

Exhibit 7

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

R2 Solutions LLC,

Plaintiff,

v.

FedEx Corporate Services, Inc.,

Defendant.

Civil Action No. 4:21-cv-00940

Jury Trial Demanded

**DECLARATION OF BILL DAVIS IN SUPPORT OF PLAINTIFF'S FIRST AMENDED
COMPLAINT FOR PATENT INFRINGEMENT**

I, Bill Davis, declare as follows:

1. I am currently a Partner and Lead Engineer at Syscall 7. My work focuses on computer information security, distributed analytical systems for network detection and response platforms, and enterprise dataflow management for indexing and search. I have over 20 years of experience building software for defense and industry applications.

2. This declaration is based on my own personal knowledge. I make this declaration in support of Plaintiff's First Amended Complaint for Patent Infringement.

3. I have been retained on behalf of Plaintiff R2 Solutions LLC ("R2") to provide a declaration in connection with the above-captioned litigation. R2 has asked me to review U.S. Patent No. 8,190,610 ("the '610 patent"), U.S. Patent No. 8,341,157 ("the '157 patent"), and U.S. Patent No. 7,698,329 (collectively "the patents-in-suit" or "asserted patents") and provide opinions on the nature of the claimed inventions (e.g., novel and improved MapReduce architecture as to the '610 patent, and novel search engine technologies as to the '157 and '329

patents), the state of the prior art as of the effective filing date of each patent, any improvements to the prior art provided by the claimed inventions of each patent and disclosed in the specification, and the benefits associated with the claimed inventions. I have knowledge of the facts set forth in this declaration based on my review of the patents-in-suit and my technical expertise.

BACKGROUND AND QUALIFICATIONS

4. In 1997, I received a Bachelor of Science with a double major in Mathematics and Computer Science from the University of Maryland, graduating as part of the Gemstone Honors Program. In 2006, I received my Master of Science in Computer Science from Johns Hopkins University. Since 2011, I have been a member of the Association for Computing Machinery.

5. I have been involved in research and technology in the area of big data analytics since 2006, building cutting-edge analytic systems for defense intelligence analysts. This included developing and deploying many classified Hadoop/MapReduce-based analytics. My experience involves designing and implementing big data analytics systems using various languages like Java, Scala, and Python, and I have also designed and implemented a number of MapReduce abstraction layers utilizing, for example, Cascading, Hive, and Pig, and a number of MapReduce alternatives such as Trino and Apache Spark. Analytics systems I have built are capable of running software across databases holding billions of events. Most recently I have been responsible for the design and implementation of an analytic architecture which operates over computer network metadata. The analytics are designed to detect and classify anomalous or malicious network activity without the use of signature-based detections. These systems rely heavily on Kafka, Hadoop HDFS, Postgres, and Apache Spark for scalable and efficient

operations. I designed and coded a framework and testing architecture to migrate these workflows to Kubernetes-based cluster deployments.

6. I have additionally been the technical lead on small teams of 5 to 10 engineers in the area of search engine technologies since 2008. I designed and built search engines and data ingest pipelines to support large scale collections of intelligence reports for the purposes of user free text and structured search queries. I have built both backend processing systems and front-end user facing applications, which included designing custom query languages to facilitate user queries. I have authored several internal papers related to graph search engines as applied to document search and network community analytics.

7. I have developed applications for the purposes of reverse engineering and software analysis. For example, since 2012, I have built, run and maintained disassembler.io, which is the leading web application for reverse engineering binaries for over 50 CPU architectures.

8. From 2018 to the present, I have also assisted multiple companies transition their enterprise Hadoop/Spark based workflows from on-premise data center deployments to modern cloud service architectures utilizing both AWS and Microsoft Azure.

9. During 2009 to 2014, I developed and deployed multiple classified information systems utilizing Java/Groovy/MySQL/MongoDB and Accumulo Column Storage for the Department of Defense.

10. As part of earlier consulting work, I developed structured search tools which are capable of indexing and storing key fields from a corpus of tens of millions of semi-structured documents. These tools relied on a suite of open-source and commercial software like ElasticSearch, Oracle, and Apache Solr to provide extensive search capabilities for data analysts.

11. My curriculum vitae, which is attached as Appendix A also includes list of all other cases in which, during at least the previous 4 years, the I have served as an expert.

LEGAL STANDARDS

12. I am not an attorney or patent agent. Accordingly, I have relied upon the legal standards provided to me by counsel for R2. These legal principles are outlined below. In forming my opinions herein, I applied this framework.

