Play, YouTube, and/or Widevine physical and/or virtual server appliance) is configured to transmit the second message (e.g., the Internet Protocol message sent from the Google Play, YouTube, and/or Widevine physical and/or virtual server appliance to a Widevine-provisioned Google Play and/or YouTube client).

Nested Classes		
class	MediaDrm.CryptoSession	In addition to supporting decryption of DASH Common Encrypted Media, the MediaDrm APIs provide the ability to securely deliver session keys from an operator's session key server to a client device, based on the factory-installed root of trust, and then perform encrypt, decrypt, sign and verify operations with the session key on arbitrary user data.
class	MediaDrm.KeyRequest	Contains the opaque data an app uses to request keys from a license server
class	MediaDrm.KeyStatus	Defines the status of a key.
class	MediaDrm.MediaDrmStateException	Thrown when an unrecoverable failure occurs during a MediaDrm operation.
interface	MediaDrm.OnEventListener	Interface definition for a callback to be invoked when a drm event occurs
interface	MediaDrm.OnExpirationUpdateListener	Interface definition for a callback to be invoked when a drm session expiration update occurs
interface	MediaDrm.OnKeyStatusChangeListener	Interface definition for a callback to be invoked when the keys in a drm session change states.
class	MediaDrm.ProvisionRequest	Contains the opaque data an app uses to request a certificate from a provisioning server

MediaDRM, ANDROID DEVELOPERS REFERENCE, available at: <http://developer.android.com/reference/android/media/MediaDrm.html> (retrieved Oct. 13, 2015).

283. On information and belief, the automated processor (e.g., the Google server automated processor in the Google Play, YouTube, and/or Widevine physical and/or virtual server appliance) does not store as part of the transcription any decrypted representation of the

encrypted communication, and the transcription key (e.g., the composite cryptographic key derived from and/or comprising at least CP- or CWIP-specific asymmetric decryption key information; client-specific asymmetric encryption key information; and session-specific cryptographic key information) is employed without revealing any secret cryptographic information usable for decrypting the first message (e.g., the Internet Protocol message sent to the Google Play, YouTube, and/or Widevine physical and/or virtual server appliance by and/or on behalf of the CP or CWIP) or the second message (e.g., the Internet Protocol message sent from the Google Play, YouTube, and/or Widevine physical and/or virtual server appliance to a Widevine-provisioned Google Play and/or YouTube client).

In this implementation Widevine DRM keys and decrypted content are never exposed to the host CPU. Only security hardware or a protected security co-processor uses clear key values and the media content is decrypted by the secure hardware. This level of security requires factory provisioning of the Widevine key-box or requires the Widevine key-box to be protected by a device key installed at the time of manufacturing. The following describes some key points to this security level:

- Device manufacturers must provide a secure bootloader. The chain of trust from the bootloader must extend through any software or firmware components involved in the security implementation, such as the ARM TrustZone protected application and any components involved in the enforcement of the secure video path.
- The Widevine key-box must be encrypted with a device-unique secret key that is not visible to software or probing methods outside of the TrustZone.
- The Widevine key-box must be installed in the factory or delivered to the device using an approved secure delivery mechanism.
- Device manufacturers must provide an implementation of the Widevine Level 1 OEMCrypto API that performs all key processing and decryption in a trusted environment.

Interfaces: DRM, ANDROID SOURCE, <http://source.android.com/devices/drm.html> (Sep. 14, 2014 version).

284. On information and belief, the Google '237 Products comprise a secure distributed access control system (e.g., a distributed Google server system for securely controlling access to sensitive Google Play and/or YouTube license, content, and/or cryptographic key information).

Widevine DRM security levels

Security is never implemented in a single place in the stack, but instead relies on the integration of hardware, software, and services. The combination of hardware security functions, a trusted boot mechanism, and an isolated secure OS for handling security functions is critical to provide a secure device.

At the system level, Android offers the core security features of the Linux kernel, extended and customized for mobile devices. In the application framework, Android provides an extensible DRM framework and system architecture for checking and enforcing digital rights. The Widevine DRM plugin integrates with the hardware platform to leverage the available security capabilities. The level of security offered is determined by a combination of the security capabilities of the hardware platform and the integration with Android and the Widevine DRM plugin. Widevine DRM security supports the three levels of security shown in the table below.

Security Level	Secure Bootloader	Widevine Key Provisioning	Security Hardware or ARM Trust Zone	Widevine Keybox and Video Key Processing	Hardware Video Path
Level 1	Yes	Factory provisioned Widevine Keys	Yes	Keys never exposed in clear to host CPU	Hardware protected video path
Level 2	Yes	Factory provisioned Widevine Keys	Yes	Keys never exposed in clear to host CPU	Hardware protected video path

Interfaces: DRM, ANDROID SOURCE, available at: <http://source.android.com/devices/drm.html> (Sep. 14, 2014 version).

285. On information and belief, the secure distributed access control system (e.g., the distributed server system for securely controlling access to sensitive Google Play and/or YouTube license, content, and/or cryptographic key information) comprises a communication interface device (e.g., a Google Play, YouTube, and/or Widevine server-to-server Internet Protocol communication interface physical and/or virtual appliance) configured to communicate with a plurality of independently operating servers (e.g., a plurality of independently operating CP and/or CWIP servers), each communicating server encrypted information (e.g., CP and/or CWIP server-encrypted Google Play and/or YouTube license, content, and/or cryptographic key information), wherein the server encrypted information is in an encrypted form negotiated between an respective server (e.g., a respective CP and/or CWIP server) and an intermediary (e.g., a Google Play, YouTube, and/or Widevine packaging, provisioning, and/or license server, which serves as a secure access control intermediary between the CP/CWIP and a Google Play and/or YouTube client).

A. Encryption

1. Licensee (or its Approved Secure Streaming Provider) will always stream Included Programs to Customers in encrypted form.
2. The DRM Technology will only decrypt streamed Included Programs temporarily for the purpose of decoding and rendering such content.
3. Included Programs will be encrypted using standard, nonproprietary, time-tested cryptographic primitives and algorithms and offer effective security equivalent to or better than the encryption standard AES 128.
4. Encryption will be applied to a reasonable portion of audiovisual data given performance weighed against security risk.
5. Each content file containing an Included Program will be encrypted at least once with a cryptographic key which is unique within a large number set.
6. Passwords, cryptographic keys or any other information that is critical to the cryptographic strength of the DRM Technology will never be transmitted or stored outside of Licensee data centers in non-obfuscated form.
7. Playback Licenses, revocation certificates, and security-critical data will be cryptographically protected against tampering, forging, and spoofing.

B. Authentication, Playback and Storage

1. A valid Playback License (containing the unique cryptographic key(s) and other information necessary to decrypt a file of Included Program content and the set of usage rules associated with such content) will be required in order to decrypt and play a specific *instance* of Included Program content.
2. Each Playback License will be keyed to work only on a specific Customer's authorised device or client and will be designed to be incapable of being transferred between unauthorised devices or clients.
3. In the event that the DRM Technology includes client side software, each installation of the DRM Technology client software on an end user device will be individualized to such device and thus uniquely identifiable. As a result, if such software is copied or transferred to another device, the content will be designed to not play on the subsequent device without such subsequent device being authorized by a valid Playback License. Although the current industry standard is to individualize DRM software to devices, Licensee may elect to individualize its DRM Technology client software to a different concept (e.g., by browser, key card, Customer) as the industry standard evolves.
4. The DRM Technology will be upgradeable, allow for backward compatibility for a period of time (where the length of such period of time is determined by Licensee in its sole discretion) if desired, and allow for integration of new rules and business models.

Schedule C: Content Protection Requirements and Obligations, GOOGLE/YOUTUBE VOD LICENSE (Mar. 17, 2011).

286. On information and belief, the intermediary (e.g., the Google Play, YouTube, and/or Widevine packaging, provisioning, and/or license server) has an automated processor (e.g., a Google server automated processor) configured to communicate with a network (e.g., a client-facing VPN and/or the Internet) using network encrypted information, wherein the network encrypted information is in an encrypted form negotiated between a network endpoint (e.g., a Google Play and/or YouTube client device and/or a Google Services client proxy device) and the intermediary (e.g., the Google Play, YouTube, and/or Widevine packaging, provisioning, and/or license server).

287. On information and belief, for respective information (e.g., respective Google Play and/or YouTube license, content, and/or cryptographic key information), the automated processor (e.g., the Google server automated processor) transmits between the server encrypted information (e.g., the information encrypted in a format negotiated between Google Play, YouTube, and/or Widevine packaging, provisioning, and/or license server and a respective CP and/or CWIP server) and the network encrypted information (e.g., the information encrypted in a form negotiated between the Google Play, YouTube, and/or Widevine packaging, provisioning, and/or license server and a Google Play and/or YouTube client device and/or Google Services client proxy device), substantially without an intermediate representation of the information in a decrypted form.

288. On information and belief, the secure distributed access control system (e.g., the distributed Google server system for securely controlling access to sensitive Google Play and/or YouTube license, content, and/or cryptographic key information) comprises at least one audit database configured to log usage of at least one of the plurality of independently operating servers and the activity of the intermediary (e.g., at least one audit database configured to securely log access to and/or usage of (and/or attempted access to and/or usage of) protected Google Play and/or YouTube license, content, and/or cryptographic key information and/or access to and/or usage of (and/or attempted access to and/or usage of) privileged DRM encryption/decryption routines, methods, libraries, and/or frameworks required to effect such access to and/or usage of the protected Google Play and/or YouTube license, content, and/or cryptographic key information).

```
public List<byte[]> getSecureStops ()
```

Added in API level 18

A means of enforcing limits on the number of concurrent streams per subscriber across devices is provided via SecureStop. This is achieved by securely monitoring the lifetime of sessions.

Information from the server related to the current playback session is written to persistent storage on the device when each MediaCrypto object is created.

In the normal case, playback will be completed, the session destroyed and the Secure Stops will be queried. The app queries secure stops and forwards the secure stop message to the server which verifies the signature and notifies the server side database that the session destruction has been confirmed. The persisted record on the client is only removed after positive confirmation that the server received the message using `releaseSecureStops()`.

MediaDRM, ANDROID DEVELOPERS REFERENCE, available at: <http://developer.android.com/reference/android/media/MediaDrm.html> (retrieved Oct. 13, 2015).

289. On information and belief, Google (*e.g.*, through operation of the Google ‘237 Products) performs a method of secure distributed information access control.

III. OTHER FEATURES OF THE GOOGLE SECURITY SYSTEM

In addition to the DRM Technology, Licensee will also apply the following security measures as part of its overall security system designed to protect the Included Programs from unauthorized access during the Term (the “Licensee Security System”):

A. Time-Limited URLs (for Streaming only)

Licensee and/or its designated CDN will use commercially reasonable efforts to implement time and usage limited URLs. The URL address from which Included Program streams can be obtained will be valid for a limited period of time, authorized for a single Customer only, and will contain a statistically unique and unpredictable element or a cryptographic signature to verify authenticity of the URL.

B. Anti-Piracy Cooperation between parties

Without limiting any other provision of the Agreement, the parties acknowledge and agree that it is in their mutual interest to take affirmative measures, acting in good faith cooperation, to combat the unauthorized distribution of copyrighted content. Hence, the parties have entered into the Content Identification and Management Agreement (“CIMA”) or Content Hosting Services Agreement (“CHSA”), as applicable, as an important initiative to combat the unauthorized distribution of copyrighted content.

D. Customer Account Authorization.

Content Delivery. Content shall only be delivered from a network service to a single Customer with an account using verified credentials. Customer Account credentials must be transmitted securely to ensure privacy and protection against attacks.

Services requiring user authentication:

The credentials shall consist of at least a User ID and password of sufficient length to prevent brute force attacks.

Licensee shall take steps to prevent Customers from sharing account access. In order to prevent unwanted sharing of such access, account credentials may provide access to any of the following (by way of example):

- purchasing capability (*e.g.* access to the Customer’s active credit card or other financially sensitive information)
- personal information
- administrator rights over the Customer’s account (*e.g.* including the ability to change passwords, register/de-register devices)

Schedule C: Content Protection Requirements and Obligations, GOOGLE/YOUTUBE VOD LICENSE (Mar. 17, 2011).

290. For example, on information and belief, Google (*e.g.*, through operation of the Google ‘237 Products) communicates with a plurality of independently operating servers (*e.g.*, a plurality of independently operating CP and/or CWIP origin servers) using server encrypted information, wherein the server encrypted information is in an encrypted form negotiated between a respective server (*e.g.*, a respective CP and/or CWIP origin server) and an

intermediary (e.g., a Google, YouTube, and/or Widevine packaging, provisioning, and/or license server intermediary).

291. For example, on information and belief, Google (e.g., through operation of the Google '237 Products) communicates with a network (e.g., a client-facing VPN and/or the Internet) using network encrypted information, wherein the network encrypted information is in a form negotiated between a network endpoint (e.g., a respective CP and/or CWIP origin server) and an intermediary (e.g., a Google, YouTube, and/or Widevine packaging, provisioning, and/or license server intermediary).

292. For example, on information and belief, Google (e.g., through operation of the Google '237 Products) transmits between the server encrypted information and the network encrypted information with an automated processor (e.g., a Google server automated processor), substantially without an intermediate representation of the information in a decrypted form.

293. For example, on information and belief, Google (e.g., through operation of the Google '237 Products) logs in a database (e.g., a Google/Widevine audit/metering database) usage of at least one of the plurality of independently operating servers (e.g., at least one of the plurality of independently operating CP and/or CWIP origin servers) and the activity of the intermediary (e.g., the Google, YouTube, and/or Widevine packaging, provisioning, and/or license server intermediary).

294. On information and belief, the Google '237 Products are made, sold, and/or offered for sale by and/or on behalf of Google to entities (e.g., businesses, schools, and other organizations) and individuals throughout the United States.

295. On information and belief, the Google '237 Products are made, sold, and offered for sale by and/or on behalf of Google to entities (e.g., businesses, schools, and other organizations) and individuals located in the Eastern District of Texas.

296. On information and belief, the Google '237 Products are used by Google (e.g., by and/or on behalf of Google employees) throughout the United States.

297. On information and belief, the Google ‘237 Products are used by Google (e.g., by and/or on behalf of Google employees) within the Eastern District of Texas.

298. On information and belief, Google has directly infringed and continues to directly infringe the ‘237 patent by, among other things, making, using, offering for sale, and/or selling secure three-party communications products and/or services, including but not limited to the Google ‘237 Products, which comprise infringing secure encryption and access control technologies. Such products and/or services include, by way of example and without limitation, the Google Play platform, the YouTube platform, and the Widevine DRM platform.

299. By making, using, offering for sale, and/or selling infringing secure encryption and access control products and services, including but not limited to the Google ‘237 Products, Google has injured St. Luke and is liable to St. Luke for directly infringing one or more claims of the ‘237 patent, including at least claims 1, 18 and 19, pursuant to 35 U.S.C. § 271(a).

300. On information and belief, Google also indirectly infringes the ‘237 patent by actively inducing infringement under 35 USC § 271(b).

301. For example, on information and belief, Google has had knowledge of the ‘237 patent since at least November 2014, when the USPTO expressly cited the ‘237 patent as prior art to a then-pending Google Inc. patent application, U.S. Patent Application No. 13/463,668 (now U.S. Patent No. 8,978,093).

* A		Document Number		Date	Name	Classification
		Country Code-Number-Kind Code	MM-YYYY			
	A	US-2005/0160161 A1	07-2005	Barrett et al.	709/223	
	B	US-2007/0016597 A1	01-2007	Beadles et al.	707/100	
	C	US-2010/0268771 A1	10-2010	Kulakowski et al.	709/203	
	D	US-2010/0322255 A1	12-2010	Hao et al.	370/398	
	E	US-2011/0103586 A1	05-2011	Nobre, Tacito Pereira	380/270	
	F	US-2011/0119729 A1	05-2011	Bergeson et al.	726/1	
	G	US-2011/0235595 A1	09-2011	Mehta et al.	370/329	
	H	US-2011/0252230 A1	10-2011	Segre et al.	713/155	
	I	US-2011/0314145 A1	12-2011	Raleigh et al.	709/224	
	J	US-2012/0069727 A1	04-2012	Raleigh et al.	709/224	
	K	US-2012/0180135 A1	07-2012	Hodges et al.	726/26	
	L	US-2012/0215911 A1	08-2012	Raleigh et al.	709/224	
	M	US-8,316,237 B1	11-2012	Felsher et al.	713/171	

302. Alternatively, on information and belief, at least since service of this Complaint or shortly thereafter, Google has known of the '237 patent and has known about infringement of the '237 patent by Google itself and by third-party Google customers, end-users, developers, and/or integrators/partners of the Google '237 Products.

303. On information and belief, beginning in November 2014 (and in no event later than the date of service of this Complaint), Google has intentionally performed acts that induce infringement of the '237 patent by third parties (e.g., Google '237 Product customers, end-users, developers, and/or integrators/partners), knowing that these acts would induce third-party infringement of the '237 patent and/or with willful blindness to this fact.

304. For example, on information and belief, Google provides products and services (e.g., the Google '237 Products) capable of infringing one or more claims of the '237 patent, including at least claims 1, 18, and 19.

305. On information and belief, Google configures these products and services (e.g., the Google '237 Products) to infringe at least one claim of the '237 patent in normal operation by Google customers, end-users, developers, and/or integrators/partners.

306. For example, on information and belief, Google instructs and directs customers, end-users, developers, and/or integrators/partners to make and/or use the Google '237 Products in an infringing manner and/or configuration (e.g., through creation and dissemination of Google '237 Product documentation, training materials, SDKs, client libraries, and API products and services that not only facilitate, but effectively mandate, third-party infringement of the '237 patent by Google customers, end-users, developers, and/or integrators/partners).

307. Accordingly, Google has actively induced and continues to actively induce infringement of the '237 patent by Google '237 Product customers, end-users, developers, and/or integrators/partners.

308. To the extent applicable, the requirements of 35 U.S.C. § 287(a) have been met with respect to the '237 patent.

309. As a result of Google's infringement of the '237 patent, St. Luke has suffered monetary damages, and seeks recovery in an amount adequate to compensate for Google's infringement, but in no event less than a reasonable royalty for the use made of the '237 patent inventions by Google, together with interest and costs as fixed by the Court.

COUNT II
INFRINGEMENT OF U.S. PATENT NO. 7,181,017

310. St. Luke references and incorporates by reference the preceding paragraphs of this Complaint as if fully set forth herein.

311. Google designs, makes, sells, offers to sell, imports, and/or uses in the United States products and/or services for secure three-party communications.

312. Google designs, makes, sells, offers to sell, imports, and/or uses in the United States the Google Play Store platform, including but not limited to Google Play and/or Widevine packaging and license/key management server hardware and/or programming; Google Play web and mobile applications and services (e.g., Google Play Movies & TV applications and services for Android, iOS, and Chrome); the Google Play website; and Google Play/Widevine secure client hardware and/or programming on end-user devices (e.g., Android smartphones, tablets, and streaming media devices; Chrome OS computers and streaming media devices; Widevine CDM and/or plugin provisioned computers and mobile devices with Chrome Browser software and services) (collectively, "Google Play").

313. Google designs, makes, sells, offers to sell, imports, and/or uses in the United States the YouTube platform, including but not limited to YouTube and/or Widevine packaging and license/key management server hardware and/or programming; YouTube web and mobile applications and services (e.g., YouTube applications and services for Android, iOS, and Chrome); the YouTube website; and YouTube/Widevine secure client hardware and/or programming on end-user devices (e.g., Android smartphones, tablets, and streaming media devices; Chrome OS computers and streaming media devices; Widevine CDM and/or plugin

provisioned computers and mobile devices with Chrome Browser software and services) (collectively, “YouTube”).

314. Google designs, makes, sells, offers to sell, imports, and/or uses in the United States the Widevine DRM platform, including but not limited to Widevine packaging, provisioning, and license/key management server hardware and/or programming; Widevine secure client libraries and programming embedded in and/or integrated with web and mobile applications and services (e.g., YouTube and Google Play Movies & TV applications and services for Android, iOS, and Chrome); and Widevine secure client hardware and/or programming on end-user devices (e.g., Android smartphones, tablets, and streaming media devices; Chrome OS computers and streaming media devices; Widevine CDM and/or plugin provisioned computers and mobile devices with Chrome Browser software and services) (collectively, “Widevine”).

315. Google designs, makes, sells, offers to sell, imports, and/or uses in the United States Google Play, YouTube, and Widevine (collectively, the “Google ‘017 Products”).

316. On information and belief, Google performs, through operation of the Google ‘017 Products, a method for processing information.

317. On information and belief, Google receives, through operation of the Google ‘017 Products, information to be processed. For example, an intermediary Widevine physical and/or virtual server appliance (“the Widevine transcryption intermediary”) receives Google Play and/or YouTube license, content, and/or cryptographic key information (“the VOD transcryption information”) communicated from a server appliance controlled by, dedicated to, and/or associated with a Google Content Partner and/or Certified Widevine Integration Partner (“the CP appliance”). The VOD transcryption information is to be cryptographically processed by the Widevine transcryption intermediary (as described below, for example) before being communicated to a specific Widevine-provisioned Google Play and/or YouTube client (“the Google VOD client”).

318. On information and belief, Google defines, through operation of the Google ‘017 Products, a cryptographic comprehension function for the information, adapted for making at least a portion of the information incomprehensible. For example, the Widevine transcription intermediary defines, through negotiation with the CP appliance, a “back end” cryptographic comprehension function (e.g., cryptographic protocol, key, and initialization information) for the VOD transcription information. The back-end cryptographic comprehension function is adapted for making at least a portion of the VOD transcription information incomprehensible.

319. On information and belief, Google receives, through operation of the Google ‘017 Products, asymmetric cryptographic key information, comprising at least asymmetric encryption key information and asymmetric decryption key information. For example, the Widevine transcription intermediary receives (from the Google VOD client and/or a separate keystore appliance such as a physical or virtual HSM) asymmetric cryptographic key information associated with a specific Widevine-provisioned Google VOD client. The asymmetric key information comprises a globally unique RSA public-private asymmetric key pair for a single Widevine-provisioned Google VOD client, and includes at least asymmetric encryption key information (e.g., a public key portion of the Widevine-provisioned client-specific RSA keypair) and asymmetric decryption key information (e.g., a private key portion of the Widevine-provisioned client-specific RSA keypair).

320. On information and belief, Google negotiates, through operation of the Google ‘017 Products, a new cryptographic comprehension function between two parties to a communication using an intermediary. For example, the Widevine transcription intermediary negotiates between a sending party (e.g., the CP appliance) and a receiving party (e.g., the Google VOD client) a new, “front end” cryptographic comprehension function (e.g., cryptographic protocol, key, and initialization information) for the VOD transcription information.

321. On information and belief, Google processes, through operation of the Google ‘017 Products, the information to invert the cryptographic comprehension function and impose

the new cryptographic comprehension function in an integral process, in dependence on at least the symmetric cryptographic key information, without providing the intermediary with sufficient asymmetric cryptographic key information to decrypt the processed information. For example, the Widevine transcription server processes the VOD transcription information to invert the back end cryptographic comprehension function and impose the new, front end cryptographic comprehension function in an integral process, in dependence on at least the client-specific Widevine-provisioned RSA keypair. For example, the Widevine transcription appliance uses the public key portion of the client-specific Widevine-provisioned RSA keypair as an encryption input to generate processed (transcripted) VOD transcription information for external communication to the Google VOD client. The cleartext private key portion of the client-specific Widevine-provisioned RSA keypair is necessary to decrypt the processed (transcripted) VOD transcription information, but the Widevine intermediary has no unencrypted access to the private key portion of the client-specific Widevine-provisioned RSA keypair because the cryptographic wrapper on the private key portion is uniquely and persistently tied to a non-exportable root of trust on the Google VOD client. As a result, the Widevine transcription intermediary is not provided with sufficient asymmetric cryptographic key information to decrypt the processed information (e.g., the transcripted VOD transcription information).