13. I understand that patent eligibility is defined by statute, including 35 U.S.C. §101, with definitions provided by 35 U.S.C. §100. I understand that, generally, patent claims covering any new and useful process, machine, manufacture, or composition of matter, or any new and useful improvements thereof, are patent eligible. I understand that there are judicial exceptions to this rule including laws of nature, natural phenomena, and abstract ideas. It is my understanding that issued patent claims are presumed patent eligible and that it is a challenger's burden to demonstrate otherwise by clear and convincing evidence.

14. When determining whether or not a patent claim is patent eligible, I understand that courts look to the language of the statute itself, and then if they seek to apply a judicial exception to patent eligibility, they apply a two-step analysis. First, the courts determine whether the claims at issue are "directed to" an abstract idea. I understand that as part of this step the claim limitations should be fairly characterized and not overly simplified.

15. Second, if the claim is directed to an abstract idea, I understand the next step is to determine whether the claims nevertheless recite an "inventive concept," such as an unconventional element or combination of elements, which is sufficient to ensure that the patent amounts to significantly more than a patent on the abstract idea. I understand that this second step is a question of fact that involves determining the state of the prior art at the time of

invention, as well as understanding precisely what, at the time of the invention, went beyond the prior art (such that aspects of the subject invention were not considered “well-known, routine, and conventional”).

16. I further understand that implementing a well-known technique with particular devices in a specific combination can be inventive. Further, it is my understanding that claims that use an environment, such as a computer, mobile phone, etc., to do significantly more than simply carry out an abstract idea are patent eligible.

17. I further understand that one factor that courts consider in determining whether patent claims are patent eligible is whether the invention improves the operation of a device and/or provides a technological solution to a technological problem.

18. I further understand that patent eligibility is different than, and distinct from, a determination of novelty under 35 U.S.C. § 102 and non-obviousness under 35 U.S.C. § 103. Those are separate inquiries.

PATENT ANALYSIS

19. I understand that R2 has asserted claims 1-5 and 17-21 of the '610 patent; claims 1-5 and 7-10 of the '157 patent; and claims 1, 4-5, 8, and 11-12 of the '329 patent. I believe that these claims are directed to patent eligible subject matter. I also understand that this Court previously construed terms of these patents in *R2 Sols. LLC v. Walmart Inc.*, No. 4:21-cv-00091, Dkt. 54 (E.D. Tex. Jan. 4, 2022).

20. In my opinion, the inventions recited in the claims of the asserted patents are directed to specific improvements in computer capabilities in the fields of distributed database processing (with respect to the '610 patent) and search engine design and application (with respect to the '157 and '329 patents). It is further my opinion that the asserted claims define

technical enhancements to the ways computers operated at the times of the inventions, and that each claim of the asserted patents recites unconventional combinations that were previously unheard of to improve upon technology existing in the patents' respective fields. In my view, the asserted claims each bring together numerous unconventional concepts and features and set forth solutions that are necessarily rooted in computer technology—indeed, the specific computer technologies covered by the asserted claims overcome problems arising solely in the realm of computer-implemented distributed database processing and search engine operation (as applicable), as opposed to problems arising out of, e.g., human interactions. It is my understanding that computerized inventions such as these constitute patentable subject matter.

21. I address each of these patents below, in turn.

I. '610 Patent


22. In my opinion, the inventions claimed in the '610 patent constitute specific technological advancements in distributed database processing. These advancements address and overcome significant problems with prior technologies known in 2006. The '610 patent claims a novel MapReduce architecture that can be applied to affect mapping and reducing across diverse data sets, including data sets comprising different schema. Providing such a MapReduce architecture adapted to receive and process non-uniform input data greatly expanded the applicability of MapReduce to dramatically augment big data analytics. Such improvement was significant at the time of the '610 patent, enhancing the overall capabilities of computer functionality as it relates to distributed database processing.

A. Introduction to MapReduce

23. As laid out in the specification, the concept of MapReduce essentially consists of a “‘map’ function [that] maps key-value pairs to new (intermediate) key-value pairs,” and a

“‘reduce’ function [that] represents all mapped (intermediate) key-value pairs sharing the same key to a single key-value pair or a list of values.” ’610 patent at 1:17-20. At its core, MapReduce is designed to facilitate the generation and/or processing of large data sets via parallel processing.

24. In a simple example, MapReduce can be used to perform a word count on a collection of text:




<https://medium.com/edureka/mapreduce-tutorial-3d9535ddbe7c>.

25. In the above example, the input data consists of a collection of words. Applying traditional MapReduce in the “Splitting” step, the input data is split and supplied to different nodes (e.g., other processors that are part of the MapReduce architecture). This splits the processing load amongst several processors, which is much more important with larger data sets as compared to this example. Next, in the “Mapping” step, a list of key-value pairs is created. A key-value pair works much like a run-of-the-mill dictionary, i.e., the words are the “keys” and the definitions are the “values.” In the example above, the mapping function generates a list of key-value pairs where the words are the “keys” and the values are “1” to indicate that each word, as a discrete item, occurs one time. Obviously, some of the words occur multiple times, and that will be accounted for later in the process. At the “Shuffling” step, sorting occurs such that all

key-value pairs having the same key are grouped together to facilitate passing to the same reducers (recall that reducers reduce based on a common key). Thus, at the “Reducing” step, the reducers can determine how many occurrences of that particular key exist for the received data and can generate another key-value pair, accordingly. Ultimately, the output is a list of all of the key-value pairs from each reducer that represents the number of times the word appears in the input text. This is a very simple application of traditional MapReduce.

26. MapReduce can also be used to accomplish other tasks, such as “joins” of data and/or lists. For example:



<https://courses.cs.washington.edu/courses/cse490h/08au/lectures/MapReduceDesignPatterns-UW2.pdf> at Slide 8.


Here, the input data is already provided as a list of key-value pairs. In this example, the input data is a list of features in a graph (e.g., roads and intersections) that are numbered; the numbers are the

keys and the feature types are the values. In the mapping phase, new keys are assigned that correspond to the intersection that the particular feature participates in (e.g., roads 1, 2, and 5, and intersection 3, all are involved with intersection 3). Thus, the key in the mapping phase is the intersection ID, and the value is the key-value pair corresponding to the graph feature. The values are shuffled and passed to the reducers based on their common key. Then, the reducers “reduce” the data to yield the output data, i.e., the reducers can consolidate all of the features pertaining to a particular common key, such that all of the features participating in intersection 3 fall with intersection 3, and all of the features participating in intersection 6 fall with intersection 6.

27. As is apparent from the above examples, MapReduce offers a way to take large, unwieldy datasets and extricate quantifiable and useful information algorithmically. As discussed further below, however, as of 2006 significant limitations existed relative to MapReduce that limited its usefulness.

B. Shortcomings of Prior Technologies


28. The '610 patent discloses shortcomings in the prior art and then explains, in detail, the technical way the claimed inventions resolve or overcome those shortcomings. The specification discusses that “conventional MapReduce implementations do not have facility to efficiently process data from heterogeneous sources” and that “it is impractical to perform joins over two relational tables that have different schemas.” '610 patent at 3:9-20. “Schema” can refer to, e.g., “fields” in a particular collection of data. For example, following along with the simple example in the '610 specification at FIG. 3, the “All Emp” table (302) and “All Dept” table (304) are different “schema,” because the “All Emp” table has the fields of “DeptID” and “EmpName,” while the “All Dept” table has the fields of “DeptID” and “EmpName”:

**Fig. 3**

'610 patent at FIG. 3; 3:19-20.

29. With respect to this simple example, the specification explains that it would be inefficient for traditional MapReduce to be implemented to join the “All Emp” table (302) and “All Dept” table (304) to achieve the “All Emp and Dept” table (306):

30. The impracticality of applying traditional MapReduce to join data of different schema is apparent from extrapolation of the traditional MapReduce examples explained above. For example, applying traditional MapReduce to the “All Emp” and “All Dept” tables together could look something like the following:



31. In this example, the “key” for the mapping phase can be considered to be the “DeptID, name” followed by a value of “1” (such as in the word counting example above).¹ As can be seen here, even though there is a key-in-common amongst the two “All Emp” and “All Dept” data sets at the input, there is no “joining of the data in any usable way. Instead, the output is a jumbled list that shows both employees and department names for a single key instead of indicating, e.g., DeptID, employee name, and the associated department name, such as can be seen in the FIG. 3 “All Emp and Dept” table. For instance, the output lists Smith, Robinson, Jasper, and Clerical for the key 34, providing little useful information.

32. To cure this deficiency with traditional MapReduce, it is necessary to include significant additional steps, such as additional map phases that construct new key values pairs and include particular metadata. Such steps can then be chained together in a consolidated

¹ Another way of notating this is to ignore, in the schematic, the keys generated at the mapping phase. This may be desired if there is already a key in common in the data prior to mapping that could be used by at the reduce phase. For example, each of the entries has a DeptID key, and, thus, an entry in the mapping phase could simply be “34, Smith,” as opposed to “(34, Smith), 1.” Such notation is used throughout the ’610 patent because reduction is ultimately affected with keys already present in the data prior to the mapping phase. In other words, the key-in-common is one the data comes with.