322. On information and belief, Google outputs, through operation of the Google '017 Products, the processed information. For example, the Google transcription appliance outputs the transcripted Google transcription information (now obscured with the new cryptographic comprehension function) for transmission to the Widevine-provisioned Google Play and/or YouTube client.

323. On information and belief, the ability of the asymmetric decryption key information (e.g., the private key portion of the Widevine-provisioned client-specific asymmetric key pair) to decrypt the processed information (e.g., the transcripted Google Play, YouTube, and/or Widevine license, content, and/or cryptographic key information) changes dynamically.

324. On information and belief, the Google '017 Products are made, sold, and/or offered for sale by and/or on behalf of Google to entities (e.g., businesses, schools, and other organizations) and individuals throughout the United States.

325. On information and belief, the Google '017 Products are made, sold, and offered for sale by and/or on behalf of Google to entities (e.g., businesses, schools, and other organizations) and individuals located in the Eastern District of Texas.

326. On information and belief, the Google '017 Products are used by Google (e.g., by and/or on behalf of Google employees) throughout the United States.

327. On information and belief, the Google '017 Products are used by Google (e.g., by and/or on behalf of Google employees) within the Eastern District of Texas.

328. On information and belief, Google has directly infringed and continues to directly infringe the '017 patent by, among other things, directly performing the method of claim 1 of the '017 patent through operation of at least the Google '017 Products.

329. By making, using, offering for sale, and/or selling infringing secure encryption and access control products and services, including but not limited to the Google '017 Products, Google has injured St. Luke and is liable to St. Luke for directly infringing one or more claims of the '017 patent, including at least claim 1, pursuant to 35 U.S.C. § 271(a).

330. On information and belief, Google also indirectly infringes the '017 patent by actively inducing infringement under 35 USC § 271(b).

331. On information and belief, at least since service of this Complaint or shortly thereafter, Google has known of the '017 patent and has known about infringement of the '017 patent by Google itself and by third-party Google customers, end-users, developers, and/or integrators/partners of the Google '017 Products.

332. On information and belief, beginning no later than the date of service of this Complaint, Google has intentionally performed acts that induce infringement of the '017 patent by third parties (e.g., Google '017 Product customers, end-users, developers, and/or

integrators/partners), knowing that these acts would induce third-party infringement of the '017 patent and/or with willful blindness to this fact.

333. For example, on information and belief, Google provides products and services (e.g., the Google '017 Products) capable of infringing one or more claims of the '017 patent, including at least claim 1.

334. For example, on information and belief, Google configures these products and services (e.g., the Google '017 Products) to infringe at least one claim of the '017 patent in normal operation by Google customers, end-users, developers, and/or integrators/partners.

335. For example, on information and belief, Google instructs and directs customers, end-users, developers, and/or integrators/partners to make and/or use the Google '017 Products in an infringing manner and/or configuration (e.g., through creation and dissemination of Google '017 Product documentation, training materials, SDKs, client libraries, and API products and services that not only facilitate, but effectively mandate, third-party infringement of the '017 patent by Google customers, end-users, developers, and/or integrators/partners).

336. Accordingly, Google has actively induced and continues to actively induce infringement of the '017 patent by Google '017 Product customers, end-users, developers, and/or integrators/partners.

337. To the extent applicable, the requirements of 35 U.S.C. § 287(a) have been met with respect to the '017 patent. Among other things, the '017 patent includes only method claims.

338. As a result of Google's infringement of the '017 patent, St. Luke has suffered monetary damages, and seeks recovery in an amount adequate to compensate for Google's infringement, but in no event less than a reasonable royalty for the use made of the '017 patent inventions by Google, together with interest and costs as fixed by the Court.

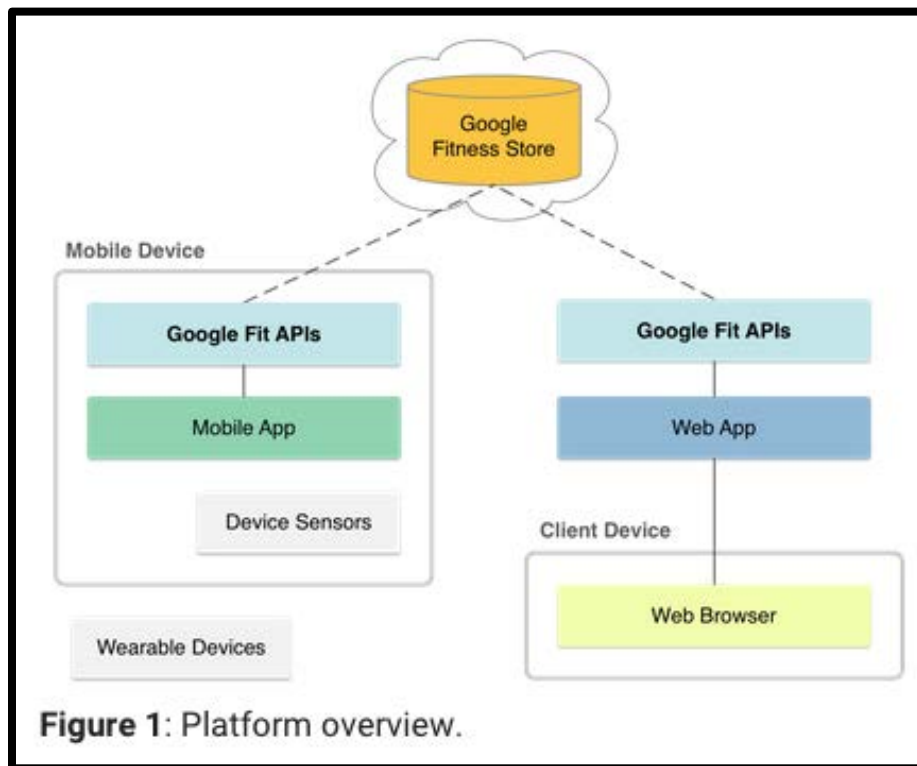
COUNT III
INFRINGEMENT OF U.S. PATENT NO. 7,805,377

339. St. Luke references and incorporates by reference the preceding paragraphs of this Complaint as if fully set forth herein.

340. Google designs, makes, sells, offers to sell, imports, and/or uses in the United States products and/or services for managing access to protected data.

341. Google makes, sells, offers to sell, imports, and/or uses the Google Fit services platform (“Google Fit” or “Google ‘377 Products”).

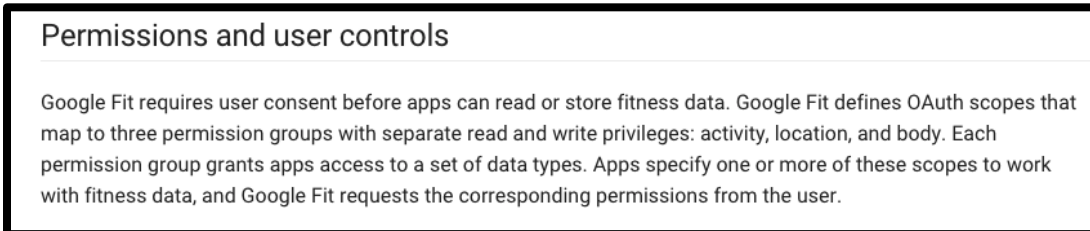
342. On information and belief, Google Fit comprises a system adapted to control access to a patient medical record (e.g., a Google Fit user’s health dataset) hosted by at least one medical record repository (e.g., the Google Fitness Store).



Platform Overview: Components, GOOGLE DEVELOPERS: GOOGLE FIT, <https://developers.google.com/fit/overview> (last accessed Oct. 12, 2015).

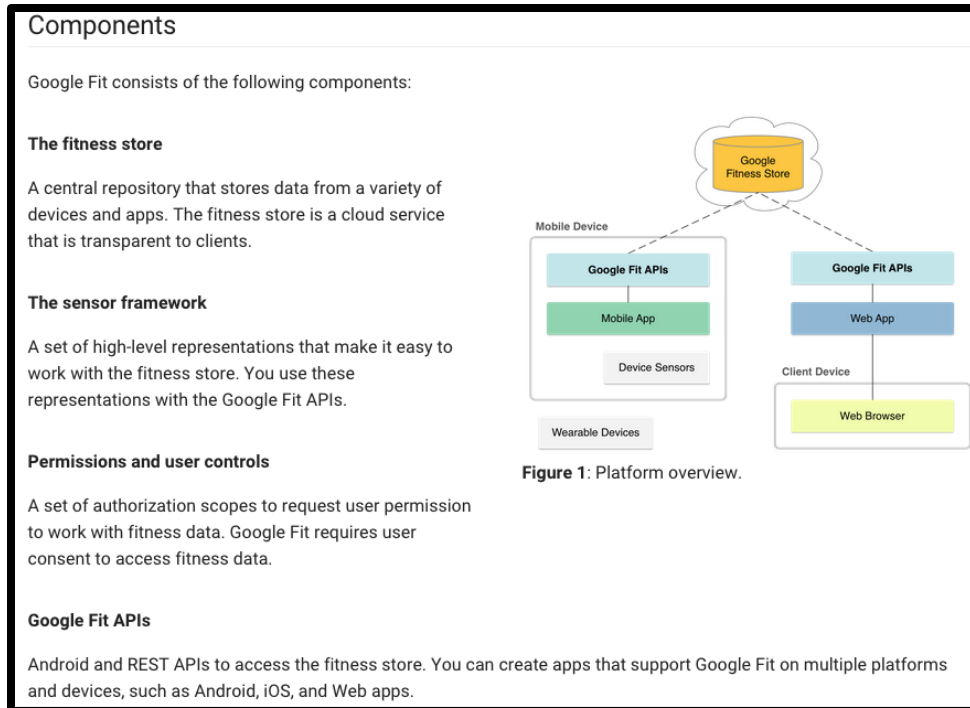
343. On information and belief, the patient medical record (e.g., the Google Fit user’s health dataset) comprises a plurality of record portions (e.g., a plurality of access types and app-specific activity, body, location, and/or nutrition record portions), each record portion being

associated with different patient-controlled access control criteria (e.g., different Google Fit user-controlled authorization scopes required for the respective access types and app-specific activity, body, location, and/or nutrition record portions).




Platform Overview: Permissions and user controls, GOOGLE DEVELOPERS: GOOGLE FIT, <https://developers.google.com/fit/overview> (last accessed Oct. 12, 2015).

344. On information and belief, Google Fit comprises an automated processor (e.g., a Google server automated processor); a database adapted to store information for authenticating requestors (e.g., a Google Fit API authentication database adapted to store digital certificate SHA1 fingerprint information for Google Fit-enabled apps); a database adapted to store information for determining patient-controlled access control criteria for respective record portions of a patient medical record (e.g., a Google Fit API OAuth 2.0 authorization scope database); and a computer network interface (e.g., an Internet and/or VPN wired and/or wireless computer network interface).



Platform Overview: Components, GOOGLE DEVELOPERS: GOOGLE FIT, <https://developers.google.com/fit/overview> (last accessed Oct. 12, 2015).

345. On information and belief, the Google Fit automated processor (e.g., the Google server automated processor) is controlled by instructions stored on a computer readable storage medium (e.g., Google Fit computer code stored in Google server-accessible non-volatile memory) to control access to a patient medical record (e.g., a Google Fit user’s health dataset) comprising a plurality of record portions (e.g., a plurality of access types and app-specific activity, body, and/or location record portions), each record portion being associated with different patient-controlled access control criteria (e.g., different Google Fit user-controlled authorization scopes required for the requested access type and app-specific activity, body, location, and/or nutrition record portions).



Platform Overview: The fitness store, GOOGLE DEVELOPERS: GOOGLE FIT, <https://developers.google.com/fit/overview> (last accessed Oct. 12, 2015).

Here's the OAuth 2.0 scope information for the Fitness API:

Scope	Data Types
https://www.googleapis.com/auth/fitness.activity.read	com.google.activity.sample com.google.activity.segment com.google.activity.summary com.google.calories.consumed com.google.calories.expended
https://www.googleapis.com/auth/fitness.activity.write	com.google.cycling.pedaling.cadence com.google.power.sample com.google.step_count.cadence com.google.step_count.delta
https://www.googleapis.com/auth/fitness.body.read	com.google.heart_rate.bpm com.google.heart_rate.summary com.google.height
https://www.googleapis.com/auth/fitness.body.write	com.google.weight com.google.weight.summary
https://www.googleapis.com/auth/fitness.location.read	com.google.cycling.wheel_revolution.cumulative com.google.cycling.wheel.revolutions com.google.distance.delta com.google.location.sample
https://www.googleapis.com/auth/fitness.location.write	com.google.location.bounding_box com.google.speed com.google.speed.summary

To request access using OAuth 2.0, your application needs the scope information, as well as information that Google supplies when you register your application (such as the client ID and the client secret).

Fit REST API: Platform Basics - Authorization, GOOGLE DEVELOPERS: GOOGLE FIT, <https://developers.google.com/fit/rest/v1/authorization> (last accessed Oct. 12, 2015).

346. On information and belief, the Google Fit automated processor (e.g., the Google server automated processor) is controlled by instructions stored on a computer readable storage medium (e.g., Google Fit computer code stored in Google server-accessible non-volatile memory) to receive a request (e.g., a specially formatted HTTP request) for a medical record (e.g., a Google Fit health record) from a requestor (e.g., a Google Fit-enabled web or mobile application and/or service).

Fitness Data Types ☆☆☆☆☆

Google Fit provides a set of fitness data types under the `com.google` namespace. You can also define custom data types in your own namespace.

Data types define the format of the values inside data points. A data point can represent:

- An instantaneous reading or observation
- An aggregate with statistics over a time interval

Google Fit defines data types for instantaneous observations and data types for aggregate data. Data points consist of values for the fields of a data type and timestamp information. Points that represent instantaneous observations include a timestamp, and points of an aggregate data type also include the start time for the interval.

Google Fit enables you to define new data types in your app and to share your new data types with other apps.

Fit API for Android: Platform Basics – Fitness Data Types, GOOGLE DEVELOPERS: GOOGLE FIT, <https://developers.google.com/fit/android/data-types> (last accessed Oct. 12, 2015).

```
Get a particular data source
-----
This example demonstrates how to retrieve the data source from the previous example.

HTTP method
GET

Request URL
https://www.googleapis.com/fitness/v1/users/me/dataSources/derived:com.google.step_count
.delta:1234567890:Example%20Manufacturer:ExampleTablet:1000001
```

Fit REST API: Store and Access Data – Manage Data Sources, GOOGLE DEVELOPERS: GOOGLE FIT, <https://developers.google.com/fit/rest/v1/data-sources> (last accessed Oct. 12, 2015).

```
Get a dataset
-----
This example demonstrates how to get the contents of a dataset.

HTTP method
GET

Request URL
https://www.googleapis.com/fitness/v1/users/me/dataSources/
derived:com.google.step_count.delta:1234567890:Example%20Manufacturer:ExampleTablet:10000
01
/datasets/1397513334728708316-1397515179728708316
```

Fit REST API: Store and Access Data – Work with Datasets, GOOGLE DEVELOPERS: GOOGLE FIT, <https://developers.google.com/fit/rest/v1/datasets> (last accessed Oct. 12, 2015).

347. On information and belief, the request comprises a medical record identifier (e.g., a Google Fit dataset identifier comprising and/or derived from data source and/or data point identification information); a requestor identifier (e.g., a GUID comprising, constituting, and/or derived from a Google Developers Console API client ID for the requesting Google Fit-enabled web or mobile application and/or service); requestor authentication information (e.g., cryptographic authentication information comprising, constituting, and/or derived from a Google Developers Console API client secret and/or developer digital certificate for the Google Fit-enabled web or mobile application and/or service); and patient-provided access control authorization (e.g., cryptographically tokenized OAuth 2.0 access scope information comprising, constituting, and/or derived from Google Fit user-provided consent information).

Authorizing Requests ☆☆☆☆☆

Every request your application sends to the Fitness API must include an authorization token. The token also identifies your application to Google.

About authorization protocols

Your application must use [OAuth 2.0](#) to authorize requests. No other authorization protocols are supported. If your application uses [Google+ Sign-In](#), some aspects of authorization are handled for you.

Authorizing requests with OAuth 2.0

All requests to the Fitness API must be authorized by an authenticated user.

The details of the authorization process, or "flow," for OAuth 2.0 vary somewhat depending on what kind of application you're writing. The following general process applies to all application types:

1. When you create your application, you register it using the [Google Developers Console](#). Google then provides information you'll need later, such as a client ID and a client secret.
2. Activate the Fitness API in the Google Developers Console. (If the API isn't listed in the Developers Console, then skip this step.)
3. When your application needs access to user data, it asks Google for a particular **scope** of access.
4. Google displays a **consent screen** to the user, asking them to authorize your application to request some of their data.
5. If the user approves, then Google gives your application a short-lived **access token**.
6. Your application requests user data, attaching the access token to the request.
7. If Google determines that your request and the token are valid, it returns the requested data.

Fit REST API: Platform Basics - Authorization, GOOGLE DEVELOPERS: GOOGLE FIT, <https://developers.google.com/fit/rest/v1/authorization> (last accessed Oct. 12, 2015).

Authorization on Android

User consent is always required before your app can read or write fitness data. To obtain authorization:

- Register your Android app with a project in the Google Developers Console.
- Specify a scope of access when connecting to the fitness service.

In Google Fit, scopes are strings that determine what kinds of fitness data an app can access and the level of access to this data.

Authorization flow

The authorization flow is the following:

1. Your app requests a connection to the fitness service with one or more scopes of access.
2. Google Fit prompts the user to grant your app the required permissions.
3. If the user consents, your app can then access fitness data of the types defined by the scope.

The specific permissions requested to the user depend on the scopes that you specify when connecting to the service.

Fitness scopes

The scopes for Google Fit on Android are defined as public fields of the `Scopes` class. Their field names start with the `FITNESS_` prefix. Each scope provides access to a set of fitness data types. Some scopes provide read-only access to fitness data, while other scopes provide read and write access to fitness data. The scopes are listed in table 1.

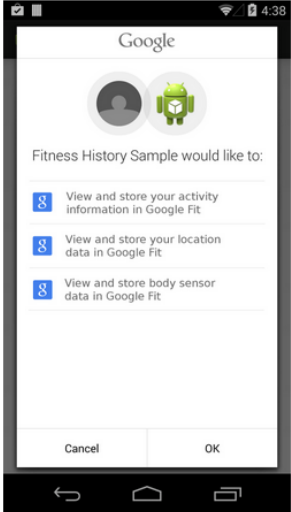


Figure 1: The consent screen.

Fit API for Android: Platform Basics – Authorization, GOOGLE DEVELOPERS: GOOGLE FIT, <https://developers.google.com/fit/android/authorization> (last accessed Oct. 12, 2015).

348. On information and belief, the Google Fit automated processor (e.g., the Google server automated processor) is controlled by instructions stored on a computer readable storage medium (e.g., Google Fit computer code stored in Google server-accessible non-volatile memory) to process the request for the medical record (e.g., the specially formatted HTTP request for the Google Fit health record), to authenticate the requestor (e.g., the requesting Google Fit-enabled web or mobile application and/or service) and determine sufficiency of the patient-provided access control authorization (e.g., cryptographically tokenized OAuth 2.0 access scope information comprising, constituting, and/or derived from Google Fit user-provided consent information) to meet the patient-controlled access control criteria for each respective record portion encompassed by the request (e.g., the Google Fit user-controlled authorization

scope required for the requested access type and app-specific activity, body, location, and/or nutrition record portions).

Permissions and user controls

Google Fit requires user consent before apps can read or store fitness data. Google Fit defines OAuth scopes that map to three permission groups with separate read and write privileges: activity, location, and body. Each permission group grants apps access to a set of data types. Apps specify one or more of these scopes to work with fitness data, and Google Fit requests the corresponding permissions from the user.

Platform Overview: Permissions and user controls, GOOGLE DEVELOPERS: GOOGLE FIT, <https://developers.google.com/fit/overview> (last accessed Oct. 12, 2015).

Permission	Scope	Type of Access	Data Types
Activity	FITNESS_ACTIVITY_READ	Read	com.google.activity.sample com.google.activity.segment com.google.activity.summary (deprecated) com.google.calories.consumed com.google.calories.expended
	FITNESS_ACTIVITY_READ_WRITE	Read and Write	com.google.cycling.pedaling.cadence com.google.power.sample com.google.step_count.cadence com.google.step_count.delta com.google.activity.exercise
Body	FITNESS_BODY_READ	Read	com.google.heart_rate.bpm com.google.heart_rate.summary com.google.height
	FITNESS_BODY_READ_WRITE	Read and Write	com.google.weight com.google.weight.summary
Location	FITNESS_LOCATION_READ	Read	com.google.cycling.wheel_revolution. cumulative com.google.cycling.wheel.revolutions com.google.distance.delta
	FITNESS_LOCATION_READ_WRITE	Read and Write	com.google.location.sample com.google.location.bounding_box com.google.speed com.google.speed.summary
Nutrition	FITNESS_NUTRITION_READ	Read	com.google.nutrition.item com.google.nutrition.summary
	FITNESS_NUTRITION_READ_WRITE	Read and Write	

Table 1: Scopes for Google Fit

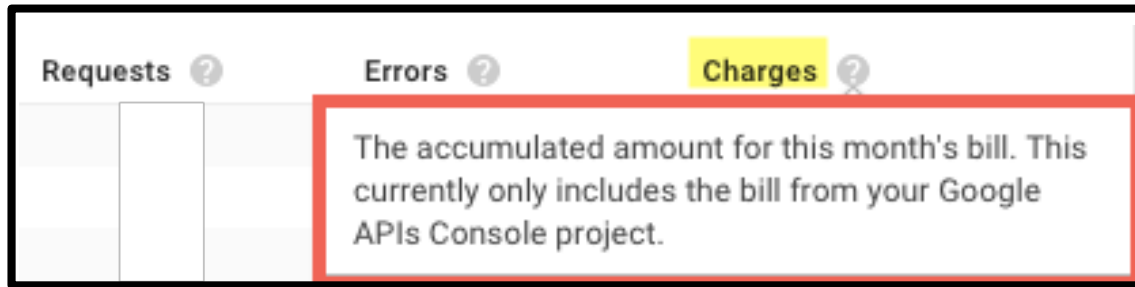
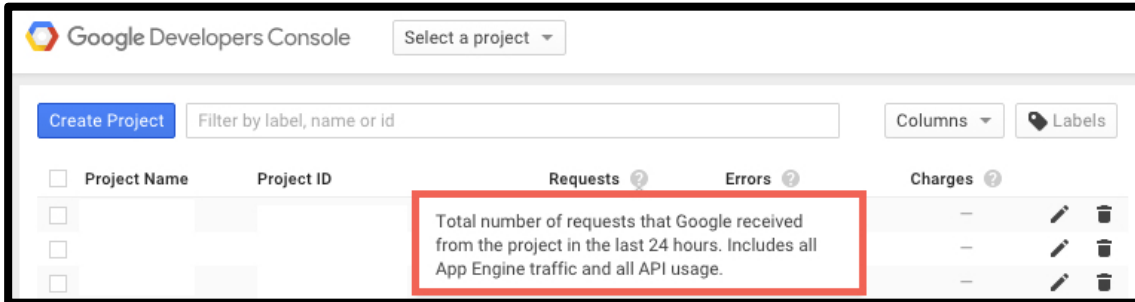
Fit API for Android: Platform Basics – Authorization, GOOGLE DEVELOPERS: GOOGLE FIT, <https://developers.google.com/fit/android/authorization> (last accessed Oct. 12, 2015).