MapReduce program, using this metadata in future reduce steps to avoid outputting undesired values. However, these steps can be costly in terms of compute time, especially if the input contains a large number of terms.

33. Such drawbacks emphasize the limited utility and application of traditional MapReduce as of 2006. Merely being able to map and/or reduce data having the same schema severely restricts the sources of data from which input data can be drawn, resulting in substantially-restricted output data that has inherently-limited usefulness. Indeed, a user can have a user profile with Business A that includes business name, username, password, first name, last name, and email. The same user can have a user profile with Business B that includes business name, username, password, first name, last name, email, *and* birthday. Despite the fact that the two profiles belong to the same user, they nevertheless have different “schema.” Thus, these two profiles cannot be reduced together using traditional MapReduce without the use of additional steps in processing.

34. This simplistic example epitomizes the foundational issue with MapReduce as of 2006—if data is collected with different fields, traditional MapReduce cannot effectively “reduce” the data after mapping. The purpose of big data analytics, and, specifically, MapReduce (as mentioned above), is to elucidate useful trends and other information from massive amounts of data. From a statistical standpoint, using billions or trillions, as opposed to, e.g., hundreds or thousands, of data points increases the fidelity of the data—recognized trends become more and more reliable as they emerge from more and more data. Thus, the ability to analyze data from a plurality of sources paramount in implementing big data analytics.

C. The '610 Patent Describes an Improvement to MapReduce.

35. The '610 patent offers an improvement over traditional MapReduce to specifically address these issues discussed above. The specification summarizes these enhancements:

In accordance with an aspect, an input data set is treated as a plurality of grouped sets of key/value pairs, which enhances the utility of the MapReduce programming methodology. Utilizing such grouping, map processing is carried out independently on two or more related datasets (e.g., related by each being characterized by a schema with a key in common). The intermediate results of the map processing (key/value pairs) for a particular key are processed together in a single reduce function by applying a different iterator to intermediate values for each group. Different iterators can be composed inside reduce functions in ways however desired.

Thus, for example, the enhanced MapReduce programming methodology may be easily employed to carry out distributed relational database processing.

'610 patent at 1:31-44.

36. The specification goes on to explain how implementation of, e.g., “data groups” realizes these improvements:

In general, partitioning the data sets into data groups enables a mechanism to associate (group) identifiers with data sets, map functions and iterators (useable within reduce functions to access intermediate data) and, also, to produce output data sets with (group) identifiers. It is noted that the output group identifiers may differ from the input/intermediate group identifiers.

'610 patent at 3:58-64.

37. In my opinion, consistent with the summary above, the '610 patent provides a clear technological improvement to existing MapReduce systems by describing and implementing a novel MapReduce architecture where mapping and reduce functions can be

applied to data from heterogeneous data sources (i.e., data sources having different schema) to accomplish the merger of heterogeneous data based on a key in common between or among the heterogeneous data. As the specification explains:

[T]he MapReduce concept may be *utilized to carry out map processing independently on two or more related datasets* (e.g., related by being characterized by a common key) *even when the related data sets are heterogeneous with respect to each other*, such as data tables organized according to different schema. The intermediate results of the map processing (key/value pairs) for a particular key can be processed together in a single reduce function by applying a different iterator to intermediate values for each group. In this way, operations on the two or more related datasets may be carried out more efficiently or in a way not even possible with the conventional MapReduce architecture.

'610 patent at 8:47-58 (emphasis added).

38. The technical features that facilitate this improved MapReduce architecture include, for example, “data groups,” selective mapping based off of said “data groups,” and selective processing in the reduction phase based off of said “data groups.” These technical features find ample support and explanation in the specification. For example, the '610 patent explicitly discusses implementing “data groups” to enable “map processing [to] be carried out independently on two or more related datasets (e.g., related by being characterized by a common key).” '610 patent at 2:2-4. This is true “even when the related data sets are heterogeneous with respect to each other, such as data tables organized according to different schema.” '610 patent at 8:49-51. The specification gives an example that “[t]he intermediate results of the map processing (key/value pairs) for a particular key can be processed together in a single reduce function by applying a different iterator to intermediate values for each group.” '610 patent at 8:51-54.

39. The specification also provides specific examples of how a “data group”-oriented MapReduce architecture can be deployed in programming. The ’610 patent shows, for example, that a “data group” identifier can be passed as a parameter to a mapping function and thereby travel with the resulting intermediate data, followed by a reduce function that employs the “data group” identifier to call particular programming objects (e.g., iterators) associated with specific “data groups”:

```
map(string group, string key, string val) {
  // group: "emp" or "dept"
  // key: DeptID
  // val: LastName for "emp", DeptName for "dept"
  DeptID = key;
  emit_to_group(group, DeptID, val);
}
```

```
reduce(hashtable iterators) {
  // iterators: one iterator for "emp", another for "dept"
  DeptID = iterators.get_key( );
  emp_iter = iterators["emp"].get_iter( );
  dept_iter = iterators["dept"].get_iter( );
  for each LastName in emp_iter do {
    for each DeptName in dept_iter do {
      emit_to_group("emp_dept", DeptID, (LastName,
      DeptName));
    }
  }
}
```

’610 patent at 6:29-33, 39-48.

40. Following a similar scheme, another specific embodiment is offered in the context of FIGS. 3-5, again showing how a “data group” identifier can be indicated to the mapping and/or reduce functions, such that different “data groups” can be mapped differently and/or processed in a way that corresponds to a particular “data group”:

```
map("emp", 34, "Smith") map("dept", 34, "Clerical")
```

```
reduce((34, ("emp", ("Smith", "Robinson", "Jasper")), ("dept",
("Clerical"))))
```

’610 patent at 7:45, 55, 64-65.

D. The Claims of the ’610 Patent Are Directed Towards Improving MapReduce.

41. The technological solution described in the specification is embodied, for example, in Claim 1 of the '610 patent, reproduced below in a bulleted-list format (for clarity) with particular emphasis on key technological features:

A method of processing data of a data set over a distributed system, wherein the data set comprises a ***plurality of data groups***, the method comprising:

- partitioning the data of each one of the data groups into a plurality of data partitions that each have a plurality of key-value pairs and
- ***providing each data partition to a selected one of a plurality of mapping functions*** that are each user-configurable to
 - independently output a plurality of lists of values for each of a set of keys found in such map function's corresponding data partition to form corresponding ***intermediate data for that data group and identifiable to that data group***,
 - wherein ***the data of a first data group has a different schema than the data of a second data group*** and
 - ***the data of the first data group is mapped differently than the data of the second data group*** so that different lists of values are output for the corresponding different intermediate data,
 - ***wherein the different schema and corresponding different intermediate data have a key in common***; and
- reducing the intermediate data for the data groups to at least one output data group, including
- ***processing the intermediate data for each data group in a manner that is defined to correspond to that data group***,
- ***so as to result in a merging of the corresponding different intermediate data based on the key in common***,
- wherein the mapping and reducing operations are performed by a distributed system.

(emphasis added).

42. In my view, Claim 1 succinctly claims an improved MapReduce architecture that embodies the exact technological improvements outlined in the specification. A crucial feature of Claim 1's MapReduce architecture is the implementation of what is referred to as a "data group."

I understand that this Court previously construed the term “data group” to mean “a group of data and a mechanism for identifying data from that group.”

43. The concept of “data groups” as found in Claim 1 of the ’610 patent in the context of MapReduce, thus, attains a novel and technological improvement in computer capabilities. As outlined in Claim 1, by utilizing “data groups,” a diverse data set can be fed to a collection of mapping functions within the same MapReduce architecture to ultimately be reduced and/or merged in spite of the diversity, and this is facilitated by a character of each “data group” (i.e., the “mechanism for identifying data from that group”) of the diverse data set following the data through the mapping. In this manner, the improved MapReduce architecture in the reducing phase is able to selectively employ specialized processing based on the “data group” from which the data being reduced originated. This specialized processing, enables the MapReduce architecture in the reducing phase to accomplish the merger of intermediate data hailing from different data groups.

44. These features are explicitly claimed in Claim 1, which requires that, e.g., “the data of the first data group is mapped differently than the data of the second data group” and “the data of [the] first data group has a different schema than the data of a second data group,” and further requires that the reducing step include “processing the intermediate data for each data group in a manner that is defined to correspond to that data group” to accomplish “a merging of the corresponding different intermediate data based on the key in common.” In short, Claim 1 expressly claims the exact key technological solutions that culminate in improvement of the functioning of a computer process (i.e., MapReduce).

45. The technological significance of the improved MapReduce’s application of “data groups” is apparent from the claims and specification of the ’610 patent. Throughout Claim 1,

the particular “data group” and/or data produced therefrom (e.g., “intermediate data”) plays a pivotal role in the improved MapReduce architecture. For example, a particular “data group” provides an identity of sorts to “intermediate data” (“intermediate data...[is] identifiable to that data group”); mapping is different from one “data group” to another; and processing during the reducing step corresponds to a particular “data group.”