349. On information and belief, the Google Fit automated processor (e.g., the Google server automated processor) is controlled by instructions stored on a computer readable storage medium (e.g., Google Fit computer code stored in Google server-accessible non-volatile memory) to selectively communicate through the computer network interface (e.g., the Internet and/or VPN wired and/or wireless computer network interface) to the at least one medical record repository (e.g., the Google Fitness Store), an identification of each record portion (e.g.,

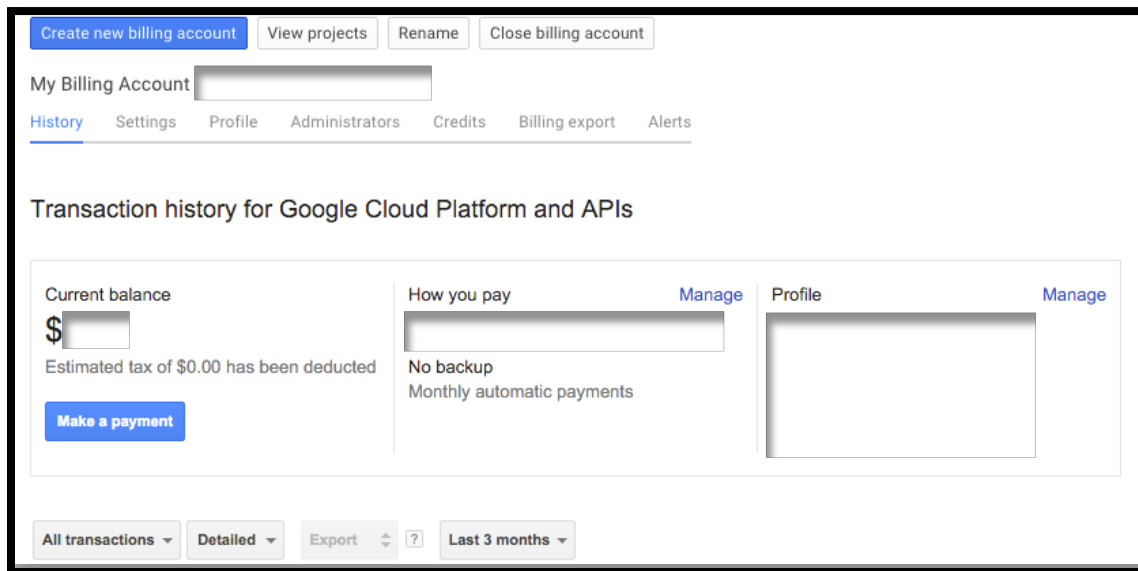
requested access type and app-specific activity, body, location, and/or nutrition information) for which access control criteria (e.g., the Google Fit user-controlled authorization scope required for the requested access type and app-specific activity, body, location, and/or nutrition record portions) are determined to be sufficient for access by the requestor (e.g., the requesting Google Fit-enabled web or mobile application and/or service).

350. On information and belief, the Google Fit automated processor (e.g., the Google server automated processor) is controlled by instructions stored on a computer readable storage medium (e.g., Google Fit computer code stored in Google server-accessible non-volatile memory) to generate an electronic payment authorization associated with the request (e.g., an electronic Google API usage record associated with Google Fit API request), for compensation of at least one of the system and the at least one medical record repository (e.g., to compensate Google for the requesting Google Fit-enabled web or mobile application and/or service's use of Google Fit (and of the broader Google API ecosystem)).

351. For example, on information and belief, Google API requests, including Google Fit API requests, are individually metered to enforce daily, weekly, and monthly API request quotas related to compensating Google for applications' (e.g., Google Fit-enabled web or mobile applications and/or services) use of Google services, including Google Fit. On information and belief, Google requires that developers provide at least one form of automated payment prior to authorizing API use, in order to facilitate automated payment to Google for applications' metered usage of Google API, including Google Fit APIs.



Project, GOOGLE DEVELOPERS CONSOLE, <https://console.developers.google.com/project> (accessed Oct. 3, 2015).



Billing, GOOGLE DEVELOPERS CONSOLE, <https://console.developers.google.com/billing> (accessed Oct. 3, 2015).

352. On information and belief, Google performs (e.g., via Google Fit) a method for controlling access to a medical record or a patient hosted by at least one medical record repository.

353. On information and belief, the method performed by Google (e.g., via Google Fit) comprises receiving, by an intermediary, a request for a medical record comprising a plurality of record portions, each record portion having an associated different patient-controlled access control criteria, from a requestor, said request comprising a medical record identifier, a requestor identifier, and a requestor authentication information.

354. On information and belief, the method performed by Google (e.g., via Google Fit) comprises automatically processing, by an automated processor associated with the intermediary, the request for the medical record to authenticate the requestor.

355. On information and belief, the method performed by Google (e.g., via Google Fit) comprises receiving, by the intermediary, a patient-provided access control authorization associated with a request for a medical record from a requestor.

356. On information and belief, the method performed by Google (e.g., via Google Fit) comprises automatically processing, by the automated processor associated with the intermediary, the request for the medical record, to further determine sufficiency of the patient-provided access control authorization to meet the patient-controlled access control criteria for each respective record portion encompassed by the request.

357. On information and belief, the method performed by Google (e.g., via Google Fit) comprises selectively communicating through an automated communication network, from the intermediary to the at least one medical record repository, an authenticated request for access to the medical record by the requestor and an identification of each record portion for which access control criteria are determined to be sufficient for access by the requestor.

358. On information and belief, the method performed by Google (e.g., via Google Fit) comprises generating an electronic payment message associated with the request, relating to compensation of at least one of the intermediary and a medical record repository.

359. On information and belief, Google Fit is made, sold, and/or offered for sale by and/or on behalf of Google to entities (e.g., businesses, schools, and other organizations) and individuals throughout the United States.

360. On information and belief, Google Fit is made, sold, and offered for sale by and/or on behalf of Google to entities (e.g., businesses, schools, and other organizations) and individuals located in the Eastern District of Texas.

361. On information and belief, Google Fit is used by Google (e.g., by and/or on behalf of Google employees) throughout the United States.

362. On information and belief, Google Fit is used by Google (e.g., by and/or on behalf of Google employees) within the Eastern District of Texas.

363. On information and belief, Google has directly infringed and continues to directly infringe the '377 patent by, among other things, directly performing the method of claim 13 of the patent through operation of at least Google Fit.

364. By making, using, offering for sale, and/or selling infringing products and services for managing access to protected data, including but not limited to Google Fit, Google has injured St. Luke and is liable to St. Luke for directly infringing one or more claims of the '377 patent, including at least claims 7 and 13, pursuant to 35 U.S.C. § 271(a).

365. On information and belief, Google also indirectly infringes the '377 patent by actively inducing infringement under 35 USC § 271(b).

366. On information and belief, at least since service of this Complaint or shortly thereafter, Google has known of the '377 patent and has known about infringement of the '377 patent by Google itself and by third-party Google customers, end-users, developers, and/or integrators/partners of Google Fit.

367. On information and belief, beginning no later than the date of service of this Complaint, Google has intentionally performed acts that induce infringement of the '377 patent by third parties (e.g., Google Fit customers, end-users, developers, and/or integrators/partners), knowing that these acts would induce third-party infringement of the '377 patent and/or with willful blindness to this fact.

368. For example, on information and belief, Google provides products and services (e.g., Google Fit) capable of infringing one or more claims of the '377 patent, including at least claims 7 and 13.

369. For example, on information and belief, Google configures these products and services (e.g., Google Fit) to infringe at least one claim of the '377 patent in normal operation by Google customers, end-users, developers, and/or integrators/partners.

370. For example, on information and belief, Google instructs and directs customers, end-users, developers, and/or integrators/partners to make and/or use Google Fit in an infringing manner and/or configuration (e.g., through creation and dissemination of Google Fit documentation, training materials, SDKs, client libraries, and API products and services that not only facilitate, but effectively mandate, third-party infringement of the '377 patent by Google customers, end-users, developers, and/or integrators/partners).

371. Accordingly, Google has actively induced and continues to actively induce infringement of the '377 patent by Google Fit customers, end-users, developers, and/or integrators/partners.

372. To the extent applicable, the requirements of 35 U.S.C. § 287(a) have been met with respect to the '377 patent.

373. As a result of Google's infringement of the '377 patent, St. Luke has suffered monetary damages, and seeks recovery in an amount adequate to compensate for Google's infringement, but in no event less than a reasonable royalty for the use made of the '377 patent inventions by Google, together with interest and costs as fixed by the Court.

COUNT IV
INFRINGEMENT OF U.S. PATENT NO. 7,587,368

374. St. Luke references and incorporates by reference the preceding paragraphs of this Complaint as if fully set forth herein.

375. Google designs, makes, sells, offers to sell, imports, and/or uses in the United States products and/or services for managing access to protected data.

376. Google designs, makes, sells, offers to sell, imports, and/or uses in the United States the Google Play Store platform, including but not limited to Google Play and/or Widevine packaging and license/key management server hardware and/or programming; Google Play web and mobile applications and services (e.g., Google Play Movies & TV applications and services for Android, iOS, and Chrome); the Google Play website; and Google Play/Widevine secure client hardware and/or programming on end-user devices (e.g., Android smartphones, tablets, and streaming media devices; Chrome OS computers and streaming media devices; Widevine CDM and/or plugin provisioned computers and mobile devices with Chrome Browser software and services) (collectively, “Google Play”).

377. Google designs, makes, sells, offers to sell, imports, and/or uses in the United States the YouTube platform, including but not limited to YouTube and/or Widevine packaging and license/key management server hardware and/or programming; YouTube web and mobile applications and services (e.g., YouTube applications and services for Android, iOS, and Chrome); the YouTube website; and YouTube/Widevine secure client hardware and/or programming on end-user devices (e.g., Android smartphones, tablets, and streaming media devices; Chrome OS computers and streaming media devices; Widevine CDM and/or plugin provisioned computers and mobile devices with Chrome Browser software and services) (collectively, “YouTube”).

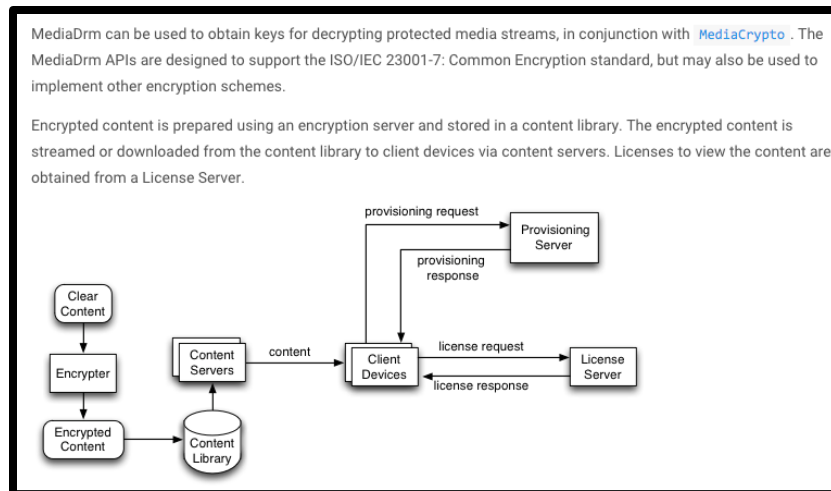
378. Google designs, makes, sells, offers to sell, imports, and/or uses in the United States the Widevine DRM platform, including but not limited to Widevine packaging, provisioning, and license/key management server hardware and/or programming; Widevine secure client libraries and programming embedded in and/or integrated with web and mobile applications and services (e.g., YouTube and Google Play Movies & TV applications and services for Android, iOS, and Chrome); and Widevine secure client hardware and/or programming on end-user devices (e.g., Android smartphones, tablets, and streaming media devices; Chrome OS computers and streaming media devices; Widevine CDM and/or plugin

provisioned computers and mobile devices with Chrome Browser software and services) (collectively, “Widevine”).

379. Google designs, makes, sells, offers to sell, imports, and/or uses in the United States Google Play, YouTube, and Widevine (collectively, the “Google ‘368 Products’”).

380. On information and belief, the Google ‘368 Products comprise a database system (e.g., a Google/Widevine content key/license server database system).

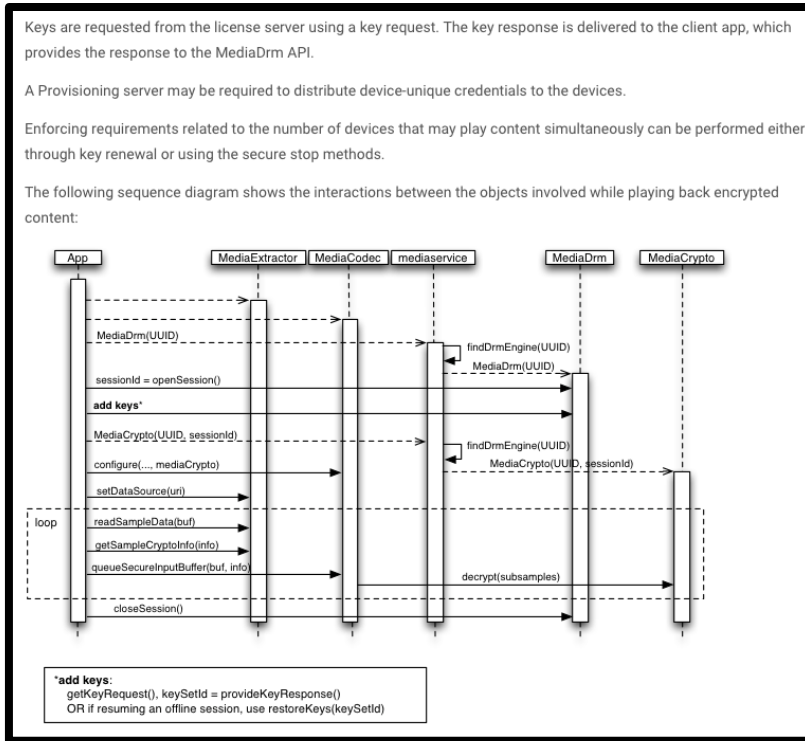
381. On information and belief, the database system (e.g., the Google/Widevine content key/license server database system) comprises a plurality of digital records (e.g., Google/Widevine content key/license digital records associated with respective Google Play and/or YouTube content items), each digital record having an associated set of access rules, stored in a computer memory associated with a server system (e.g., non-volatile flash and/or SSD memory associated with a Google/Widevine content key/license server system).



MediaDRM, ANDROID DEVELOPERS REFERENCE, <http://developer.android.com/reference/android/media/MediaDrm.html> (retrieved Oct. 13, 2015).

382. On information and belief, the database system (e.g., the Google/Widevine content key/license server database system) comprises an interface computer (e.g., a client-facing web server) in communication with a remote computer (e.g., a remote Google Play and/or YouTube client device) for receiving a request (e.g., a specially formatted HTTP client request) for access from the remote computer to access a digital record (e.g., a Google/Widevine content

key/license digital record) stored in the computer memory (e.g., the Google/Widevine content key/license server memory).



The app first constructs `MediaExtractor` and `MediaCodec` objects. It accesses the DRM-scheme-identifying UUID, typically from metadata in the content, and uses this UUID to construct an instance of a `MediaDrm` object that is able to support the DRM scheme required by the content. Crypto schemes are assigned 16 byte UUIDs. The method `isCryptoSchemeSupported(UUID)` can be used to query if a given scheme is supported on the device.

The app calls `openSession()` to generate a `sessionId` that will uniquely identify the session in subsequent interactions. The app next uses the `MediaDrm` object to obtain a key request message and send it to the license server, then provide the server's response to the `MediaDrm` object.

Once the app has a `sessionId`, it can construct a `MediaCrypto` object from the UUID and `sessionId`. The `MediaCrypto` object is registered with the `MediaCodec` in the `configure(MediaFormat, Surface, MediaCrypto, int)` method to enable the codec to decrypt content.

MediaDRM, ANDROID DEVELOPERS REFERENCE, <http://developer.android.com/reference/android/media/MediaDrm.html> (retrieved Oct. 13, 2015).

383. On information and belief, the database system (e.g., the Google/Widevine content key/license server database system) comprises an automated processor (e.g., a Google server automated processor) associated with the server system (e.g., the Google/Widevine content key/license server system).

384. On information and belief, the automated processor (e.g., the Google server automated processor) associated with the server system (e.g., the Google/Widevine content key/license server system) is configured to validate the received request (e.g., specially formatted HTTP client request) to access the digital record (e.g., the Google/Widevine content key/license digital record associated with a respective Google Play and/or YouTube content item) by applying a respective set of access rules for the digital record stored in the computer memory (e.g., the Google/Widevine content key/license server memory).

385. On information and belief, the automated processor (e.g., the Google server automated processor) associated with the server system (e.g., the Google/Widevine content key/license server system) is configured to retrieve a public key (e.g., the public key portion of a Google/Widevine-provisioned client-specific RSA or Elliptical Curve asymmetric key pair) having an associated private key (e.g., the private key portion of a Google/Widevine-provisioned client-specific RSA or Elliptical Curve asymmetric key pair), and to associate a wrapper having a respective session key (e.g., an ephemeral/PFS transport key) with the digital record (e.g., the Google/Widevine content key/license digital record associated with a respective Google Play and/or YouTube content item), after validating the received request.

The app first constructs `MediaExtractor` and `MediaCodec` objects. It accesses the DRM-scheme-identifying UUID, typically from metadata in the content, and uses this UUID to construct an instance of a `MediaDrm` object that is able to support the DRM scheme required by the content. Crypto schemes are assigned 16 byte UUIDs. The method `isCryptoSchemeSupported(UUID)` can be used to query if a given scheme is supported on the device.

The app calls `openSession()` to generate a `sessionId` that will uniquely identify the session in subsequent interactions. The app next uses the `MediaDrm` object to obtain a key request message and send it to the license server, then provide the server's response to the `MediaDrm` object.

Once the app has a `sessionId`, it can construct a `MediaCrypto` object from the UUID and `sessionId`. The `MediaCrypto` object is registered with the `MediaCodec` in the `configure(MediaFormat, Surface, MediaCrypto, int)` method to enable the codec to decrypt content.

When the app has constructed `MediaExtractor`, `MediaCodec` and `MediaCrypto` objects, it proceeds to pull samples from the extractor and queue them into the decoder. For encrypted content, the samples returned from the extractor remain encrypted, they are only decrypted when the samples are delivered to the decoder.

`MediaDrm` methods throw `MediaDrm.MediaDrmStateException` when a method is called on a `MediaDrm` object that has had an unrecoverable failure in the DRM plugin or security hardware. `MediaDrm.MediaDrmStateException` extends `IllegalStateException` with the addition of a developer-readable diagnostic information string associated with the exception.

In the event of a mediaserver process crash or restart while a `MediaDrm` object is active, `MediaDrm` methods may throw `MediaDrmResetException`. To recover, the app must release the `MediaDrm` object, then create and initialize a new one.

As `MediaDrmResetException` and `MediaDrm.MediaDrmStateException` both extend `IllegalStateException`, they should be in an earlier `catch()` block than `IllegalStateException` if handled separately.

MediaDRM, ANDROID DEVELOPERS REFERENCE, available at: <http://developer.android.com/reference/android/media/MediaDrm.html> (retrieved Oct. 13, 2015).

DRM Info

`DrmInfo` is a wrapper class that wraps the protocol for communicating with the DRM server. Server registration, deregistration, license acquisition, or any other server-related transaction can be achieved by processing an instance of `DrmInfo`. The protocol should be described by the plug-in in XML format. Each DRM plug-in would accomplish the transaction by interpreting the protocol. The DRM framework defines an API to retrieve an instance of `DrmInfo` called `acquireDrmInfo()`.

```
DrmInfo* acquireDrmInfo(int uniqueId, const DrmInfoRequest* drmInfoRequest);
```

Retrieves necessary information for registration, deregistration or rights acquisition information. See [DrmInfoRequest](#) for more information.

```
DrmInfoStatus* processDrmInfo(int uniqueId, const DrmInfo* drmInfo);
```

`processDrmInfo()` behaves asynchronously and the results of the transaction can be retrieved either from `OnEventListener` or `OnErrorListener`.

Interfaces: DRM, ANDROID SOURCE, available at: <http://source.android.com/devices/drm.html> (Sep. 14, 2014 version).

386. On information and belief, the session key (e.g., the ephemeral/PFS transport key associated with the wrapper) is distinct from the public key and the private key (e.g., the respective portions of the Google/Widevine-provisioned client-specific RSA or Elliptical Curve asymmetric key pair).


```
public MediaDrm.ProvisionRequest getProvisionRequest ()
```

Added in API level 18

A provision request/response exchange occurs between the app and a provisioning server to retrieve a device certificate. If provisioning is required, the EVENT_PROVISION_REQUIRED event will be sent to the event handler. `getProvisionRequest` is used to obtain the opaque provision request byte array that should be delivered to the provisioning server. The provision request byte array is returned in `ProvisionRequest.data`. The recommended URL to deliver the provision request to is returned in `ProvisionRequest.defaultUrl`.

MediaDrm, ANDROID DEVELOPERS REFERENCE, <http://developer.android.com/reference/android/media/MediaDrm.html> (retrieved Oct. 13, 2015).

387. On information and belief, the automated processor (e.g., the Google server automated processor) associated with the server system (e.g., the Google/Widevine content key/license server system) is configured to encrypt and send the requested digital record (e.g., the requested Google/Widevine content key/license digital record associated with a respective Google Play and/or YouTube content item) that has been validated, using the public key (e.g., the public key portion of a Google/Widevine-provisioned client-specific RSA or Elliptical Curve asymmetric key pair) and the session key (e.g., the ephemeral/PFS transport key associated with the wrapper) to encrypt the digital record, through the interface computer (e.g., the client-facing Google/Widevine web service interface physical and/or virtual appliance).

388. On information and belief, the automated processor (e.g., the Google server automated processor) associated with the server system (e.g., the Google/Widevine content key/license server system) is configured to receive, through the interface computer (e.g., the client-facing Google/Widevine web service interface physical and/or virtual appliance), a logging event (e.g., a usage audit/metering event) from the remote computer (e.g., the remote Google Play and/or YouTube client device) based on an operation of the wrapper and the session key (e.g., a decryption-, license-, playback- and/or other access-related operation on the remote Google Play and/or YouTube client device).

Public Methods	
void	<code>closeSession(byte[] sessionId)</code> Close a session on the MediaDrm object that was previously opened with <code>openSession()</code> .
MediaDrm.CryptoSession	<code>getCryptoSession(byte[] sessionId, String cipherAlgorithm, String macAlgorithm)</code> Obtain a CryptoSession object which can be used to encrypt, decrypt, sign and verify messages or data using the session keys established for the session using methods <code>getKeyRequest(byte[], byte[], String, int, HashMap)</code> and <code>provideKeyResponse(byte[], byte[])</code> using a session key server.
MediaDrm.KeyRequest	<code>getKeyRequest(byte[] scope, byte[] init, String mimeType, int keyType, HashMap<String, String> optionalParameters)</code> A key request/response exchange occurs between the app and a license server to obtain or release keys used to decrypt encrypted content.
byte[]	<code>getPropertyByteArray(String propertyName)</code> Read a DRM engine plugin byte array property value, given the property name string.
String	<code>getPropertyString(String propertyName)</code> Read a DRM engine plugin String property value, given the property name string.
MediaDrm.ProvisionRequest	<code>getProvisionRequest()</code> A provision request/response exchange occurs between the app and a provisioning server to retrieve a device certificate.
byte[]	<code>getSecureStop(byte[] ssid)</code> Access secure stop by secure stop ID.
List<byte[]>	<code>getSecureStops()</code> A means of enforcing limits on the number of concurrent streams per subscriber across devices is provided via SecureStop.

MediaDRM, ANDROID DEVELOPERS REFERENCE, <http://developer.android.com/reference/android/media/MediaDrm.html> (retrieved Oct. 13, 2015).