46. More specifically, Claim 1 describes that “data groups” are partitioned and the resulting partitions are provided to “a selected one of a plurality of mapping functions.” Thus, the improved MapReduce architecture is capable of selective mapping based on a particular “data group,” which is advantageous in that particular mapping functions can correspond to particular “data groups” and account for, e.g., a specific schema of a particular “data group” in said mapping. Discrete mapping to obtain “intermediate data” on a per-“data group” basis can avoid, e.g., jumbling of data having different schemas into the same map—if data were so jumbled, it could ultimately hinder the reducing phase, such as by causing different values to be assigned to the same keys, thus causing reduction of the “intermediate data” to yield extraneous output data. As discussed further below, illustrative examples in the specification lend additional context to this advantage, showing mapping functions corresponding to particular “data groups.” *See, e.g.*, ’610 patent at 6:28-34; 7:36-60. This individualized mapping also sets the foundation for the combined reducing phase, such as can be seen in FIG. 5 and the accompanying discussion. ’610 patent at FIG. 5, nos. 504, 506; 8:14-45.

47. According to the claims, after data is provided to the mapping functions, the mapping functions output “intermediate data” for a particular “data group” that is “identifiable to that data group.” This is also an important technical feature—as revealed later in the claim, the MapReduce architecture “process[es] the intermediate data for each data group in a manner that

is defined to correspond to that data group.” Such association of the “intermediate data” with a particular “data group” is therefore key to realizing an improved MapReduce architecture, because in the reducing phase, the “intermediate data” (distinguishable based on the parental “data group”) is processed selectively as may be dictated by the particular “data group” from which the “intermediate data” came. So, for example, in the reducing phase, the MapReduce architecture can recognize that received “intermediate data” stems from particular “data groups,” and can therefore employ processing that corresponds to each “data group” to accomplish a combined reducing phase. For instance, the specification discusses that the reducing phase of one embodiment can employ “data group”-specific iterators to accomplish merger, and this is only possible because of the improved MapReduce architecture’s implementation of “data groups.” Without “intermediate data” that is identifiable to a particular “data group,” the MapReduce architecture in the reducing phase would not be able to distinguish one type of “intermediate data” from another, which would either obstruct the reducing phase completely or impair the fidelity of the data output by the reducing phase.

48. Hence, “data groups” and the subsequent association of “intermediate data” there with accomplish significant technical advantages, such as, for example, tailored processing for “intermediate data” of multiple “data groups.” This tailored processing can, for example, account for different “data group” schema to accomplish a reduction and merger of data having different schema. Vitally, these technological advantages can ultimately facilitate “a merging of the corresponding different intermediate data based on the key in common,” despite the fact that the “data of a first data group has a different schema than the data of a second data group.” Thus,

these technical features found in Claim 1, and discussed above, facilitate a technological solution to the problem of applying MapReduce to diverse data (i.e., data with different schema).²

49. In sum, Claim 1's use of "data groups," such as exemplified above, enables the MapReduce architecture to effect unique and/or customizable operations for a particular "data group," ultimately achieving merger of heterogenous data. As explained in the specification,

² While Claim 1 outlines a specific enhanced MapReduce architecture, the '610 is not limited to such improved architecture, and instead contemplates additional inventions germane to improved MapReduce. For example, Claim 33 is also directed towards an improved technological process (i.e., a different improved MapReduce architecture as compared to Claim 1), describing a specific process that accomplishes reducing of different sets of "intermediate data" together when the sets of "intermediate data" have different schemas. Claim 33 specifies that particular "data sets" that belong to particular "data groups" having particular schema are partitioned and provided to "a selected one of a plurality of mapping functions" to form first and second "intermediate data sets" with sets of "resulting key-value pairs" that are different from one another. Like Claim 1, providing the partitions to "a selected one of a plurality of mapping functions" can be advantageous in that, for example, particular mapping functions can correspond to particular "data groups" and account for a specific schema of a particular "data group" in said mapping. Further, while the schemas of the overarching "data groups" are different from one another, the schemas and the "first and second set of resulting key-value pairs" have a key in common." Such feature can, for example, enable the reducing step to employ the common key if so desired. In the reducing step, the "intermediate data sets" are reduced together "so as to form an output data set," with the reducing being "accomplished by iterating on at least one key from the first and second intermediate data set," and the "output data set has a different schema than the first and second schema."

These technological elements seek to solve the same problem as Claim 1, but the claimed improved architecture is distinct. Claim 33 specifies a different programming architecture that begins with overarching "data groups" having particular schema being split into "data sets," but the schema of the "data groups" (as well as the key-value pairs resulting from the mapping) have a key in common. Then, in the reducing step, it is explicit that the reduction is accomplished via "iterating on at least one key" of the "intermediate data sets" (that, per the previous limitation, have a key in common), and such process generates an output data set that "has a different schema than the first and second schema." All of these limitations, together, constitute a technological solution to realize an improved MapReduce architecture capable of working with data of different schemas, such as by, e.g., maintaining a key in common between the overarching "data groups" and reducing via iteration "on at least one key from the first and second intermediate data set." Via these technological features, Claim 33 claims another embodiment of an improved MapReduce architecture that implements different mechanisms as compared to Claim 1 to accomplish a similar goal, e.g., to address the technological problem in the art of not being able to apply MapReduce to data of different schemas.

such task was ineffective using traditional MapReduce architectures, meaning that the inventions in the '610 patent constitute, in my opinion, unconventional technological improvements to MapReduce technologies available in 2006. It is my understanding that such improvements are patent-eligible, and thus, in my view, the claims of the '610 is patent-eligible.

E. The Inventions of the '610 Patent Offer Unconventional Technological Features That Improve the Functioning of a Computer.

50. The inventions described and claimed in the '610 patent (such as in Claim 1) improve the speed, efficiency, effectiveness, and functionality of computer systems. The inventions provide an improvement in computer functionality rather than improvement in performance of an economic task or other tasks for which a computer is used merely as a tool. For example, the '610 patent states that the disclosed inventions “enhance[] the utility of the MapReduce programming methodology.” '610 patent at Abstract, 1:31–33, 1:66–2:2. All of the advantages of the improved MapReduce architecture reflect enhancements in computer functionality, unlocking the potential for MapReduce to be applied to mammoth data sets from numerous sources regardless of the heterogeneity of such data. Notably, the '610 patent schematically illustrates “improved MapReduce architecture” and processing thereof in FIGS. 4 and 5 and the descriptions thereof:




FIG. 4; see also '610 patent at 1:54-57.




FIG. 5; *see also* '610 patent at 1:58-62.

51. The claimed inventions provide for more dynamic, customizable, and efficient processing of large sets of data, including optimization of traditional MapReduce processes, which increases efficiency and reduces processor execution time. *See, e.g.*, '610 patent at 8:54-57. The '610 patent specifically discusses other types of enhancements to MapReduce architecture that are claimed in dependent claims:

In the improved MapReduce, the operations can collect information about the data that they are processing. They can also save the information as metadata of the corresponding data. For example, the partitioning, grouping, and sorting operations can keep the smallest and largest keys in each partition processed. This information can be saved as part of the metadata of each partition. When an iterator is created to scan through the records of a partition, this metadata can be made available to the reduce function through the iterators. One example benefit is that this metadata can speed up comparisons. For example, in a processing of the data to find all records whose keys are smaller than a particular key, an entire partition can be skipped if its smallest key is larger than the particular key.

'610 patent at 5:19-32.

52. In another example, the specification discusses a combiner function that “helps reduce the network traffic and speed up the total execution time.” '610 patent at 3:1–8.

53. Thus, in my opinion, the entirety of the '610 patent and the claimed inventions therein are necessarily directed to improvements in computer functionality. The focus of the claims is to improve upon an existing technological process (i.e., MapReduce) to address shortcomings present in such existing technological process. My understanding is that claims that improve an existing technological process constitute patent eligible subject matter, and it is therefore my opinion that the claims of the '610 patent are patent eligible.

54. It is my understanding that FedEx has argued that the '610 patent is directed towards “manipulating data using a known algorithm” and fails to recite significantly more than the abstract idea, and further “encompasses human activity.” Dkt. 13 at 31, 35-36. However, in my opinion, this position greatly oversimplifies the inventions of the claims and ignores the technological features discussed above.

55. The '610 patent claims much more than “manipulating data using a known algorithm.” Instead, and as discussed extensively above, the '610 patent claims a novel and improved computer architecture, i.e. an improved MapReduce architecture, which applies a number of technological features to greatly enhance the state-of-the-art as of 2006. While the inventions in the '610 patent certainly involve data manipulation, the claims themselves are directed towards a technical solution to the problems outlined in the specification that constitutes a specific asserted improvement in computer capabilities. As discussed above, the claims specifically address *how* the technological problem of being unable to apply traditional MapReduce to heterogenous data is solved, e.g., via “data groups,” selective mapping, “intermediate data” with particular characteristics, data of “different schema” that nevertheless have a common key, specialized reducing steps capable of selective processing based on a particular “data group” and/or “intermediate data” stemming therefrom, etc. These are key technical features that outline a significant improvement to an existing technological process, meaning that the inventions of the '610 patent are necessarily rooted in computer functionality.