389. On information and belief, the automated processor (e.g., the Google server automated processor) associated with the server system (e.g., the Google/Widevine content key/license server system) is configured to record the logging event (e.g., the usage audit/metering event) in an access log (e.g., a concurrency monitoring log).

`public List<byte[]> getSecureStops ()` Added in API level 18

A means of enforcing limits on the number of concurrent streams per subscriber across devices is provided via SecureStop. This is achieved by securely monitoring the lifetime of sessions.

Information from the server related to the current playback session is written to persistent storage on the device when each MediaCrypto object is created.

In the normal case, playback will be completed, the session destroyed and the Secure Stops will be queried. The app queries secure stops and forwards the secure stop message to the server which verifies the signature and notifies the server side database that the session destruction has been confirmed. The persisted record on the client is only removed after positive confirmation that the server received the message using `releaseSecureStops()`.

MediaDRM, ANDROID DEVELOPERS REFERENCE, <http://developer.android.com/reference/android/media/MediaDrm.html> (retrieved Oct. 13, 2015).

390. On information and belief, server-implemented, event-based usage/access logging is critical to concurrency limiting—a key secure access control technology Google Play and YouTube rely on to enforce CP requirements for premium streaming media content.

391. On information and belief, the Google '368 Products are made, sold, and/or offered for sale by and/or on behalf of Google to entities (e.g., businesses, schools, and other organizations) and individuals throughout the United States.

392. On information and belief, the Google '368 Products are made, sold, and offered for sale by and/or on behalf of Google to entities (e.g., businesses, schools, and other organizations) and individuals located in the Eastern District of Texas.

393. On information and belief, the Google '368 Products are used by Google (e.g., by and/or on behalf of Google employees) throughout the United States.

394. On information and belief, the Google '368 Products are used by Google (e.g., by and/or on behalf of Google employees) within the Eastern District of Texas.

395. On information and belief, Google has directly infringed and continues to directly infringe the '368 patent by, among other things, directly performing the method of claim 1 of the patent through operation of at least the Google '368 Products.

396. By making, using, offering for sale, and/or selling infringing products and services for managing access to protected data, including but not limited to the Google '368 Products, Google has injured St. Luke and is liable to St. Luke for directly infringing one or more claims of the '368 patent, including at least claims 1, 78, 133, and 140, pursuant to 35 U.S.C. § 271(a).

397. On information and belief, Google also indirectly infringes the '368 patent by actively inducing infringement under 35 USC § 271(b).

398. On information and belief, at least since service of this Complaint or shortly thereafter, Google has known of the '368 patent and has known about infringement of the '368 patent by Google itself and by third-party Google customers, end-users, developers, and/or integrators/partners of the Google '368 Products.

399. On information and belief, beginning no later than the date of service of this Complaint, Google has intentionally performed acts that induce infringement of the '368 patent by third parties (e.g., Google '368 Product customers, end-users, developers, and/or

integrators/partners), knowing that these acts would induce third-party infringement of the '368 patent and/or with willful blindness to this fact.

400. For example, on information and belief, Google provides products and services (e.g., the Google '368 Products) capable of infringing one or more claims of the '368 patent, including at least claims 1, 78, 133, and 140.

401. For example, on information and belief, Google configures these products and services (e.g., the Google '368 Products) to infringe at least one claim of the '368 patent in normal operation by Google customers, end-users, developers, and/or integrators/partners.

402. For example, on information and belief, Google instructs and directs customers, end-users, developers, and/or integrators/partners to make and/or use the Google '368 Products in an infringing manner and/or configuration (e.g., through creation and dissemination of Google '368 Product documentation, training materials, SDKs, client libraries, and API products and services that not only facilitate, but effectively mandate, third-party infringement of the '368 patent by Google customers, end-users, developers, and/or integrators/partners).

403. Accordingly, Google has actively induced and continues to actively induce infringement of the '368 patent by Google '368 Product customers, end-users, developers, and/or integrators/partners.

404. To the extent applicable, the requirements of 35 U.S.C. § 287(a) have been met with respect to the '368 patent.

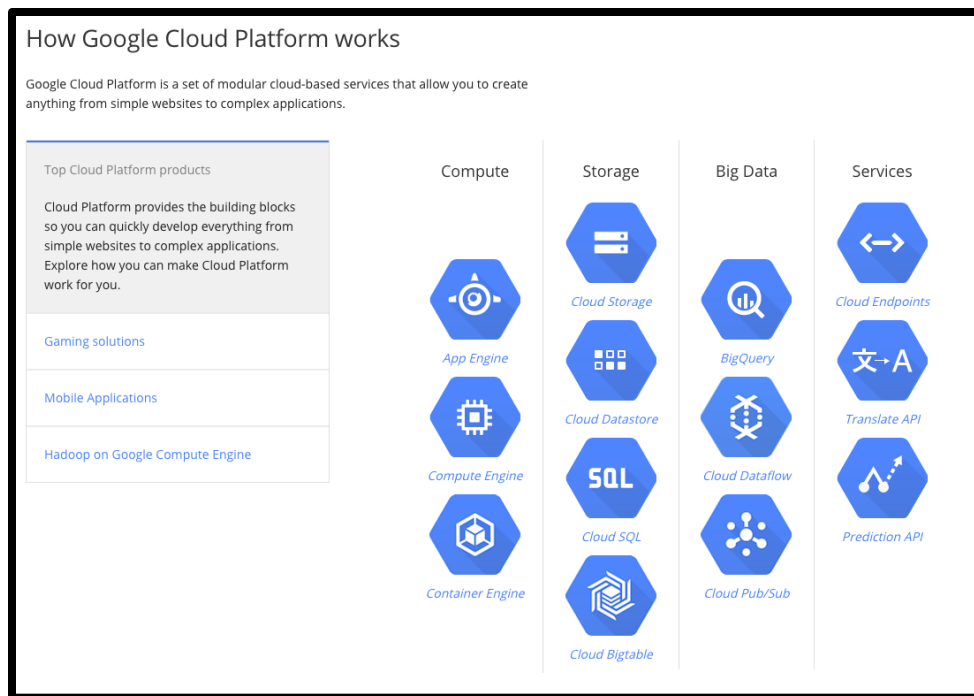
405. As a result of Google's infringement of the '368 patent, St. Luke has suffered monetary damages, and seeks recovery in an amount adequate to compensate for Google's infringement, but in no event less than a reasonable royalty for the use made of the '368 patent inventions by Google, together with interest and costs as fixed by the Court.

COUNT V
INFRINGEMENT OF U.S. PATENT NO. 8,498,941

406. St. Luke references and incorporates by reference the preceding paragraphs of this Complaint as if fully set forth herein.

407. Google designs, makes, sells, offers to sell, imports, and/or uses in the United States products and/or services for managing access to protected data.

408. Google designs, makes, sells, offers to sell, imports, and/or uses in the United States the Google Cloud Platform, including but not limited to Google Cloud Compute products and services (e.g., Google App Engine, Google Compute Engine, Google Container Engine); Google Cloud Storage products and services (e.g., Google Cloud Storage, Google Cloud Datastore, Google Cloud SQL, Google Cloud Bigtable); Google Cloud Big Data products and services (e.g., Google Cloud BigQuery, Google Cloud Dataflow, Google Cloud Pub/Sub); and Google Cloud Services products and services (e.g., Google Cloud Endpoints, Google Translate API, Google Prediction API) (collectively, “Google Cloud”).

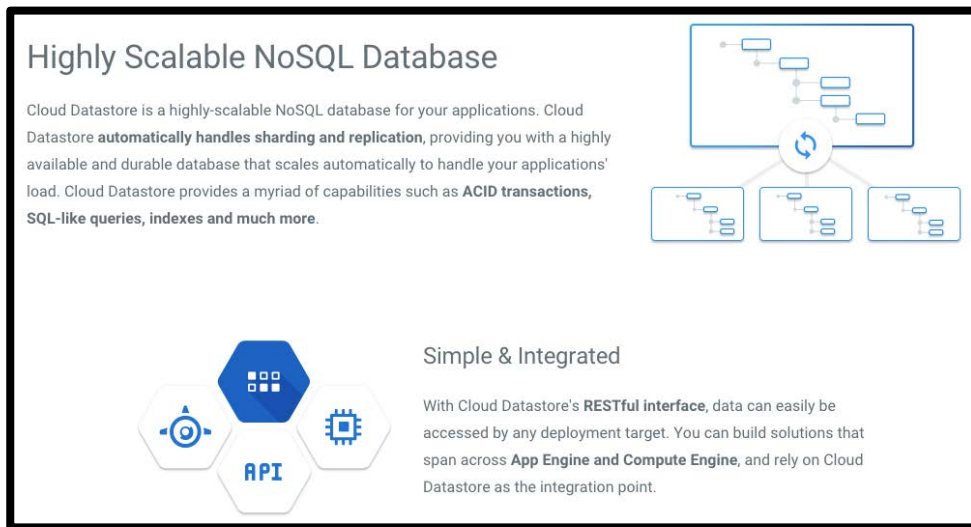


How Google Cloud Platform Works, <https://cloud.google.com> (retrieved Oct. 19, 2015).

409. Google designs, makes, sells, offers to sell, imports, and/or uses in the United States the Google App Engine platform-as-a-service, including but not limited to Google App Engine products and services and server-side and client-side products and services running on

and/or interfacing with Google App Engine products and/or services (collectively, “Google App Engine”).

410. Google designs, makes, sells, offers to sell, imports, and/or uses in the United States the Google Cloud Datastore platform, including but not limited to Google Cloud Datastore products and services and server-side and client-side products and services running on and/or interfacing with Google Cloud Datastore products and/or services (collectively, “Google Cloud Datastore”).



The image is a screenshot of a marketing page for Google Cloud Datastore. It features a white background with blue text and icons. At the top left, the heading "Highly Scalable NoSQL Database" is followed by a paragraph describing the database's scalability and features like sharding, replication, ACID transactions, and SQL-like queries. To the right of this text is a diagram showing a central node with a refresh icon connected to three smaller nodes, each containing a tree-like structure. Below the main text, there are three hexagonal icons: a target, a server rack, and a chip, with the letters "API" in the center. To the right of these icons is the heading "Simple & Integrated" and a paragraph explaining the RESTful interface and integration with App Engine and Compute Engine.

Highly Scalable NoSQL Database

Cloud Datastore is a highly-scalable NoSQL database for your applications. Cloud Datastore **automatically handles sharding and replication**, providing you with a highly available and durable database that scales automatically to handle your applications' load. Cloud Datastore provides a myriad of capabilities such as **ACID transactions**, **SQL-like queries**, **indexes** and **much more**.

Simple & Integrated

With Cloud Datastore's **RESTful interface**, data can easily be accessed by any deployment target. You can build solutions that span across **App Engine** and **Compute Engine**, and rely on Cloud Datastore as the integration point.

Cloud Datastore, GOOGLE CLOUD PLATFORM, <https://cloud.google.com/datastore> (accessed Oct. 19, 2015).

Easy to Use Query Language

Datastore is a schemaless database, which allows you to worry less about making changes to your underlying data structure as your application evolves. Datastore **provides a powerful query engine** that allows you to search for data across multiple properties and sort as needed.

```
1.
2. // List Google companies with less than 400 employees.
3. var companies = query.filter('name =', 'Google').filter('size <', 400);
```

Server-Side Encryption

Cloud Datastore automatically encrypts all data before it is written to disk, at no additional charge. There is no setup or configuration required, no need to modify the way you access the service and no visible performance impact. The data is automatically and transparently decrypted when read by an authorized user.

With server-side encryption, Google manages the cryptographic keys on your behalf using the same hardened key management systems that we use for our own encrypted data, including strict key access controls and auditing. Each Datastore object's data and metadata is encrypted under the [256-bit Advanced Encryption Standard](#), and each encryption key is itself encrypted with a regularly rotated set of master keys.

Server-side encryption can be used in combination with client-side encryption. In client-side encryption, you manage your own encryption keys and encrypt data before writing it to Datastore. In this case, your data is encrypted twice, once with your keys and once with Google's keys.

Cloud Datastore: Datastore Concepts Overview, GOOGLE CLOUD PLATFORM, <https://cloud.google.com/datastore/docs/concepts/overview> (accessed Oct. 19, 2015).

411. Google designs, makes, sells, offers to sell, imports, and/or uses in the United States the Google Cloud Bigtable platform, including but not limited to Google Cloud Bigtable products and services and server-side and client-side products and services running on and/or interfacing with Google Cloud Bigtable products and services (collectively, "Google Cloud Bigtable").

CLOUD BIGTABLE FEATURES

Cloud Bigtable is a fast, fully managed, massively scalable NoSQL database service.


<p>High Performance Cloud Bigtable has a higher performance under high load than alternative products. What this means is that large applications and workflows are faster, more reliable, and more efficient running on Bigtable.</p>	<p>Redundant Autoscaling Storage Cloud Bigtable is built with a redundant internal storage strategy for high durability. You don't need to configure separate storage or disks, and you only pay for the amount of storage you are using.</p>
<p>Security & Permissions All data is encrypted both in-flight and at rest. You have full control over who has access to the data stored in Cloud Bigtable.</p>	<p>Scaling During operation without the need for a restart, allowing efficient use of resources and helping your applications and workflows stay up and running.</p>
<p>Low Latency Storage Cloud Bigtable utilizes a low-latency storage stack, enabling single-digit millisecond latency at the 99th percentile, compared to more than 50x that latency with alternative products.</p>	<p>Industry Standard API Cloud Bigtable is offered through the same open source API as HBase, the native Hadoop database. This enables portability of applications between HBase and Bigtable.</p>
<p>Global Availability Cloud Bigtable is available in regions around the world, allowing you to place your service and data exactly where you want it.</p>	<p>Seamless Cluster Cloud Bigtable cluster nodes can be dynamically added and removed.</p>
<p>Fully Managed Cloud Bigtable is offered as a fully managed service, meaning you spend your time developing valuable applications instead of configuring and tuning your database for performance and scalability. In addition, Google's own Bigtable operations team monitors the service to ensure issues are addressed quickly.</p>	

Cloud Bigtable, GOOGLE CLOUD PLATFORM, <https://cloud.google.com/bigtable> (accessed Oct. 19, 2015).

412. Google designs, makes, sells, offers to sell, imports, and/or uses in the United States the Google Cloud SQL platform, including but not limited to Google Cloud SQL products and services and server-side and client-side products and services running on and/or interfacing with Google Cloud SQL products and/or services (collectively, “Google Cloud SQL”).

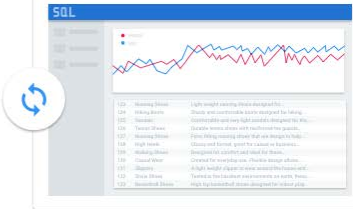
Security & Reliability

Your data is **automatically encrypted and replicated** in many geographic locations and failover between copies are handled automatically. This means your data is protected and your database is available even in the event of a major failure. Google manages your backups, making it easy for you to restore when needed, including point-in-time recovery. Cloud SQL is ISO/IEC 27001 compliant.



A Cloud MySQL Database

Google Cloud SQL is a **fully-managed database service** that makes it easy to set-up, maintain, manage and administer your relational MySQL databases in the cloud. Cloud SQL allows you to focus on your applications rather than administering your databases. Hosted on Google Cloud Platform, Cloud SQL **provides a database infrastructure for applications running anywhere.**



Cloud SQL, GOOGLE CLOUD PLATFORM, <https://cloud.google.com/sql/> (accessed Oct. 19, 2015).

413. Google designs, makes, sells, offers to sell, imports, and/or uses in the United States Google-branded application products and services running on and/or interfacing with the Google Cloud Datastore, Google Cloud SQL, and/or Google Cloud Bigtable cloud database products and/or services (e.g., Google Analytics, Gmail, YouTube) (collectively, “Google Cloud Applications”).

414. Google designs, makes, sells, offers to sell, imports, and/or uses in the United States Google Cloud, Google App Engine, Google Cloud Datastore, Google Cloud Bigtable, Google Cloud SQL, and Google Cloud Applications (collectively, the “Google ‘941 Products”).

415. On information and belief, Google performs (e.g., through operation of the Google ‘941 Products) at least one method for controlling access to a plurality of records provided within a plurality of automated electronic databases, each record having an associated set of access rules. For example, the Google Cloud Platform and/or Google App Engine PaaS

identity and access management broker (the “Google Cloud Broker”) controls access to a plurality of Google Cloud database records provided within a plurality of automated electronic databases (e.g., Google Cloud Datastore NoSQL tables; Google Cloud Bigtable NoSQL tables; and/or Google Cloud SQL instances), each record having an associated set of access rules (e.g., account-level Google Cloud Platform and/or Google App Engine PaaS API access rules; database-level Cloud Datastore, Cloud Bigtable, and/or CloudSQL native access rules).

Server and Software Stack Security

At Google we run tens of thousands of identical, custom-built servers. We've built everything from hardware and networking to the custom Linux software stack with security in mind. Homogeneity, combined with ownership of the entire stack, greatly reduces our security footprint and allows us to react to threats faster.

[Learn more about server and software stack security](#)

Data Access

Google has controls and practices to protect the security of customer information. The layers of the Google application and storage stack require that requests coming from other components are authenticated and authorized. Access by production application administrative engineers to production environments is also controlled. A centralized group and role management system is used to define and control engineers' access to production services, using a security protocol that authenticates engineers through the use of short-lived personal public key certificates; issuance of personal certificates is in turn guarded by two-factor authentication.

[Learn more about data access](#)

Benefits and features

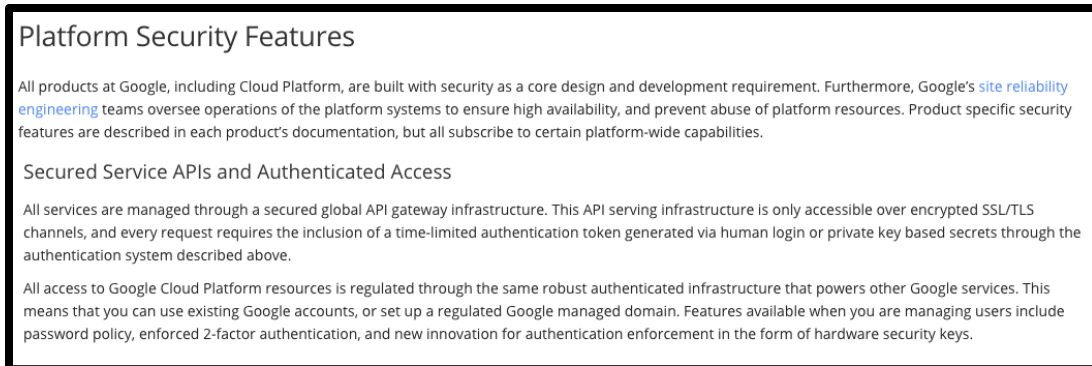
Cloud IAM includes the following features:

- **Single access control interface.** Cloud IAM provides a simple and consistent access control interface for all Cloud Platform services. You can learn one access control interface and apply that knowledge to all Cloud Platform services.
- **Resource-level access control.** You can assign roles to users to access resources at a granularity *finer than the project level*. For example, you can create a policy that assigns the *subscriber* role to a user for a particular Pub/Sub topic.
- **Flexible roles.** Prior to Cloud IAM, you could only assign owner, editor, or viewer roles to users. Cloud Platform products now expose additional flexible roles. For example, the Pub/Sub service exposes *publisher* and *subscriber* roles in addition to the owner, editor, and viewer roles.
- **UI and REST APIs.** You can create and manage Cloud IAM policies using the Google Developers Console or the Cloud IAM APIs.
- **Google account support.** Cloud IAM supports standard Google accounts. You can create Cloud IAM policies to grant permission to a *Google group*, a *Google-hosted domain*, a *service account*, or specific *Google account* holders. You can centrally manage users and groups through the *Google Apps Admin Console*.
- **Available free of charge.** Cloud IAM is offered at no additional charge for all Cloud Platform customers. You will be charged only for use of other Cloud Platform products. For information about the pricing of other Cloud Platform products, see the *Google Cloud Platform Pricing Calculator*.

Cloud Identity and Access Management: What is Cloud IAM?, GOOGLE CLOUD PLATFORM, <https://cloud.google.com/iam> (accessed Oct. 19, 2015).

416. On information and belief, at least one method for controlling access to a plurality of records provided within a plurality of automated electronic database comprises Google (e.g., through operation of the Google ‘941 Products) receiving a request from a requestor, the requestor having at least one attribute. For example, the Google Cloud Broker receives a Google

Cloud API (e.g., RESTful) request from a requestor (e.g., “RequestApp”, a Google Cloud-enrolled web or mobile application), the requestor having at least one attribute (e.g., RequestApp’s account, user, role, geographic location, and/or client platform information for the request).



Google Cloud Platform Security, GOOGLE CLOUD PLATFORM, <https://cloud.google.com/security/> (accessed Oct. 19, 2015).

417. On information and belief, at least one method for controlling access to a plurality of records provided within a plurality of automated electronic database comprises Google (e.g., through operation of the Google ‘941 Products) searching the plurality of automated electronic databases to find records relating to an entity corresponding to the request, and records having connections to records corresponding to the request, relating to transactions, relationships or communications between the entity and another entity. For example, the Google Cloud Broker searches the plurality of automated electronic databases (e.g., Cloud Datastore tables; Cloud Bigtable tables; CloudSQL instances) to find records (e.g., cloud-synced RequestApp data) relating to an entity (e.g., a unique RequestApp end user, for example, “JaneDoe555”) corresponding to the request, and records having connections to records corresponding to the request, relating to transactions, relationships, or communications between the entity and another entity. For example, where the Google Cloud API request includes and/or references “JaneDoe555” in the field “userID,” the Google Cloud Broker searches the plurality of automated electronic databases to find (as particular examples): all cloud-synced RequestApp data corresponding to (e.g., created and/or owned by) JaneDoe555; and/or selected (e.g., allowed

via RequestApp-dependent sharing and/or privacy permissions) cloud-synced RequestApp data corresponding to (e.g., created and/or owned by) RequestApp entities related to JaneDoe555—for example, RequestApp users within JaneDoe555’s same RequestApp account (e.g., ExampleEnterpriseLLP), group (e.g., Accounting), and/or role.

Entities

Objects in the Datastore are known as *entities*. An entity has one or more named *properties*, each of which can have one or more values. Property values can belong to a variety of data types, including integers, floating-point numbers, strings, dates, and binary data, among others. A query on a property with multiple values tests whether any of the values meets the query criteria. This makes such properties useful for membership testing.

Note: Datastore entities are *schemaless*: unlike traditional relational databases, the Datastore does not require that all entities of a given kind have the same properties or that all of an entity's values for a given property be of the same data type. If a formal schema is needed, the application itself is responsible for ensuring that entities conform to it.

Kinds, keys, and identifiers

Each Datastore entity is of a particular *kind*, which categorizes the entity for the purpose of queries; for instance, a human resources application might represent each employee at a company with an entity of kind `Employee`. In addition, each entity has its own *key*, which uniquely identifies it. The key consists of the following components:

- The entity's kind
- An *identifier*, which can be either
 - a *key name* string
 - an integer *ID*
- An optional *ancestor path* locating the entity within the Datastore hierarchy

The identifier is assigned when the entity is created. Because it is part of the entity's key, it is associated permanently with the entity and cannot be changed. It can be assigned in either of two ways:

- Your application can specify its own key name string for the entity.
- You can have the Datastore automatically assign the entity an integer numeric ID.

Cloud Datastore: Datastore Concepts Overview, GOOGLE CLOUD PLATFORM, <https://cloud.google.com/datastore/docs/concepts/overview> (accessed Oct. 19, 2015).

418. On information and belief, at least one method for controlling access to a plurality of records provided within a plurality of automated electronic database comprises Google (e.g., through operation of the Google ‘941 Products) searching the plurality of automated electronic databases to find records in dependence on the request and on connections between respective records. For example, the Google Cloud Broker searches the plurality of automated electronic databases (e.g., Cloud Datastore tables; Cloud Bigtable tables; CloudSQL instances) to find records (e.g., cloud-synced RequestApp data) relating to an entity (e.g., a unique RequestApp

end user, for example, “JaneDoe555”) corresponding to the request, and respective records having respective connections to respective records corresponding to the request. For example, where the Google Cloud API request includes and/or references “JaneDoe555” in the field “userID,” the Google Cloud Broker searches the plurality of automated electronic databases to find (as particular examples): all cloud-synced RequestApp data corresponding to (e.g., created and/or owned by) JaneDoe555; and/or selected (e.g., allowed via RequestApp-dependent sharing and/or privacy permissions) cloud-synced RequestApp data corresponding to (e.g., created and/or owned by) RequestApp entities connected to JaneDoe555—for example, RequestApp users within JaneDoe555’s same RequestApp account (e.g., ExampleEnterpriseLLP), group (e.g., Accounting), and/or role.

Ancestor paths

Entities in the Datastore form a hierarchically structured space similar to the directory structure of a file system. When you create an entity, you can optionally designate another entity as its *parent*; the new entity is a *child* of the parent entity (note that unlike in a file system, the parent entity need not actually exist). An entity without a parent is a *root entity*. The association between an entity and its parent is permanent, and cannot be changed once the entity is created. The Datastore will never assign the same numeric ID to two entities with the same parent, or to two root entities (those without a parent).

An entity's parent, parent's parent, and so on recursively, are its *ancestors*; its children, children's children, and so on, are its *descendants*. A root entity and all of its descendants belong to the same *entity group*. The sequence of entities beginning with a root entity and proceeding from parent to child, leading to a given entity, constitute that entity's *ancestor path*. The complete key identifying the entity consists of a sequence of kind-identifier pairs specifying its ancestor path and terminating with those of the entity itself:

```
[Person:GreatGrandpa, Person:Grandpa, Person:Dad, Person:Me]
```

For a root entity, the ancestor path is empty and the key consists solely of the entity's own kind and identifier:

```
[Person:GreatGrandpa]
```

This concept is illustrated by the following diagram:

A query can also include an *ancestor filter* limiting the results to just the entity group descended from a specified ancestor. Such a query is known as an *ancestor query*. By default, ancestor queries return *strongly consistent* results, which are guaranteed to be up to date with the latest changes to the data. Non-ancestor queries, by contrast, can span the entire Datastore rather than just a single entity group, but are only *eventually consistent* and may return stale results. If strong consistency is important to your application, you may need to take this into account when structuring your data, placing related entities in the same entity group so they can be retrieved with an ancestor rather than a non-ancestor query; see [Structuring Data for Strong Consistency](#) for more information.

Every Datastore query computes its results using one or more *indexes*, tables containing entities in a sequence specified by the index's properties and, optionally, the entity's ancestors. The indexes are updated incrementally to reflect any changes the application makes to its entities, so that the correct results of all queries are available with no further computation needed.

The Datastore predefines a simple index on each property of an entity. You can define further custom indexes in an *index configuration file* named `index.yaml`. The development server automatically adds suggestions to this file as it encounters queries that cannot be executed with the existing indexes. You can tune indexes manually by editing the file before uploading the application.

Cloud Datastore: Datastore Concepts Overview, GOOGLE CLOUD PLATFORM, <https://cloud.google.com/datastore/docs/concepts/overview> (accessed Oct. 19, 2015).

419. On information and belief, at least one method for controlling access to a plurality of records provided within a plurality of automated electronic database comprises Google (e.g., through operation of the Google '941 Products) applying a set of access rules associated with each found record by at least one automated processor, to produce a set of accessible records, at a server device. For example, the Google Cloud Broker applies a set of access rules (e.g., account-level Google Cloud Platform and/or Google App Engine PaaS API access rules; database-level Cloud Datastore, Cloud Bigtable, and/or CloudSQL native access rules)

associated with each found record (e.g., piece of cloud-synced RequestApp data) by at least one automated processor (e.g., Google server automated processor), to produce a set of accessible records (e.g., accessible cloud-synced RequestApp data owned, controlled, and/or shared with Jane555) at a Google Cloud server device.

Projects and resources

Google Cloud Platform uses [projects](#) as the unit of resource ownership. All Google Cloud Platform resources — such as App Engine apps, Compute Engine instances, Cloud Storage buckets, API keys, and service accounts — must be owned by projects. A project can be used by multiple users working in collaboration.

Authentication and authorization

Authentication is the process of determining the identity of a client, which is typically a user account or a service account. Authorization is the process of determining what permissions an authenticated identity has on a set of specified resources.

In terms of Google Cloud APIs, there can be no authorization without authentication, therefore, you may see the term authentication (auth) used to refer to both authentication and authorization.

OAuth scopes

Scopes is an OAuth feature that limits the permissions of an OAuth credential. A user account or a service account may have a wide range of permissions on a resource, for example, Storage Read and Storage Write on a Cloud Storage bucket. For security and privacy reasons, your application may not need all these permissions. The OAuth scope feature lets you limit the credential to a subset of these permissions/scopes (for example, Storage Read). You can find the list of scopes that a Google API method requires in its reference documentation. See the Google Cloud Pub/Sub [projects.topics.list method](#) for an example.

Application Default Credentials

[Application Default Credentials](#) (ADC) is part of Google APIs client libraries that simplifies authentication to Google Cloud APIs. It abstracts authentication across different environments, allowing you to authenticate with just a single line of code. It is designed for use cases where application requests have the same authenticated identity and authorization scope, independent of user.

Developer workflow

All scenarios boil down to the same workflow. Your application obtains an OAuth credential and then makes a request to a Google Cloud API, passing in an access token derived from the credential. The API uses the access token to decide whether to allow your request.

Google Cloud Platform Auth Guide, GOOGLE CLOUD PLATFORM, <https://cloud.google.com/docs/authentication> (accessed Oct. 19, 2015).

420. On information and belief, at least one method for controlling access to a plurality of records provided within a plurality of automated electronic database comprises Google (e.g., through operation of the Google ‘941 Products) linking the set of accessible records (e.g., piece of cloud-synced RequestApp data determined to be accessible by the Google Cloud Broker upon application of the account-level Google Cloud Platform and/or Google App Engine PaaS API

access rules and/or database-level Cloud Datastore, Cloud Bigtable, and/or CloudSQL native access rules) into an information polymer using a server device (e.g., Google Cloud server device).

421. On information and belief, at least one method for controlling access to a plurality of records provided within a plurality of automated electronic database comprises Google (e.g., through operation of the Google '941 Products) applying at least one compensation rule by at least one automated processor, dependent on the at least one attribute of the requestor. For example, the Google Cloud Broker applies at least one Google Cloud Platform and/or Google App Engine PaaS compensation rule relating to the API request and/or data retrieval. The at least one compensation rule is dependent on the at least one attribute of the requestor—for example, Google Cloud Platform and/or Google App Engine PaaS account, group, and/or user information in the request determines both who Google will charge (e.g., RequestApp's developer and/or RequestApp's MBaaS provider) and the amount Google will charge for the request (e.g., through identification and application of RequestApp's associated Google Cloud Platform and/or Google App Engine PaaS quotas and billing rates).

422. On information and belief, at least one method for controlling access to a plurality of records provided within a plurality of automated electronic database comprises Google (e.g., through operation of the Google '941 Products) logging at least the request for access by at least one automated processor. For example, the Google Cloud Broker logs at least RequestApp's Google Cloud API request by at least one Google server automated processor.

423. On information and belief, at least one method for controlling access to a plurality of records provided within a plurality of automated electronic database comprises Google (e.g., through operation of the Google '941 Products) communicating the set of accessible records and/or the information polymer to the requestor. For example, the Google Cloud Broker communicates the set of accessible records and/or the information polymer (e.g., accessible cloud-synced RequestApp data owned, controlled, and/or shared with Jane555 and/or an

information polymer comprising the foregoing data) to RequestApp and/or a designated intermediary server.

424. On information and belief, Google designs, makes, sells, offers to sell, imports and/or uses in the United States at least one apparatus adapted for performing at least one method for controlling access to a plurality of records provided within a plurality of automated electronic databases, each record having an associated set of access rules. For example, the Google '941 Products comprise at least one apparatus (e.g., a Google Cloud Broker physical and/or virtual appliance) adapted (e.g., through specially designed and/or programmed hardware and/or software components) for controlling access to a plurality of Google Cloud database records provided within a plurality of automated electronic databases (e.g., Google Cloud Datastore NoSQL tables; Google Cloud Bigtable NoSQL tables; and/or Google Cloud SQL instances), each record having an associated set of access rules (e.g., account-level Google Cloud Platform and/or Google App Engine PaaS API access rules; database-level Cloud Datastore, Cloud Bigtable, and/or CloudSQL native access rules).

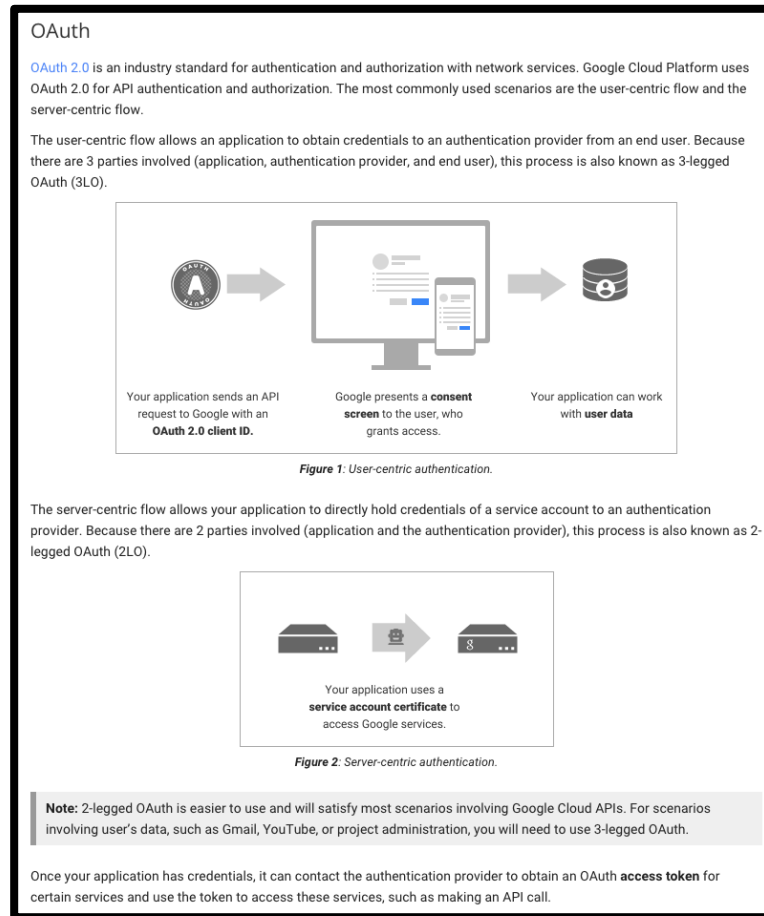
425. On information and belief, the at least one apparatus (e.g., Google Cloud Broker physical and/or virtual appliance) comprises an input port (e.g., a physical and/or virtual network communications port) configured (e.g., through specially designed and/or programmed hardware and/or software components) to receive a request (e.g., Google Cloud API request) for access to one or more records associated with at least one entity (e.g., one or more Google Cloud database records associated with at least one entity) from a requestor (e.g., "RequestApp," a Google Cloud-enrolled web or mobile application) having at least one attribute (e.g., RequestApp's account, user, role, geographic location, and/or client platform information for the request).

426. On information and belief, the at least one apparatus (e.g., Google Cloud Broker physical and/or virtual appliance) comprises at least one processor (e.g., Google server automated processor) configured (e.g., through specially designed and/or programmed hardware and/or software components) to automatically define a query (a Google Cloud Datastore, Google

Cloud Bigtable, and/or Google Cloud SQL query) based on the received Google Cloud API request.


427. On information and belief, the at least one apparatus (e.g., Google Cloud Broker physical and/or virtual appliance) comprises at least one processor (e.g., Google server automated processor) configured (e.g., through specially designed and/or programmed hardware and/or software components) to send the query (e.g., Google Cloud Datastore, Google Cloud Bigtable, and/or Google Cloud SQL query) to the plurality of automated electronic databases (e.g., Google Cloud Datastore tables; Google Cloud Bigtable tables; Google Cloud SQL instances).

428. On information and belief, the at least one apparatus (e.g., Google Cloud Broker physical and/or virtual appliance) comprises at least one processor (e.g., Google server automated processor) configured (e.g., through specially designed and/or programmed hardware and/or software components) to determine a set of records (e.g., Google Cloud Datastore, Google Cloud Bigtable, and/or Google Cloud SQL records comprising cloud-synced RequestApp data) associated with the at least one entity (e.g., a unique RequestApp end user, for example, “JaneDoe555”) contained in each respective automated electronic database based on the query (e.g., Google Cloud Datastore tables; Google Cloud Bigtable tables; Google Cloud SQL instance) and connections between respective records identified by the query (for example, selected (e.g., allowed via RequestApp-dependent sharing and/or privacy permissions) cloud-synced RequestApp data corresponding to (e.g., created and/or owned by) RequestApp entities related to JaneDoe555—for example, RequestApp users within JaneDoe555’s same RequestApp account (e.g., ExampleEnterpriseLLP), group (e.g., Accounting), and/or role).



Google Cloud Platform Auth Guide, GOOGLE CLOUD PLATFORM, <https://cloud.google.com/docs/authentication> (accessed Oct. 19, 2015).

429. On information and belief, the at least one apparatus (e.g., Google Cloud Broker physical and/or virtual appliance) comprises at least one processor (e.g., Google server automated processor) configured (e.g., through specially designed and/or programmed hardware and/or software components) to implement at least one compensation rule (at least one Google Cloud Platform and/or Google App Engine PaaS compensation rule) dependent on the at least one attribute of the requestor (Google Cloud Platform and/or Google App Engine PaaS account, group, and/or user information in the request).



Pay-per-use Billing

Our pay-per-use option makes it economical to get started. If you're running a lightly or sporadically used database, you'll **save money by only paying for the time you access your data**. The package option allows you to control your costs for more heavily loaded instances.

Cloud SQL: Pay-per-use Billing, GOOGLE CLOUD PLATFORM, <https://cloud.google.com/sql/> (accessed Oct. 19, 2015).

CLOUD DATASTORE PRICING
 Cloud Datastore is a highly-scalable NoSQL database for your web and mobile applications

	FREE LIMIT PER DAY	PRICE ABOVE FREE LIMIT
Stored data	1 GB storage	\$0.18 / GB / month
Read Operations	50k	\$0.06/100k operations
Write Operations	50k	\$0.06/100k operations
Small Operations	50k	Free

Cloud Datastore: Cloud Datastore Pricing, GOOGLE CLOUD PLATFORM, <https://cloud.google.com/datastore/> (accessed Oct. 19, 2015).

430. On information and belief, the at least one apparatus (e.g., Google Cloud Broker physical and/or virtual appliance) comprises at least one processor (e.g., Google server automated processor) configured (e.g., through specially designed and/or programmed hardware and/or software components) to aggregate the determined set of records (e.g., Google Cloud Datastore, Google Cloud Bigtable, and/or Google Cloud SQL records comprising cloud-synced RequestApp data, determined to be responsive by the Google Cloud Broker as described in preceding paragraphs of this Count) based on the received responses (e.g., the responses received by the Google Cloud Broker from the plurality of automated electronic Google Cloud Datastore tables, Google Cloud Bigtable tables, and/or Google Cloud SQL instances).

431. On information and belief, the at least one apparatus (e.g., Google Cloud Broker physical and/or virtual appliance) comprises at least one processor (e.g., Google server automated processor) configured (e.g., through specially designed and/or programmed hardware and/or software components) to communicate, through a communications port (e.g., a Google

server physical and/or virtual network communications port), the aggregation of the determined set of records (e.g., Google Cloud Datastore, Google Cloud Bigtable, and/or Google Cloud SQL records comprising cloud-synced RequestApp data), available to the requestor (e.g., RequestApp) in dependence on the access rules (e.g., account-level Google Cloud Platform and/or Google App Engine PaaS API access rules; database-level Cloud Datastore, Cloud Bigtable, and/or CloudSQL native access rules), from the plurality of automated electronic databases (e.g., Google Cloud Datastore tables, Google Cloud Bigtable tables, and/or Google Cloud SQL instances).

432. On information and belief, the at least one apparatus (e.g., Google Cloud Broker physical and/or virtual appliance) comprises at least one processor (e.g., Google server automated processor) configured (e.g., through specially designed and/or programmed hardware and/or software components) to log at least the request for access to the one or more records.

Logging

All platform API requests, such as web requests, storage bucket access, and user account access, are logged. With Cloud Platform tools, you can read operations and access logs for [Compute Engine](#), [App Engine](#), [BigQuery](#), [Cloud SQL](#), [Deployment Manager](#), [Cloud VPN](#), and [Cloud Storage](#).

Google Cloud Platform Security: Platform Security Features, GOOGLE CLOUD PLATFORM, <https://cloud.google.com/security/> (accessed Oct. 19, 2015).

433. On information and belief, the Google '941 Products are made, sold, and/or offered for sale by and/or on behalf of Google to entities (e.g., businesses, schools, and other organizations) and individuals throughout the United States.

434. On information and belief, the Google '941 Products are made, sold, and offered for sale by and/or on behalf of Google to entities (e.g., businesses, schools, and other organizations) and individuals located in the Eastern District of Texas.

435. On information and belief, the Google '941 Products are used by Google (e.g., by and/or on behalf of Google employees) throughout the United States.

436. On information and belief, the Google '941 Products are used by Google (e.g., by and/or on behalf of Google employees) within the Eastern District of Texas.

437. On information and belief, Google has directly infringed and continues to directly infringe the '941 patent by, among other things, directly performing the method of claims 1 and 16 of the patent through operation of at least the Google '941 Products.

438. By making, using, offering for sale, and/or selling infringing products and services for managing access to protected data, including but not limited to the Google '941 Products, Google has injured St. Luke and is liable to St. Luke for directly infringing one or more claims of the '941 patent, including at least claims 1, 8, and 16, pursuant to 35 U.S.C. § 271(a).

439. On information and belief, Google also indirectly infringes the '941 patent by actively inducing infringement under 35 USC § 271(b).

440. On information and belief, at least since service of this Complaint or shortly thereafter, Google has known of the '941 patent and has known about infringement of the '941 patent by Google itself and by third-party Google customers, end-users, developers, and/or integrators/partners of the Google '941 Products.

441. On information and belief, beginning no later than the date of service of this Complaint, Google has intentionally performed acts that induce infringement of the '941 patent by third parties (e.g., Google '941 Product customers, end-users, developers, and/or integrators/partners), knowing that these acts would induce third-party infringement of the '941 patent and/or with willful blindness to this fact.

442. For example, on information and belief, Google provides products and services (e.g., the Google '941 Products) capable of infringing one or more claims of the '941 patent, including at least claims 1, 8, and 16.

443. For example, on information and belief, Google configures these products and services (e.g., the Google '941 Products) to infringe at least one claim of the '941 patent in normal operation by Google customers, end-users, developers, and/or integrators/partners.

444. For example, on information and belief, Google instructs and directs customers, end-users, developers, and/or integrators/partners to make and/or use the Google '941 Products in an infringing manner and/or configuration (e.g., through creation and dissemination of Google '941 Product documentation, training materials, SDKs, client libraries, and API products and services that not only facilitate, but effectively mandate, third-party infringement of the '941 patent by Google customers, end-users, developers, and/or integrators/partners).

445. Accordingly, Google has actively induced and continues to actively induce infringement of the '941 patent by Google '941 Product customers, end-users, developers, and/or integrators/partners.

446. To the extent applicable, the requirements of 35 U.S.C. § 287(a) have been met with respect to the '941 patent.

447. As a result of Google's infringement of the '941 patent, St. Luke has suffered monetary damages, and seeks recovery in an amount adequate to compensate for Google's infringement, but in no event less than a reasonable royalty for the use made of the '941 patent inventions by Google, together with interest and costs as fixed by the Court.

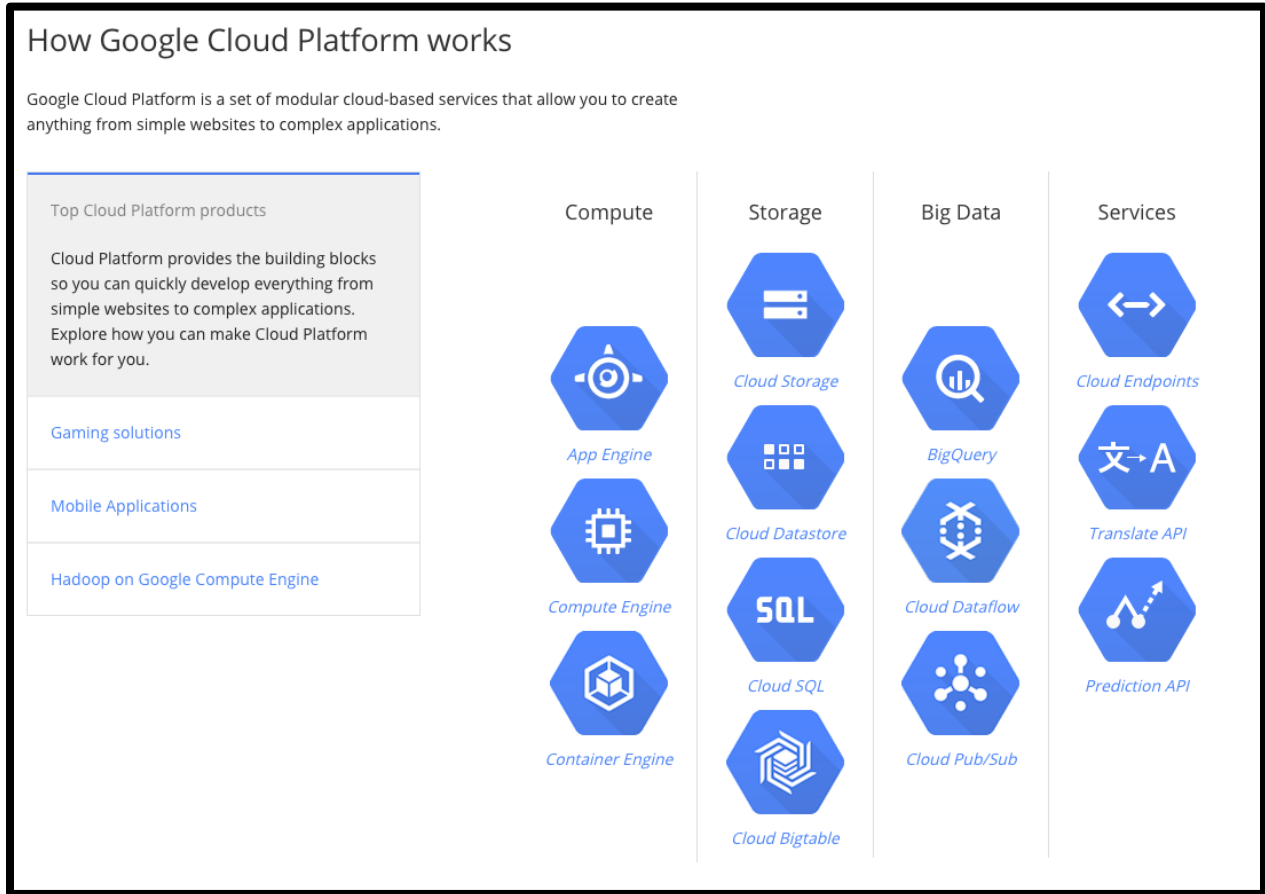
COUNT VI
INFRINGEMENT OF U.S. PATENT NO. 8,380,630

448. St. Luke references and incorporates by reference the preceding paragraphs of the Complaint as if fully set forth herein.

449. Google designs, makes, sells, offers to sell, imports, and/or uses in the United States products and/or services for managing access to protected data.

450. Google designs, makes, sells, offers to sell, imports, and/or uses in the United States the Google Cloud Platform, including but not limited to Google Cloud Compute products and services (e.g., Google App Engine, Google Compute Engine, Google Container Engine); Google Cloud Storage products and services (e.g., Google Cloud Storage, Google Cloud Datastore, Google Cloud SQL, Google Cloud Bigtable); Google Cloud Big Data products and services (e.g., Google Cloud BigQuery, Google Cloud Dataflow, Google Cloud Pub/Sub); and

Google Cloud Services products and services (e.g., Google Cloud Endpoints, Google Translate API, Google Prediction API) (collectively, “Google Cloud”).



How Google Cloud Platform Works, <https://cloud.google.com> (retrieved Oct. 19, 2015).

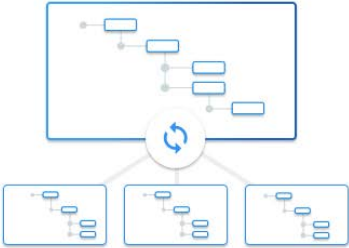
451. Google designs, makes, sells, offers to sell, imports, and/or uses in the United States the Google App Engine platform-as-a-service, including but not limited to Google App Engine products and services and server-side and client-side products and services running on and/or interfacing with Google App Engine products and/or services (collectively, “Google App Engine”).

452. Google designs, makes, sells, offers to sell, imports, and/or uses in the United States the Google Cloud Datastore platform, including but not limited to Google Cloud Datastore products and services and server-side and client-side products and services running on and/or


interfacing with Google Cloud Datastore products and/or services (collectively, “Google Cloud Datastore”).

Highly Scalable NoSQL Database

Cloud Datastore is a highly-scalable NoSQL database for your applications. Cloud Datastore **automatically handles sharding and replication**, providing you with a highly available and durable database that scales automatically to handle your applications' load. Cloud Datastore provides a myriad of capabilities such as **ACID transactions, SQL-like queries, indexes and much more.**



Simple & Integrated



With Cloud Datastore's **RESTful interface**, data can easily be accessed by any deployment target. You can build solutions that span across **App Engine and Compute Engine**, and rely on Cloud Datastore as the integration point.

Easy to Use Query Language

Datastore is a schemaless database, which allows you to worry less about making changes to your underlying data structure as your application evolves. Datastore **provides a powerful query engine** that allows you to search for data across multiple properties and sort as needed.

```
1.  
2. // List Google companies with less than 400 employees.  
3. var companies = query.filter('name =', 'Google').filter('size <', 400);
```

Cloud Datastore, GOOGLE CLOUD PLATFORM, <https://cloud.google.com/datastore> (accessed Oct. 19, 2015).

Server-Side Encryption

Cloud Datastore automatically encrypts all data before it is written to disk, at no additional charge. There is no setup or configuration required, no need to modify the way you access the service and no visible performance impact. The data is automatically and transparently decrypted when read by an authorized user.

With server-side encryption, Google manages the cryptographic keys on your behalf using the same hardened key management systems that we use for our own encrypted data, including strict key access controls and auditing. Each Datastore object's data and metadata is encrypted under the [256-bit Advanced Encryption Standard](#), and each encryption key is itself encrypted with a regularly rotated set of master keys.

Server-side encryption can be used in combination with client-side encryption. In client-side encryption, you manage your own encryption keys and encrypt data before writing it to Datastore. In this case, your data is encrypted twice, once with your keys and once with Google's keys.

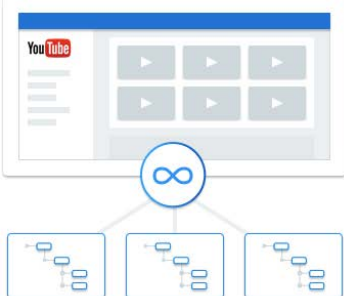
Cloud Datastore: Datastore Concepts Overview, GOOGLE CLOUD PLATFORM, <https://cloud.google.com/datastore/docs/concepts/overview> (accessed Oct. 19, 2015).

453. Google designs, makes, sells, offers to sell, imports, and/or uses in the United States the Google Cloud Bigtable platform, including but not limited to Google Cloud Bigtable products and services and server-side and client-side products and services running on and/or interfacing with Google Cloud Bigtable products and/or services (collectively, "Google Cloud Bigtable").

Massively Scalable NoSQL

Cloud Bigtable is Google's NoSQL Big Data database service. Because it is designed to handle massive workloads at **consistent low latency and high throughput**.

Bigtable is a great choice for both operational and analytical applications such as: **IoT, user analytics and financial data analysis**.



CLOUD BIGTABLE FEATURES

Cloud Bigtable is a fast, fully managed, massively scalable NoSQL database service.


<p>High Performance Cloud Bigtable has a higher performance under high load than alternative products. What this means is that large applications and workflows are faster, more reliable, and more efficient running on Bigtable.</p> <p>Security & Permissions All data is encrypted both in-flight and at rest. You have full control over who has access to the data stored in Cloud Bigtable.</p> <p>Low Latency Storage Cloud Bigtable utilizes a low-latency storage stack, enabling single-digit millisecond latency at the 99th percentile, compared to more than 50x that latency with alternative products.</p> <p>Global Availability Cloud Bigtable is available in regions around the world, allowing you to place your service and data exactly where you want it.</p> <p>Fully Managed Cloud Bigtable is offered as a fully managed service, meaning you spend your time developing valuable applications instead of configuring and tuning your database for performance and scalability. In addition, Google's own Bigtable operations team monitors the service to ensure issues are addressed quickly.</p>	<p>Redundant Autoscaling Storage Cloud Bigtable is built with a redundant internal storage strategy for high durability. You don't need to configure separate storage or disks, and you only pay for the amount of storage you are using.</p> <p>Scaling During operation without the need for a restart, allowing efficient use of resources and helping your applications and workflows stay up and running.</p> <p>Industry Standard API Cloud Bigtable is offered through the same open source API as HBase, the native Hadoop database. This enables portability of applications between HBase and Bigtable.</p> <p>Seamless Cluster Cloud Bigtable cluster nodes can be dynamically added and removed.</p>
---	--

Cloud Bigtable, GOOGLE CLOUD PLATFORM, <https://cloud.google.com/bigtable> (accessed Oct. 19, 2015).

454. Google designs, makes, sells, offers to sell, imports, and/or uses in the United States the Google Cloud SQL platform, including but not limited to Google Cloud SQL products and services and server-side and client-side products and services running on and/or interfacing with Google Cloud SQL products and/or services (collectively, “Google Cloud SQL”).

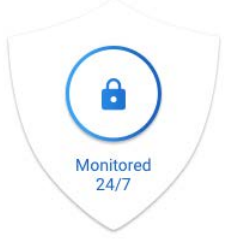
A Cloud MySQL Database

Google Cloud SQL is a **fully-managed database service** that makes it easy to set-up, maintain, manage and administer your relational MySQL databases in the cloud. Cloud SQL allows you to focus on your applications rather than administering your databases. Hosted on Google Cloud Platform, Cloud SQL **provides a database infrastructure for applications running anywhere.**



Security & Reliability

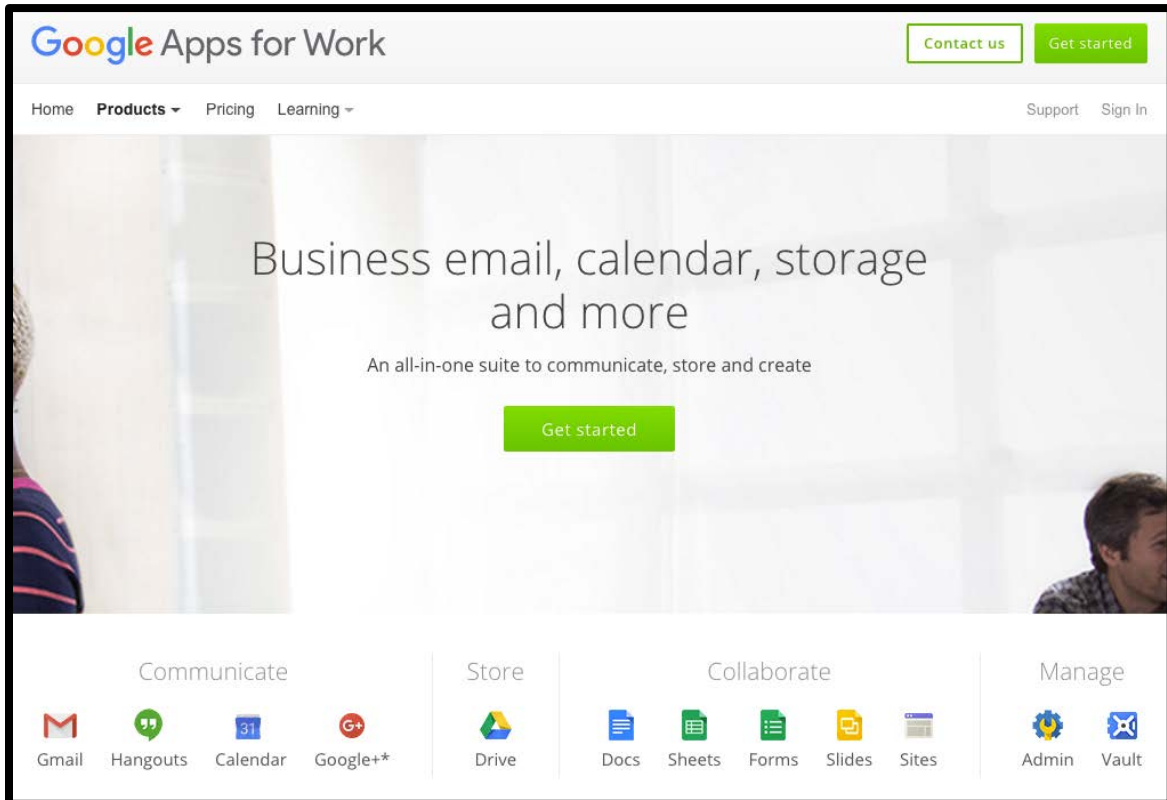
Your data is **automatically encrypted and replicated** in many geographic locations and failover between copies are handled automatically. This means your data is protected and your database is available even in the event of a major failure. Google manages your backups, making it easy for you to restore when needed, including point-in-time recovery. Cloud SQL is ISO/IEC 27001 compliant.



Cloud SQL, GOOGLE CLOUD PLATFORM, <https://cloud.google.com/sql> (accessed Oct. 19, 2015).

455. Google designs, makes, sells, offers to sell, imports, and/or uses in the United States the Google for Work platform, including but not limited to Google Apps for Work, Education, Government, and Nonprofit products and services; Google Drive for Work products and services; and Google Apps Unlimited products and services (collectively, “Google for Work”).

456. Google designs, makes, sells, offers to sell, imports, and/or uses in the United States the Google Apps for Work, Education, Government and Nonprofit platforms, including but not limited to Gmail, Google Hangouts, Google Calendar, Google+, Google Drive, Google Docs, Google Sheets, Google Forms, Google Slides, Google Sites, Google Admin, and Google Vault server-side and client-side products and services (collectively, “Google Apps”).



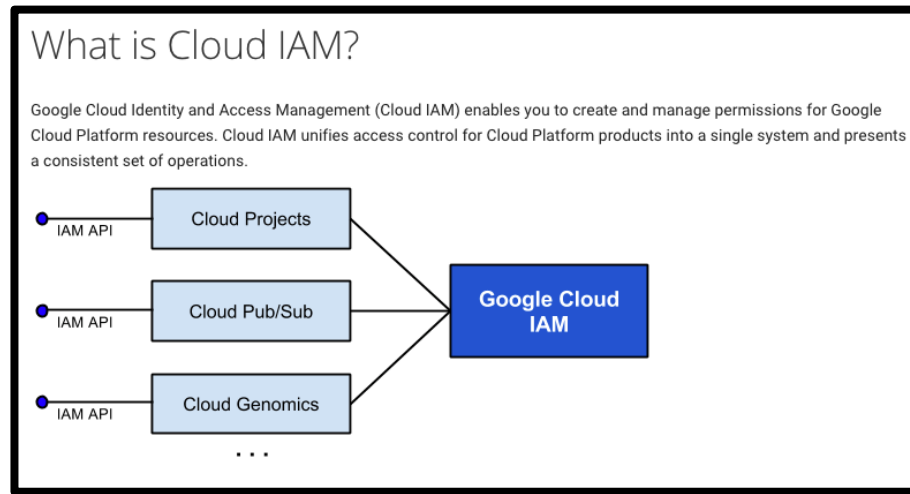
Google Apps for Work, <http://www.google.com/work/apps/business/products/> (retrieved Oct. 19, 2015).

457. Google designs, makes, sells, offers to sell, imports, and/or uses in the United States Google-branded application products and services running on and/or interfacing with the Google Cloud Datastore, Google Cloud SQL, and/or Google Cloud Bigtable cloud database products and/or services (e.g., Google Analytics, Gmail, YouTube) (collectively, “Google Cloud Applications”).

458. Google designs, makes, sells, offers to sell, imports, and/or uses in the United States Google Cloud, Google App Engine, Google Cloud Datastore, Google Cloud Bigtable, Google Cloud SQL, Google for Work, Google Apps, and Google Cloud Applications (collectively, the “Google ‘630 Products’”).

459. On information and belief, Google designs, makes, sells, offers to sell, imports, and/or uses in the United States at least one security mediator. For example, the Google ‘630 Products comprise at least one Google Cloud Platform and/or Google App Engine identity and

access management physical and/or virtual appliance security mediator (the “Google Cloud Broker”).



Cloud Identity and Access Management: What is Cloud IAM?, GOOGLE CLOUD PLATFORM, <https://cloud.google.com/iam> (accessed Oct. 19, 2015).

460. On information and belief, the Google security mediator comprises an input port configured to receive a request for information stored in a plurality of external databases (“POEDs”) from a user. For example, the Google Cloud Broker comprises an input port (e.g., a Google server physical and/or virtual network communications input port) configured to (e.g., adapted through specially-designed and/or programmed hardware and/or software components to automatically) receive a request for information stored in a POEDs (e.g., a Google API HTTP request for information stored in a plurality of Google Cloud Datastore NoSQL tables, Google Cloud Bigtable NoSQL tables, and/or Google Cloud SQL instances external to the Google Cloud Broker) from a user (e.g., “RequestApp,” a requesting Google Cloud-enrolled web or mobile application).

Platform Security Features

All products at Google, including Cloud Platform, are built with security as a core design and development requirement. Furthermore, Google's [site reliability engineering](#) teams oversee operations of the platform systems to ensure high availability, and prevent abuse of platform resources. Product specific security features are described in each product's documentation, but all subscribe to certain platform-wide capabilities.

Secured Service APIs and Authenticated Access

All services are managed through a secured global API gateway infrastructure. This API serving infrastructure is only accessible over encrypted SSL/TLS channels, and every request requires the inclusion of a time-limited authentication token generated via human login or private key based secrets through the authentication system described above.

All access to Google Cloud Platform resources is regulated through the same robust authenticated infrastructure that powers other Google services. This means that you can use existing Google accounts, or set up a regulated Google managed domain. Features available when you are managing users include password policy, enforced 2-factor authentication, and new innovation for authentication enforcement in the form of hardware security keys.

Server and Software Stack Security

At Google we run tens of thousands of identical, custom-built servers. We've built everything from hardware and networking to the custom Linux software stack with security in mind. Homogeneity, combined with ownership of the entire stack, greatly reduces our security footprint and allows us to react to threats faster.

[Learn more about server and software stack security](#)

Data Access

Google has controls and practices to protect the security of customer information. The layers of the Google application and storage stack require that requests coming from other components are authenticated and authorized. Access by production application administrative engineers to production environments is also controlled. A centralized group and role management system is used to define and control engineers' access to production services, using a security protocol that authenticates engineers through the use of short-lived personal public key certificates; issuance of personal certificates is in turn guarded by two-factor authentication.

[Learn more about data access](#)

Google Cloud Platform Security, GOOGLE CLOUD PLATFORM, <https://cloud.google.com/security/> (accessed Oct. 19, 2015).

461. On information and belief, the Google security mediator comprises an automated centralized index ("ACI"), stored in a memory, configured to store location information and associated access rules for information stored in the POEDs. For example, the Google Cloud Broker comprises at least a Google Cloud Platform and/or Google App Engine PaaS identity and access management ACI, stored in a Google Cloud server memory, configured to (e.g., adapted through specially-designed and/or programmed hardware and/or software components to automatically) store location information (e.g., the URI/URN/URL of a respective Google Cloud Datastore, Google Cloud Bigtable, or Google Cloud SQL resource) and associated access rules (e.g., respective resource-specific account-level Google Cloud Platform and/or Google App Engine PaaS API access rules stored in the Google Cloud Platform and/or Google App Engine PaaS identity and access management ACI associated with a respective Google Cloud Datastore, Google Cloud Bigtable, or Google Cloud SQL resource) for information stored in the POEDs

(e.g., respective Google Cloud Datastore, Google Cloud Bigtable, and/or Google Cloud SQL information resources comprising and/or associated with respective pieces of RequestApp data and/or metadata).

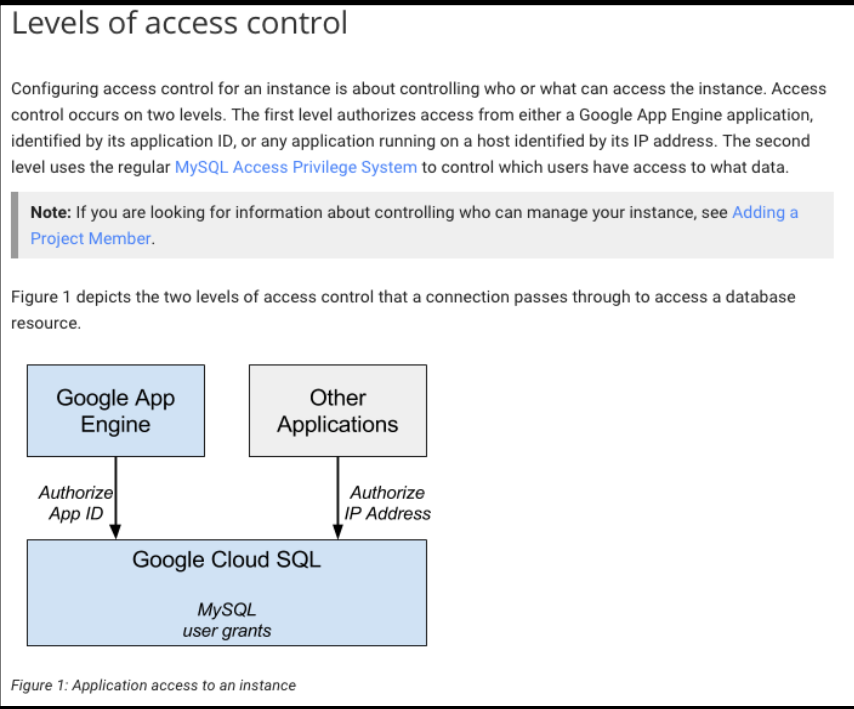
462. On information and belief, the Google security mediator comprises at least one processor configured to locate requested information. For example, the Google Cloud Broker comprises at least one Google server automated processor configured to (e.g., adapted through specially-designed and/or programmed hardware and/or software components to automatically) identify location information (e.g., URI/URN/URL of a respective Google Cloud Datastore, Google Cloud Bigtable, or Google Cloud SQL resource) for requested information (e.g., Google Cloud Datastore, Google Cloud Bigtable, and/or Google Cloud SQL information specified by and/or associated with the Google Cloud API request—for example, Google Cloud Datastore, Google Cloud Bigtable, and/or Google Cloud SQL information corresponding to and/or associated with RequestApp and/or a respective end user (e.g., the RequestApp end user with “userID” of “JaneDoe555”) specified and/or associated with content information and/or metadata of the Google Cloud API request).

463. On information and belief, the Google security mediator comprises at least one processor configured to generate a query corresponding to the request. For example, the Google Cloud Broker comprises at least one Google server automated processor configured to (e.g., adapted through specially-designed and/or programmed hardware and/or software components to automatically) generate a Google Cloud Datastore, Google Cloud Bigtable, and/or Google Cloud SQL query corresponding to the request (e.g., including, referencing, and/or associated with request content and/or metadata such as requestor identity (RequestApp); request content (e.g., *userID=JaneDoe555*); request date/time, IP address/domain, and geographic origin; etc.).

464. On information and belief, the Google security mediator comprises at least one processor configured to apply the access rules stored in the ACI to restrict access to the located requested information (“LRI”). For example, the Google Cloud Broker comprises at least one Google server automated processor configured to (e.g., adapted through specially-designed

and/or programmed hardware and/or software components to automatically), for respective LRI (e.g., respective Google Cloud Datastore record, Google Cloud Bigtable record, and/or Google Cloud SQL record determined by the Google Cloud Broker to correspond to the requested information—for example, Google Cloud API request information and/or metadata specifying and/or associated with JaneDoe555), apply the access rules stored in the ACI (e.g., the respective resource-specific account-level Google Cloud Platform and/or Google App Engine PaaS API access rules stored in the Google Cloud Platform and/or Google App Engine PaaS identity and access management ACI) to restrict access by the requesting user (e.g., RequestApp) to the LRI.

465. On information and belief, the Google security mediator comprises at least one processor configured to generate instructions to each of the POEDs storing the LRI to apply native access rules (“NARs”) of the respective POED to further restrict access to the LRI. For example, the Google Cloud Broker comprises at least one Google server automated processor configured to (e.g., adapted through specially-designed and/or programmed hardware and/or software components to automatically) generate instructions (specially-formatted HTTP requests and/or responses) to each of the POEDs storing the LRI (e.g., to each of the Google Cloud Datastore tables, Google Cloud Bigtable tables, and/or Google Cloud SQL instances external to the Google Cloud Broker storing the respective Google Cloud Datastore records, Google Cloud Bigtable records, and/or Google Cloud SQL records determined by the Google Cloud Broker to correspond to the requested information) to apply NARs of the respective POED (e.g., database-specific access rules native to the respective Google Cloud Datastore table, Google Cloud Bigtable table, or Google Cloud SQL instance external to the Google Cloud Broker) to further (e.g., in addition to access restrictions from applying ACI-stored associated access rules) restrict access to the LRI (e.g., the respective Google Cloud Datastore records, Google Cloud Bigtable records, and/or Google Cloud SQL records determined by the Google Cloud Broker to correspond to the requested information).



Application or host level access

To enable access to an instance from a Google App Engine application or an application running on a regular host, you must authorize the host application ID or IP address, respectively. You can do this by editing the instance and granting access as discussed in this page.

MySQL database access

After a connection to an instance has been negotiated, the user or application is logged in to the MySQL instance using the MySQL authorization system. After you create an instance, you must set the `root` password if you want to connect to the instance externally from locations other than Google App Engine. For more information, see [Creating users](#).

It is strongly recommended that you set a strong password for `root` and all users you create as well as create additional users to manage finer grained access to your database. See [MySQL Access Privilege System](#) for more information about managing MySQL users.

Cloud SQL Documentation: Configuring Instance Access, GOOGLE CLOUD PLATFORM, <https://cloud.google.com/sql/docs/access-control/> (accessed Oct. 19, 2015).

466. On information and belief, the Google security mediator comprises at least one processor configured to consolidate the requested information retrieved from the POEDs storing the LRI, wherein access to the LRI has not been restricted by an access rule stored in the ACI or by a NAR. For example, the Google Cloud Broker comprises at least one Google server automated processor configured to (e.g., adapted through specially-designed and/or programmed hardware and/or software components to automatically) consolidate the requested information retrieved from the POEDs storing the LRI (e.g., the respective Google Cloud Datastore tables,

Google Cloud Bigtable tables, and/or Google Cloud SQL instances external to the Google Cloud Broker storing the respective Google Cloud Datastore records, Google Cloud Bigtable records, and/or Google Cloud SQL records determined by the Google Cloud Broker to correspond to the requested information), wherein access to the LRI has not been restricted by an access rule stored in the ACI or by a NAR (e.g., for LRIs determined to be accessible to the requesting user (e.g., RequestApp) upon application of respectively applicable ACI-stored associated access rules and database-specific NARs for the respective LRI, in view of the Google Cloud API request content and metadata).

467. On information and belief, the Google security mediator comprises at least one processor configured to generate an index of POEDs storing the LRIs, wherein access to the LRI has not been restricted by an access rule stored in the ACI or by a NAR. For example, the Google Cloud Broker comprises at least one Google server automated processor configured to (e.g., adapted through specially-designed and/or programmed hardware and/or software components to automatically) generate an index of POEDs storing the LRIs, wherein access to the LRI has not been restricted by an access rule stored in the ACI or by a NAR (e.g., generate a data object specifying and/or referencing respective Google Cloud Datastore tables, Google Cloud Bigtable tables, and/or Google Cloud SQL instances external to the Google Cloud Broker storing LRIs determined to be accessible to the requesting user (e.g., RequestApp) upon application of respectively applicable ACI-stored associated access rules and database-specific NARs for the respective LRI, in view of the Google Cloud API request content and metadata).

468. On information and belief, the Google security mediator comprises a communications port configured to communicate to each of the POEDs storing the LRI, a query corresponding to the request and instructions to apply the respective NARs. For example, the Google Cloud Broker comprises a Google server physical and/or virtual network communications port configured to (e.g., adapted through specially-designed and/or programmed hardware and/or software components to automatically) communicate (e.g., transmit a specially-formatted HTTP request and/or response) to each of the POEDs storing the LRI (e.g., each of the

Google Cloud Datastore tables, Google Cloud Bigtable tables, and/or Google Cloud SQL instances external to the Google Cloud Broker storing the LRIs), a query corresponding to the request (e.g., a specially-formatted HTTP request and/or response generate a Google Cloud Datastore, Google Cloud Bigtable, and/or Google Cloud SQL query corresponding to the request (e.g., including, referencing, and/or associated with request content and/or metadata such as requestor identity (RequestApp); request content (e.g., *userID=JaneDoe555*); request date/time, IP address/domain, and geographic origin; etc.) and instructions to apply the respective NARs (e.g., a specially-formatted HTTP request and/or response configured to instruct the respective Google Cloud Datastore table, Google Cloud Bigtable table, or Google Cloud SQL instance to apply respectively-applicable database-specific access rules for information potentially responsive to the query).

469. On information and belief, the Google security mediator comprises a communications port configured to communicate to the user at least one of the consolidated index of the LRIs and the consolidation of the LRIs. For example, the Google Cloud Broker comprises a Google server physical and/or virtual network communications port configured to (e.g., adapted through specially-designed and/or programmed hardware and/or software components to automatically) communicate (e.g., via a specially-formatted HTTP response) to the user (e.g., RequestApp) at least one of the consolidated index of the LRIs and the consolidation of the LRIs.

470. On information and belief, the Google '630 Products are made, sold, and/or offered for sale by and/or on behalf of Google to entities (e.g., businesses, schools, and other organizations) and individuals throughout the United States.

471. On information and belief, the Google '630 Products are made, sold, and offered for sale by and/or on behalf of Google to entities (e.g., businesses, schools, and other organizations) and individuals located in the Eastern District of Texas.

472. On information and belief, the Google '630 Products are used by Google (e.g., by and/or on behalf of Google employees) throughout the United States.

473. On information and belief, the Google '630 Products are used by Google (e.g., by and/or on behalf of Google employees) within the Eastern District of Texas.

474. By making, using, offering for sale, and/or selling infringing products and services for managing access to protected data, including but not limited to the Google '630 Products, Google has injured St. Luke and is liable to St. Luke for directly infringing one or more claims of the '630 patent, including at least claim 16, pursuant to 35 U.S.C. § 271(a).

475. On information and belief, Google also indirectly infringes the '630 patent by actively inducing infringement under 35 USC § 271(b).

476. On information and belief, at least since service of this Complaint or shortly thereafter, Google has known of the '630 patent and has known about infringement of the '630 patent by Google itself and by third-party Google customers, end-users, developers, and/or integrators/partners of the Google '630 Products.

477. On information and belief, beginning no later than the date of service of this Complaint, Google has intentionally performed acts that induce infringement of the '630 patent by third parties (e.g., Google '630 Product customers, end-users, developers, and/or integrators/partners), knowing that these acts would induce third-party infringement of the '630 patent and/or with willful blindness to this fact.

478. For example, on information and belief, Google provides products and services (e.g., the Google '630 Products) capable of infringing one or more claims of the '630 patent, including at least claim 16.

479. For example, on information and belief, Google configures these products and services (e.g., the Google '630 Products) to infringe at least one claim of the '630 patent in normal operation by Google customers, end-users, developers, and/or integrators/partners.

480. For example, on information and belief, Google instructs and directs customers, end-users, developers, and/or integrators/partners to make and/or use the Google '630 Products in an infringing manner and/or configuration (e.g., through creation and dissemination of Google '630 Product documentation, training materials, SDKs, client libraries, and API products and

services that not only facilitate, but effectively mandate, third-party infringement of the '630 patent by Google customers, end-users, developers, and/or integrators/partners).

481. Accordingly, Google has actively induced and continues to actively induce infringement of the '630 patent by Google '630 Product customers, end-users, developers, and/or integrators/partners.

482. To the extent applicable, the requirements of 35 U.S.C. § 287(a) have been met with respect to the '630 patent.

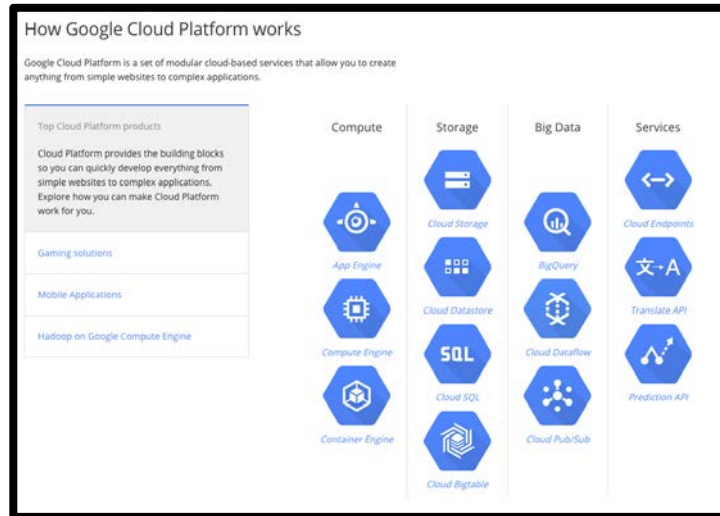
483. As a result of Google's infringement of the '630 patent, St. Luke has suffered monetary damages, and seeks recovery in an amount adequate to compensate for Google's infringement, but in no event less than a reasonable royalty for the use made of the '630 patent inventions by Google, together with interest and costs as fixed by the Court.

COUNT VII
INFRINGEMENT OF U.S. PATENT NO. 8,600,895

484. St. Luke references and incorporates by reference the preceding paragraphs of the Complaint as if fully set forth herein.

485. Google designs, makes, sells, offers to sell, imports, and/or uses in the United States products and/or services for managing access to protected data.

486. Google designs, makes, sells, offers to sell, imports, and/or uses in the United States the Google Cloud Platform, including but not limited to Google Cloud Compute products and services (e.g., Google App Engine, Google Compute Engine, Google Container Engine); Google Cloud Storage products and services (e.g., Google Cloud Storage, Google Cloud Datastore, Google Cloud SQL, Google Cloud Bigtable); Google Cloud Big Data products and services (e.g., Google Cloud BigQuery, Google Cloud Dataflow, Google Cloud Pub/Sub); and Google Cloud Services products and services (e.g., Google Cloud Endpoints, Google Translate API, Google Prediction API) (collectively, "Google Cloud").



How Google Cloud Platform Works, <https://cloud.google.com> (retrieved Oct. 19, 2015).

487. Google designs, makes, sells, offers to sell, imports, and/or uses in the United States the Google App Engine platform-as-a-service, including but not limited to Google App Engine products and services and server-side and client-side products and services running on and/or interfacing with Google App Engine products and/or services (collectively, “Google App Engine”).

488. Google designs, makes, sells, offers to sell, imports, and/or uses in the United States the Google Cloud Datastore platform, including but not limited to Google Cloud Datastore products and services and server-side and client-side products and services running on and/or interfacing with Google Cloud Datastore products and/or services (collectively, “Google Cloud Datastore”).

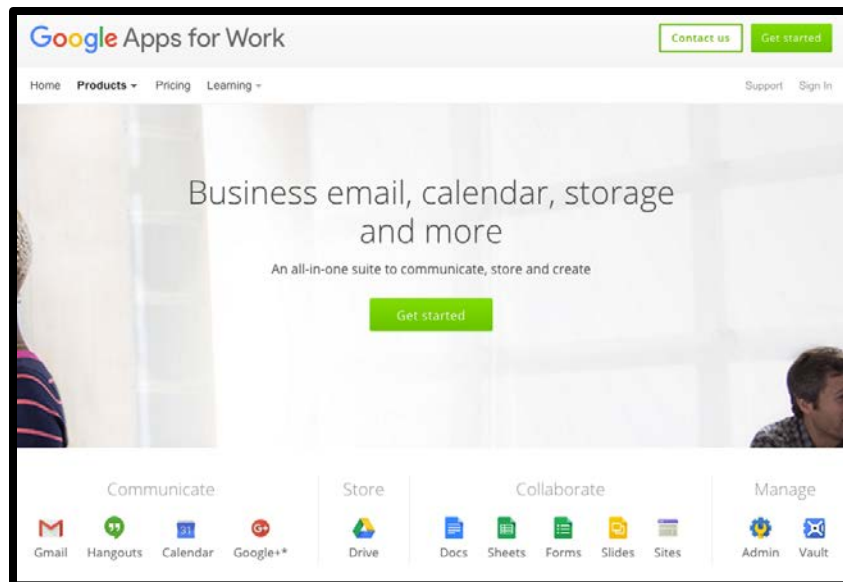
489. Google designs, makes, sells, offers to sell, imports, and/or uses in the United States the Google Cloud Bigtable platform, including but not limited to Google Cloud Bigtable products and services and server-side and client-side products and services running on and/or interfacing with Google Cloud Bigtable products and/or services (collectively, “Google Cloud Bigtable”).

490. Google designs, makes, sells, offers to sell, imports, and/or uses in the United States the Google Cloud SQL platform, including but not limited to Google Cloud SQL products

and services and server-side and client-side products and services running on and/or interfacing with Google Cloud SQL products and/or services (collectively, “Google Cloud SQL”).

491. Google designs, makes, sells, offers to sell, imports, and/or uses in the United States the Google for Work platform, including but not limited to Google Apps for Work, Education, Government, and Nonprofit products and services; Google Drive for Work products and services; and Google Apps Unlimited products and services (collectively, “Google for Work”).

492. Google designs, makes, sells, offers to sell, imports, and/or uses in the United States the Google Apps for Work, Education, Government and Nonprofit platforms, including but not limited to Gmail, Google Hangouts, Google Calendar, Google+, Google Drive, Google Docs, Google Sheets, Google Forms, Google Slides, Google Sites, Google Admin, and Google Vault server-side and client-side products and services (collectively, “Google Apps”).



Google Apps for Work, <http://www.google.com/work/apps/business/products/> (retrieved Oct. 19, 2015).

493. Google designs, makes, sells, offers to sell, imports, and/or uses in the United States Google-branded application products and services running on and/or interfacing with the Google Cloud Datastore, Google Cloud SQL, and/or Google Cloud Bigtable cloud database

products and/or services (e.g., Google Analytics, Gmail, YouTube) (collectively, “Google Cloud Applications”).

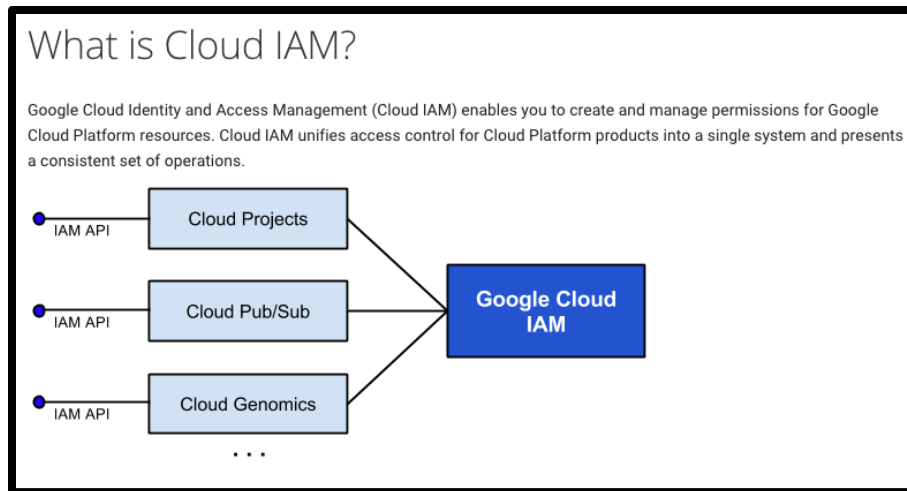
494. Google designs, makes, sells, offers to sell, imports, and/or uses in the United States Google Cloud, Google App Engine, Google Cloud Datastore, Google Cloud Bigtable, Google Cloud SQL, Google for Work, Google Apps, and Google Cloud Applications (collectively, the “Google ‘895 Products”).

495. On information and belief, Google performs (e.g., through operation of the Google ‘895 Products) at least a first method of controlling access to a plurality of records stored within a plurality of automated external databases, each record having an associated set of access rules (“ASAR”), a location identifier (“LI”), and a content identifier (“CI”) maintained in an automated centralized index (“ACI”). For example, a Google Cloud Platform and/or Google App Engine PaaS identity and access management intermediary (the “Google Cloud Broker”) performs at least a first method of controlling access to a plurality of Google Cloud database records (e.g., Google Cloud Datastore NoSQL records, Google Cloud Bigtable NoSQL records, and/or Google Cloud SQL records comprising respective pieces of cloud-synced application data) stored within a plurality of automated external databases (e.g., Google Cloud Datastore tables, Google Cloud Bigtable tables, and/or Google Cloud SQL instances), each record having an ASAR (e.g., resource-specific account-level Google Cloud Platform and/or Google App Engine PaaS API access rules), an LI (e.g., URI/URN/URL for a respective Google Cloud Datastore, Google Cloud Bigtable, or Google Cloud SQL resource), and a CI (e.g., content-related URI/URN/URL information and/or metadata for a respective Google Cloud Datastore record, Google Cloud Bigtable record, or Google Cloud SQL record) maintained in an ACI (e.g., a Google Cloud Broker identity and access management ACI).

Benefits and features

Cloud IAM includes the following features:

- **Single access control interface.** Cloud IAM provides a simple and consistent access control interface for all Cloud Platform services. You can learn one access control interface and apply that knowledge to all Cloud Platform services.
- **Resource-level access control.** You can assign roles to users to access resources at a granularity *finer than the project level*. For example, you can create a policy that assigns the *subscriber* role to a user for a particular Pub/Sub topic.
- **Flexible roles.** Prior to Cloud IAM, you could only assign owner, editor, or viewer roles to users. Cloud Platform products now expose additional flexible roles. For example, the Pub/Sub service exposes *publisher* and *subscriber* roles in addition to the owner, editor, and viewer roles.
- **UI and REST APIs.** You can create and manage Cloud IAM policies using the Google Developers Console or the Cloud IAM APIs.
- **Google account support.** Cloud IAM supports standard Google accounts. You can create Cloud IAM policies to grant permission to a [Google group](#), a [Google-hosted domain](#), a [service account](#), or specific [Google account](#) holders. You can centrally manage users and groups through the [Google Apps Admin Console](#).
- **Available free of charge.** Cloud IAM is offered at no additional charge for all Cloud Platform customers. You will be charged only for use of other Cloud Platform products. For information about the pricing of other Cloud Platform products, see the [Google Cloud Platform Pricing Calculator](#).



Cloud Identity and Access Management: What is Cloud IAM?, GOOGLE CLOUD PLATFORM, <https://cloud.google.com/iam> (accessed Oct. 19, 2015).

496. On information and belief, the first method comprises Google (e.g., through operation of the Google ‘895 Products) receiving a request, communicated from a requestor to a centralized automated security processor (“CASP”), the request containing a specified CI (“SCI”). For example, the Google Cloud Broker includes and/or constitutes a Google Cloud Broker CASP that receives, from a requesting Google Cloud-enrolled web or mobile application (“RequestApp”), a Google Cloud API request containing an SCI (e.g., Google Cloud API request information and/or metadata specifying and/or associated with a RequestApp end user—for

example, the RequestApp end user with “userID” of “JaneDoe555”). The Google Cloud API request is communicated (e.g., via HTTP over the Internet and/or a VPN) from a requestor (e.g., RequestApp) to the Google Cloud Broker CASP.

497. On information and belief, the first method comprises Google (e.g., through operation of the Google ‘895 Products) authenticating the requestor. For example, the Google Cloud Broker authenticates RequestApp by verifying Google Developer Portal-provisioned Google Cloud API credentials included in and/or referenced by the Google Cloud API request.

Users

App Engine applications can authenticate users using any one of three methods: Google Accounts, accounts on your own Google Apps domains, or OpenID identifiers. (Note that the support for OpenID is in Beta.) An application can detect whether the current user has signed in, and can redirect the user to the appropriate sign-in page to sign in or, if your app uses Google Accounts authentication, create a new account. While a user is signed in to the application, the app can access the user’s email address (or OpenID identifier if your app is using OpenID). The app can also detect whether the current user is an administrator, making it easy to implement admin-only areas of the app.

Available in [Python](#) | [Java](#) | [PHP](#) | [Go](#) [\[Back to top\]](#)

App Engine Documentation: Users Python API Overview, GOOGLE CLOUD PLATFORM, <https://cloud.google.com/appengine/docs/python/users> (accessed Oct. 19, 2015).

App Identity

The Application Identity service provides the ability for an application to identify itself, by retrieving its App ID or the hostname part of its URL. The app’s identity may be used to generate a URL or email address, or to make some run-time decision.

Many Google APIs support OAuth assertions to identify the source of a request. The App Identity API can create tokens that can be used to assert that the source of a request is the application itself. This token can be included in the HTTP headers of a call to identify the calling application. The OAuth token only works against Google systems. However you can use the underlying signing technology in the API to assert the identity of your application to other systems.

Available in [Python](#) | [Java](#) | [PHP](#) | [Go](#) [\[Back to top\]](#)

OAuth Alpha

The App Engine OAuth API uses the OAuth protocol and provides a way for your app to authenticate a user who is requesting access without asking for user credentials (username and password).

Available in [Python](#) | [Java](#) | [Go](#) [\[Back to top\]](#)

App Engine Documentation: Overview of App Engine Features, GOOGLE CLOUD PLATFORM, <https://cloud.google.com/appengine/features> (accessed Oct. 19, 2015).

498. On information and belief, the first method comprises Google (e.g., through operation of the Google ‘895 Products) querying the ACI to find entries corresponding to the SCI. For example, the Google Cloud Broker queries the ACI (e.g., Google Cloud Broker identity and access management ACI) to find entries (e.g., Google Cloud Datastore records,

Google Cloud Bigtable records, and/or Google Cloud SQL records corresponding to respective pieces of RequestApp data stored in Google Cloud Datastore tables, Google Cloud Bigtable tables, and/or Google Cloud SQL instances external to the Google Cloud Broker intermediary) corresponding to the SCI (e.g., Google Cloud API request information and/or metadata specifying and/or associated with JaneDoe555).

499. On information and belief, the first method comprises Google (e.g., through operation of the Google '895 Products), for each found entry, applying the ASAR corresponding to the LI to determine if the record stored in a respective automated external database ("AXD") of the plurality of automated external databases corresponding to the LI is accessible. For example, for each found entry (e.g., each Google Cloud Datastore record, Google Cloud Bigtable record, and/or Google Cloud SQL record determined by the Google Cloud Broker to correspond to the SCI (e.g., Google Cloud API request information and/or metadata specifying and/or associated with JaneDoe555)), the Google Cloud Broker applies the ASAR (e.g., the respective resource-specific account-level Google Cloud Platform and/or Google App Engine PaaS API access rules) corresponding to the LI (e.g., the URI/URN/URL of a respective Google Cloud Datastore, Google Cloud Bigtable, or Google Cloud SQL resource) to determine if the record stored in a respective AXD (e.g., a respective Google Cloud Datastore table, Google Cloud Bigtable table, or Google Cloud SQL instance external to the Google Cloud Broker intermediary) of the plurality of AXDs corresponding to the LI is accessible.

500. On information and belief, the first method comprises Google (e.g., through operation of the Google '895 Products), for each accessible record ("AR"), automatically communicating from the CASP to the AXD storing the AR information sufficient to determine whether the AR is releasable by the AXD storing the AR by applying a set of native access rules ("NAR") maintained by the AXD storing the AR. For example, for each AR (e.g., each Google Cloud Datastore record, Google Cloud Bigtable record, and/or Google Cloud SQL record determined by the Google Cloud Broker to correspond to the SCI and be accessible based on the ASAR), the Google Cloud Broker automatically communicates (e.g., through a specially-

formatted HTTP request between and/or among Google Cloud physical and/or virtual servers and/or data centers) from the Google Cloud Broker CASP to the AXD storing the AR (e.g., respective Google Cloud Datastore table, Google Cloud Bigtable table, or Google Cloud SQL instance external to the Google Cloud Broker intermediary storing the AR) information (e.g., request, requestor, and/or other information and/or metadata specified in and/or associated with the Google Cloud API request) sufficient to determine whether the AR is releasable by the AXD storing the AR by applying a set of NAR (e.g., database-specific access rules for a respective Google Cloud Datastore table, Google Cloud Bigtable table, or Google Cloud SQL instance) maintained by the AXD (e.g., respective Google Cloud Datastore table, Google Cloud Bigtable table, or Google Cloud SQL instance external to the Google Cloud Broker intermediary) storing the AR.

501. On information and belief, the first method comprises Google (e.g., through operation of the Google '895 Products) logically associating the releasable ARs into a linked set of releasable ARs. For example, the Google Cloud Broker logically associates the releasable ARs (e.g., the ARs determined to be releasable through application of NARs maintained by the respective AXDs storing the ARs) into a linked set of ARs.

502. On information and belief, the first method comprises Google (e.g., through operation of the Google '895 Products) communicating the linked set of releasable ARs to the requestor. For example, the Google Cloud Broker communicates (e.g., through a specially-formatted HTTP response) the linked set of releasable ARs (e.g., the logically associated set of ARs determined to be releasable upon application of NARs maintained by the respective AXDs storing the ARs) to the requestor (e.g., RequestApp).

503. On information and belief, Google performs (e.g., through operation of the Google '895 Products) at least a second method of controlling access to a plurality of records stored within a plurality of automated external databases, each record being associated with an entry maintained in an automated centralized index ("ACI"), comprising an associated set of access rules ("ASAR"), a location identifier ("LI"), and a content identifier ("CI"). For example,

the “Google Cloud Broker” performs at least a second method of controlling access to a plurality of Google Cloud database records (e.g., Google Cloud Datastore NoSQL records, Google Cloud Bigtable NoSQL records, and/or Google Cloud SQL records comprising respective pieces of cloud-synced application data) stored within a plurality of automated external databases (e.g., Google Cloud Datastore tables, Google Cloud Bigtable tables, and/or Google Cloud SQL instances), each record being associated with an entry maintained in an ACI (e.g., a Google Cloud Broker identity and access management ACI), comprising an ASAR (e.g., resource-specific account-level Google Cloud Platform and/or Google App Engine PaaS API access rules), an LI (e.g., URI/URN/URL for a respective Google Cloud Datastore, Google Cloud Bigtable, or Google Cloud SQL resource), and a CI (e.g., content-related URI/URN/URL data and/or metadata for a respective Google Cloud Datastore record, Google Cloud Bigtable record, or Google Cloud SQL record).

504. On information and belief, the second method comprises Google (e.g., through operation of the Google ‘895 Products) receiving a request, communicated from a requestor to a centralized automated security processor (“CASP”), the request containing a specified CI (“SCI”). For example, the Google Cloud Broker includes and/or constitutes a Google Cloud Broker CASP that receives, from a requesting Google Cloud-enrolled web or mobile application (“RequestApp”), a Google Cloud API request containing an SCI (e.g., Google Cloud API request information and/or metadata specifying and/or associated with a RequestApp end user—for example, the RequestApp end user with “userID” of “JaneDoe555”). The Google Cloud API request is communicated (e.g., via HTTP over the Internet and/or a VPN) from a requestor (e.g., RequestApp) to the Google Cloud Broker CASP.

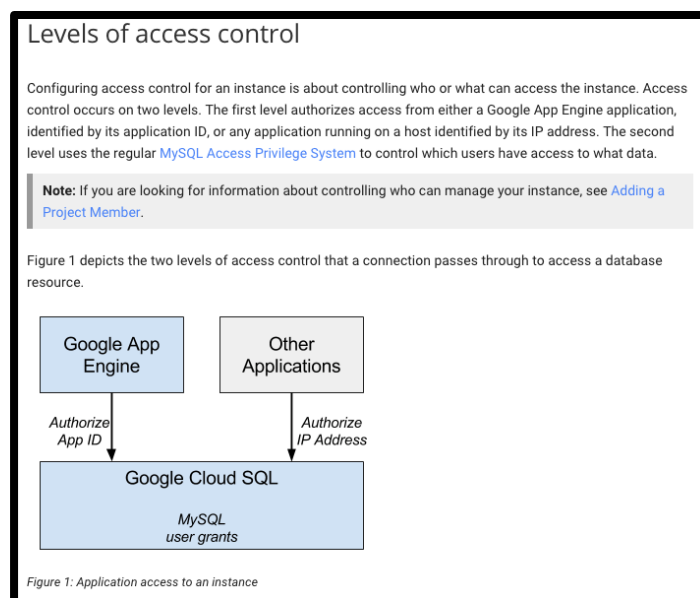
505. On information and belief, the second method comprises Google (e.g., through operation of the Google ‘895 Products) authenticating the requestor. For example, the Google Cloud Broker authenticates RequestApp by verifying Google Developer Portal-provisioned Google Cloud API credentials included in and/or referenced by the Google Cloud API request.

506. On information and belief, the second method comprises Google (e.g., through operation of the Google ‘895 Products) querying the ACI to find entries corresponding to the SCI. For example, the Google Cloud Broker queries the ACI (e.g., Google Cloud Broker identity and access management ACI) to find entries (e.g., Google Cloud Datastore records, Google Cloud Bigtable records, and/or Google Cloud SQL records corresponding to respective pieces of RequestApp data stored in Google Cloud Datastore tables, Google Cloud Bigtable tables, and/or Google Cloud SQL instances external to the Google Cloud Broker intermediary) corresponding to the SCI (e.g., Google Cloud API request information and/or metadata specifying and/or associated with JaneDoe555).

507. On information and belief, the second method comprises Google (e.g., through operation of the Google ‘895 Products), for each found entry, applying the ASAR corresponding to the LI to determine if the record stored in a respective automated external database (“AXD”) of the plurality of automated external databases corresponding to the LI is accessible. For example, for each found entry (e.g., each Google Cloud Datastore record, Google Cloud Bigtable record, and/or Google Cloud SQL record determined by the Google Cloud Broker to correspond to the SCI (e.g., Google Cloud API request information and/or metadata specifying and/or associated with JaneDoe555)), the Google Cloud Broker applies the ASAR (e.g., the respective resource-specific account-level Google Cloud Platform and/or Google App Engine PaaS API access rules) corresponding to the LI (e.g., the URI/URN/URL of a respective Google Cloud Datastore, Google Cloud Bigtable, or Google Cloud SQL resource) to determine if the record stored in a respective AXD (e.g., a respective Google Cloud Datastore table, Google Cloud Bigtable table, or Google Cloud SQL instance external to the Google Cloud Broker intermediary) of the plurality of AXDs corresponding to the LI is accessible.

508. On information and belief, the second method comprises Google (e.g., through operation of the Google ‘895 Products), for each accessible record (“AR”), communicating, from the CASP to the AXD storing the AR, information sufficient for the AXD storing the AR to apply a set of native access rules (“NAR”) it maintains to determine whether AR is releasable by

the AXD storing the AR. For example, for each AR (e.g., each Google Cloud Datastore record, Google Cloud Bigtable record, and/or Google Cloud SQL record determined by the Google Cloud Broker to correspond to the SCI and be accessible based on the ASAR), the Google Cloud Broker communicates (e.g., through a specially-formatted HTTP request between and/or among Google Cloud physical and/or virtual servers and/or data centers) from the Google Cloud Broker CASP to the AXD storing the AR (e.g., respective Google Cloud Datastore table, Google Cloud Bigtable table, or Google Cloud SQL instance external to the Google Cloud Broker intermediary storing the AR) information (e.g., request, requestor, and/or other information and/or metadata specified in and/or associated with the Google Cloud API request) sufficient for the AXD storing the AR (e.g., respective Google Cloud Datastore table, Google Cloud Bigtable table, or Google Cloud SQL instance external to the Google Cloud Broker intermediary storing the AR) to apply a set of NAR (e.g., database-specific access rules for a respective Google Cloud Datastore table, Google Cloud Bigtable table, or Google Cloud SQL instance) it maintains to determine whether the AR is releasable by the AXD storing the AR (e.g., respective Google Cloud Datastore table, Google Cloud Bigtable table, or Google Cloud SQL instance external to the Google Cloud Broker intermediary storing the AR).



Cloud SQL Documentation: Configuring Instance Access, GOOGLE CLOUD PLATFORM, <https://cloud.google.com/sql/docs/access-control/> (accessed Oct. 19, 2015).

509. On information and belief, the second method comprises Google (e.g., through operation of the Google ‘895 Products) generating, by the CASP, an index of releasable ARs. For example, the Google Cloud Broker CASP generates an index of releasable ARs (e.g., a data object logically identifying and/or referencing the ARs determined to be releasable through application of NARs maintained by the respective AXDs storing the ARs).

510. On information and belief, the second method comprises Google (e.g., through operation of the Google ‘895 Products) communicating the generated index to the requestor. For example, the Google Cloud Broker communicates to the requestor (e.g., via a specially-formatted HTTP response communicated over one or more computer networks from the Google Cloud Broker to RequestApp) the generated index (e.g., the data object generated by the Google Cloud Broker CASP logically identifying and/or referencing the ARs determined to be releasable through application of NARs maintained by the respective AXDs storing the ARs).

511. On information and belief, the second method comprises Google (e.g., through operation of the Google ‘895 Products) receiving a communication from the requestor containing a selection of at least one releasable AR in the generated index. For example, the Google Cloud Broker receives a communication from the requestor (e.g., a specially-formatted Google API HTTP request communicated over one or more computer networks from RequestApp to the Google Cloud Broker) containing a selection of at least one releasable AR in the generated index (e.g., specifying and/or referencing at least one respective AR (e.g., Google Cloud Datastore record, Google Cloud Bigtable record, and/or Google Cloud SQL record) from the releasable ARs logically identified and/or referenced in the releasable AR index generated by the Google Cloud Broker CASP).

512. On information and belief, the second method comprises Google (e.g., through operation of the Google ‘895 Products) linking the selected releasable ARs into a logically associated set of releasable ARs (“LAS”). For example, the Google Cloud Broker logically associates into a LAS the selected releasable ARs (e.g., the at least one releasable AR determined

by the Google Cloud Broker to be requestor-selected based on the content of the communication (e.g., specially-formatted Google API HTTP request) received from RequestApp).

513. On information and belief, the second method comprises Google (e.g., through operation of the Google '895 Products) communicating the LAS to the requestor. For example, the Google Cloud Broker communicates the LAS to the requestor (e.g., via a specially-formatted HTTP response communicated over one or more computer networks from the Google Cloud Broker to RequestApp).

514. On information and belief, Google designs, makes, sells, offers to sell, imports and/or uses in the United States at least one apparatus for controlling access to a plurality of records stored within a plurality of automated external databases ("AXES"). For example, the Google '895 Products comprise at least one apparatus (e.g., a Google Cloud Broker physical and/or virtual appliance) for (e.g., configured through specially designed and/or programmed hardware and/or software components for) controlling access to a plurality of Google Cloud database records (e.g., Google Cloud Datastore records, Google Cloud Bigtable records, and/or Google Cloud SQL records) provided within a plurality of AXES (e.g., Google Cloud Datastore NoSQL tables; Google Cloud Bigtable NoSQL tables; and/or Google Cloud SQL instances external to the Google Cloud Broker physical and/or virtual appliance).

515. On information and belief, the at least one apparatus (e.g., the Google Cloud Broker physical and/or virtual appliance) comprises an automated centralized index ("ACI"), stored in a memory, configured to store an entry for each record consisting of a location identifier ("LI"), an associated set of access rules ("ASAR"), and a content identifier ("CI"). For example, the Google Cloud Broker physical and/or virtual appliance comprises a Google Cloud Broker ACI, stored in a Google Cloud server memory, configured to store an entry for each record (e.g., Google Cloud Datastore record, Google Cloud Bigtable record, and/or Google Cloud SQL record) consisting of an LI (e.g., URI/URN/URL for a respective Google Cloud Datastore, Google Cloud Bigtable, or Google Cloud SQL resource), an ASAR (e.g., resource-specific account-level Google Cloud Platform and/or Google App Engine PaaS API access

rules), and a CI (e.g., content-related URI/URN/URL data and/or metadata for a respective Google Cloud Datastore record, Google Cloud Bigtable record, or Google Cloud SQL record).

516. On information and belief, the at least one apparatus (e.g., the Google Cloud Broker physical and/or virtual appliance) comprises an input port configured to receive a request from a requestor for access to one or more records stored in the plurality of AXES, wherein the request specifies a CI with which to query the ACI. For example, the Google Cloud Broker physical and/or virtual appliance comprises an input port (e.g., a Google server physical and/or virtual computer network communications port) configured to (e.g., adapted through specially-designed and/or programmed hardware and/or software components to automatically) receive a Google Cloud API request from a requestor (e.g., “RequestApp,” a Google Cloud-enrolled web or mobile application) for access to one or more records (e.g., one or more Google Cloud Datastore records, Google Cloud Bigtable records, and/or Google Cloud SQL records comprising respective pieces of application (e.g., RequestApp) data) stored in the plurality of AXES (e.g., plurality of Google Cloud Datastore tables, Google Cloud Bigtable tables, and/or Google Cloud SQL instances external to the Google Cloud Broker physical and/or virtual appliance), wherein the Google Cloud API request specifies a CI (e.g., Google Cloud API request information and/or metadata specifying and/or associated with a RequestApp end user—for example, the RequestApp end user with “userID” of “JaneDoe555”) with which to query the ACI (e.g., the Google Cloud Broker ACI).

517. On information and belief, the at least one apparatus (e.g., the Google Cloud Broker physical and/or virtual appliance) comprises at least one processor configured to generate a query based on the specified CI (“SCI”). For example, the Google Cloud Broker physical and/or virtual appliance comprises at least one Google server automated processor configured to (e.g., adapted through specially-designed and/or programmed hardware and/or software components to automatically) generate (e.g., by lexically parsing and, if necessary, translating, a received Google Cloud API request) a Google Cloud Broker identity and access management ACI query corresponding to the SCI (e.g., the Google Cloud API request information and/or

metadata specifying and/or associated with a RequestApp end user—for example, the RequestApp end user with “userID” of “JaneDoe555”).

518. On information and belief, the at least one apparatus (e.g., the Google Cloud Broker physical and/or virtual appliance) comprises at least one processor configured to find entries in the ACI containing the SCI. For example, the Google Cloud Broker physical and/or virtual appliance comprises at least one Google server automated processor configured to (e.g., adapted through specially-designed and/or programmed hardware and/or software components to automatically) find entries in the Google Cloud Broker identity and access management ACI containing the SCI (e.g., the Google Cloud API request information and/or metadata specifying and/or associated with a RequestApp end user—for example, the RequestApp end user with “userID” of “JaneDoe555”).

519. On information and belief, the at least one apparatus (e.g., the Google Cloud Broker physical and/or virtual appliance) comprises at least one processor (e.g., a Google server automated processor) configured to, for each found entry, apply the ASAR corresponding to the LI to determine if the record stored in a respective one of the AXES corresponding to the LI is accessible. For example, the Google Cloud Broker physical and/or virtual appliance comprises at least one Google server automated processor configured to (e.g., adapted through specially-designed and/or programmed hardware and/or software components to automatically), for each found entry (e.g., each Google Cloud Datastore record, Google Cloud Bigtable record, and/or Google Cloud SQL record determined by the Google Cloud Broker to correspond to the SCI (e.g., Google Cloud API request information and/or metadata specifying and/or associated with JaneDoe555)), apply the ASAR (e.g., the respective resource-specific account-level Google Cloud Platform and/or Google App Engine PaaS API access rules) corresponding to the LI (e.g., the URI/URN/URL of a respective Google Cloud Datastore, Google Cloud Bigtable, or Google Cloud SQL resource) to determine if the record stored in a respective AXD (e.g., a respective Google Cloud Datastore table, Google Cloud Bigtable table, or Google Cloud SQL instance

external to the Google Cloud Broker intermediary) of the plurality of AXDs corresponding to the LI is accessible.

520. On information and belief, the at least one apparatus (e.g., the Google Cloud Broker physical and/or virtual appliance) comprises at least one processor configured to generate a communication, for communication to the respective one of the AXES storing an accessible record (“AR”), wherein the communication contains information sufficient for the respective one of the AXES storing the AR to apply a set of native access rules (“NAR”) it maintains to determine if the AR is releasable. For example, the Google Cloud Broker physical and/or virtual appliance comprises at least one Google server automated processor configured to (e.g., adapted through specially-designed and/or programmed hardware and/or software components to automatically), generate a communication (e.g., a specially-formatted HTTP request between and/or among Google Cloud physical and/or virtual servers and/or data centers), for communication to the respective one of the AXES (e.g., the respective Google Cloud Datastore table, Google Cloud Bigtable table, or Google Cloud SQL instance external to the Google Cloud Broker physical and/or virtual appliance) storing an AR (e.g., a respective Google Cloud Datastore record, Google Cloud Bigtable record, and/or Google Cloud SQL record determined by the Google Cloud Broker to correspond to the SCI and be accessible based on the ASAR), wherein the communication (e.g., specially-formatted HTTP request between and/or among Google Cloud physical and/or virtual servers and/or data centers) contains information (e.g., request, requestor, and/or other information and/or metadata specified in and/or associated with the Google Cloud API request) sufficient for respective one of the AXES (e.g., respective Google Cloud Datastore table, Google Cloud Bigtable table, or Google Cloud SQL instance external to the Google Cloud Broker intermediary) to apply a set of NAR (e.g., database-specific access rules for a respective Google Cloud Datastore table, Google Cloud Bigtable table, or Google Cloud SQL instance) it maintains to determine if the AR is releasable.

521. On information and belief, the at least one apparatus (e.g., the Google Cloud Broker physical and/or virtual appliance) comprises at least one processor configured to form a

linked set of releasable ARs by logically associating the releasable ARs. For example, the Google Cloud Broker physical and/or virtual appliance comprises at least one Google server automated processor configured to (e.g., adapted through specially-designed and/or programmed hardware and/or software components to automatically) form a linked set of releasable ARs (e.g., a logically associated set of ARs determined to be releasable through application of NARs maintained by the respective AXDs storing the ARs) by logically associating the releasable ARs.

522. On information and belief, the at least one apparatus (e.g., the Google Cloud Broker physical and/or virtual appliance) comprises at least one processor configured to generate a communication containing the linked set of releasable ARs. For example, the Google Cloud Broker physical and/or virtual appliance comprises at least one Google server automated processor configured to (e.g., adapted through specially-designed and/or programmed hardware and/or software components to automatically) generate a communication (e.g., generate a specially-formatted HTTP response for transmission across one or more computer networks from the Google Cloud Broker to the requestor (e.g., RequestApp)) containing the linked set of releasable ARs (e.g., the logically associated set of ARs determined to be releasable through application of NARs maintained by the respective AXDs storing the ARs).

523. On information and belief, the at least one apparatus (e.g., the Google Cloud Broker physical and/or virtual appliance) comprises at least one communications port configured to communicate the generated communication to the respective one of the AXES storing the ARs. For example, the Google Cloud Broker physical and/or virtual appliance comprises at least one Google server physical and/or virtual network communications port configured to (e.g., adapted through specially-designed and/or programmed hardware and/or software components to automatically) communicate, to each respective Google Cloud Datastore table, Google Cloud Bigtable table, or Google Cloud SQL instance external to the Google Cloud Broker physical and/or virtual appliance storing an AR, the specially-formatted HTTP request between and/or among Google Cloud physical and/or virtual servers and/or data centers containing information sufficient for the respective Google Cloud Datastore table, Google Cloud Bigtable table, or

Google Cloud SQL instance storing an AR to apply a set of NAR it maintains to determine if the AR is releasable.

524. On information and belief, the at least one apparatus (e.g., the Google Cloud Broker physical and/or virtual appliance) comprises at least one communications port configured to communicate the linked set of releasable ARs. For example, the Google Cloud Broker physical and/or virtual appliance comprises at least one Google server physical and/or virtual network communications port configured to (e.g., adapted through specially-designed and/or programmed hardware and/or software components to automatically) communicate the linked set of releasable ARs (e.g., transmit via one or more computer networks from the Google Cloud Broker to the requestor (e.g., RequestApp) the specially-formatted HTTP response containing the logically associated set of ARs determined to be releasable through application of NARs maintained by the respective AXDs storing the ARs).

525. On information and belief, the Google '895 Products are made, sold, and/or offered for sale by and/or on behalf of Google to entities (e.g., businesses, schools, and other organizations) and individuals throughout the United States.

526. On information and belief, the Google '895 Products are made, sold, and offered for sale by and/or on behalf of Google to entities (e.g., businesses, schools, and other organizations) and individuals located in the Eastern District of Texas.

527. On information and belief, the Google '895 Products are used by Google (e.g., by and/or on behalf of Google employees) throughout the United States.

528. On information and belief, the Google '895 Products are used by Google (e.g., by and/or on behalf of Google employees) within the Eastern District of Texas.

529. On information and belief, Google has directly infringed and continues to directly infringe the '895 patent by, among other things, directly performing the methods of claims 1 and 8 of the patent through operation of at least the Google '895 Products.

530. By making, using, offering for sale, and/or selling infringing products and services for managing access to protected data, including but not limited to the Google '895

Products, Google has injured St. Luke and is liable to St. Luke for directly infringing one or more claims of the '895 patent, including at least claims 1, 8, and 16, pursuant to 35 U.S.C. § 271(a).

531. On information and belief, Google also indirectly infringes the '895 patent by actively inducing infringement under 35 USC § 271(b).

532. On information and belief, at least since service of this Complaint or shortly thereafter, Google has known of the '895 patent and has known about infringement of the '895 patent by Google itself and by third-party Google customers, end-users, developers, and/or integrators/partners of the Google '895 Products.

533. On information and belief, beginning no later than the date of service of this Complaint, Google has intentionally performed acts that induce infringement of the '895 patent by third parties (e.g., Google '895 Product customers, end-users, developers, and/or integrators/partners), knowing that these acts would induce third-party infringement of the '895 patent and/or with willful blindness to this fact.

534. For example, on information and belief, Google provides products and services (e.g., the Google '895 Products) capable of infringing one or more claims of the '895 patent, including at least claims 1, 8, and 16.

535. For example, on information and belief, Google configures these products and services (e.g., the Google '895 Products) to infringe at least one claim of the '895 patent in normal operation by Google customers, end-users, developers, and/or integrators/partners.

536. For example, on information and belief, Google instructs and directs customers, end-users, developers, and/or integrators/partners to make and/or use the Google '895 Products in an infringing manner and/or configuration (e.g., through creation and dissemination of Google '895 Product documentation, training materials, SDKs, client libraries, and API products and services that not only facilitate, but effectively mandate, third-party infringement of the '895 patent by Google customers, end-users, developers, and/or integrators/partners).

537. Accordingly, Google has actively induced and continues to actively induce infringement of the '895 patent by Google '895 Product customers, end-users, developers, and/or integrators/partners.

538. To the extent applicable, the requirements of 35 U.S.C. § 287(a) have been met with respect to the '895 patent.

539. As a result of Google's infringement of the '895 patent, St. Luke has suffered monetary damages, and seeks recovery in an amount adequate to compensate for Google's infringement, but in no event less than a reasonable royalty for the use made of the '895 patent inventions by Google, together with interest and costs as fixed by the Court.

PRAYER FOR RELIEF

WHEREFORE, Plaintiff St. Luke respectfully requests that this Court enter:

- A. A judgment in favor of Plaintiff St. Luke that Google has infringed, either literally and/or under the doctrine of equivalents, the '237 patent, the '017 patent, the '377 patent, the '368 patent, the '941 patent, the '630 patent, and/or the '895 patent;
- B. An award of damages resulting from Google's acts of infringement in accordance with 35 U.S.C. § 284;
- C. A judgment and order requiring Google to provide accountings and to pay supplemental damages to St. Luke, including, without limitation, prejudgment and post-judgment interest; and
- D. Any and all other relief to which St. Luke may show itself to be entitled.

JURY TRIAL DEMANDED

Pursuant to Rule 38 of the Federal Rules of Civil Procedure, St. Luke requests a trial by jury of any issues so triable by right.

Dated: October 20, 2015

Respectfully submitted,

/s/ Elizabeth L. DeRieux
Elizabeth L. DeRieux (TX Bar No. 05770585)
D. Jeffrey Rambin (TX Bar No. 00791478)
CAPSHAW DERIEUX, LLP
114 E. Commerce Ave.
Gladewater, Texas 75647
Telephone: 903-236-9800
Facsimile: 903-236-8787
E-mail: ederieux@capshawlaw.com
E-mail: jrambin@capshawlaw.com

OF COUNSEL:

Brian J. Dunne (CA SB No. 275689)
OLAVI DUNNE LLP
816 Congress Ave., Ste. 1620
Austin, Texas 78701
Telephone: 512-717-4485
Facsimile: 512-717-4495
E-mail: bdunne@olavidunne.com

Dorian S. Berger (CA SB No. 264424)
Daniel P. Hipskind (CA SB No. 266763)
OLAVI DUNNE LLP
1880 Century Park East, Ste. 815
Los Angeles, CA 90067
Telephone: 213-516-7900
Facsimile: 213-516-7910
E-mail: dberger@olavidunne.com
E-mail: dhipskind@olavidunne.com

Attorneys for St. Luke Technologies, LLC