56. In my view, to say that the claims “encompass[] human activity” reflects an extremely simplistic distillation of the technology into a concept that is merely involved in the claims, ignoring the technological solution to which the claims are directed. It is my understanding that FedEx focuses on the examples provided in the patent with respect to FIGS. 3

and 5 (which, in my view, are offered only to better communicate the technological improvements of the patent more clearly) and concludes that because a human could join the tables, this means that the claims cannot be patent eligible. Dkt. 13 at 35-36. This position does not align with my understanding of what constitutes patentable subject matter—indeed, both a construction crane and a human can lift a feather, but the human’s ability to lift the feather does not take away from the patentable technological features embodied in the crane. Clearly, the crane’s utility, enabled by a plurality of technical capabilities, becomes more and more apparent as the item to lift gets larger and larger. Such is the case with the improved MapReduce architecture claimed in the ’610 patent as input data sets get larger and larger. A human cannot complete MapReduce on a company’s big data by hand any more than he or she could hoist a steel beam to the top of a skyscraper by hand. Technology is absolutely necessary for these tasks, and technology is what is supplied by the ’610 patent.

57. In my view, FedEx’s position also ignores the simple truth that a human could *not* perform the steps of the claims, because the claims recite, specifically, actual computer functionality. To again use the crane example, even though a human could lift the feather, it would be impossible for the human to lift the feather in the exact same manner as the crane because, simply put, a human is not a machine, and a human does not function like a crane does. So the human’s ability to accomplish a task that can also be accomplished by a crane does not reflect on the patentability of the crane. The improved MapReduce architecture is accomplished only in the realm of computer functionality, implementing distributed processing, specialized data structures (e.g., “data groups”), schema-based delineation, specialized processing, iterators, etc. These are not human activities. Further, the recitation that the mapping functions are “user-configurable” does not reflect deviation from computer functionality—undoubtedly, every aspect

of computer capabilities finds a genesis in human activity, as humans build computers and write code. This limitation simply reveals that the improved MapReduce architecture leverages mapping functions that can be customized and/or configured by a user, which is still a technical feature of the improved architecture.

58. In my view, FedEx's position fails to consider the environment of the '610 patent that is ingrained in the claims. The claims are solving the technical problem presented by existing MapReduce systems, specifically, that traditional MapReduce is impractical with heterogenous data. The claims are not solving an everyday problem encountered by humans wherein a computer is used merely as a tool, but are instead improving upon a technological process. Thus, in my opinion, the improved MapReduce architecture to which the claims of the '610 patent is directed goes far beyond the use of a known algorithm, mere human activity, and/or an abstract idea.

59. Further, in my opinion, the inventions claimed by the '610 patent were far from conventional at the time of invention. Indeed, in 2006, these features and components were unconventional when taken alone, and were even more unconventional when taken combination. Utilizing "data groups" to accomplish "map processing independently on two or more related datasets" that are heterogenous, and subsequently performing selective processing within a reducing step to process "intermediate data" from said map processing together to ultimately affect MapReduce on diverse data cannot be characterized, at least in combination, as being well-known, routine, and conventional in at least October of 2006. I am unaware of any software or other service that put these elements into practice at that time. Thus, these elements cannot be ignored in determining the nature of the claims.

60. Importantly, and as outlined above, Claim 1 of the '610 patent explicitly provides for *how* the improved MapReduce architecture is accomplished, reciting process steps that delineate how a “data group” drives reducing of data of different schema together based on “data group” origin and a common key. In my opinion, these limitations define a specialized process executed by a novel programming architecture that improves the way a computer can store, retrieve, generate, and/or use data. This goes far beyond simply using conventional method of “a known algorithm in MapReduce.” Dkt. 13 at 32. Therefore, in my opinion, these claims recite a specific improvement over prior systems, resulting in an improved MapReduce architecture. It is my understanding that such improvements are patent-eligible.

61. Further, it is my understanding and opinion, given my study and research relating to distributed processing and MapReduce, that the claimed improved MapReduce architecture constitutes a specific structure inextricably intertwined with a computing device itself, particularly as it relates to the function of MapReduce technologies. To enable MapReduce to be used on diverse data (e.g., data of different schema), the '610 patent claims a specific approach, designing both the programming architecture (which includes “data groups,” “intermediate data” associated therewith, different schema with a common key, mapping and/or reducing on a per-“data group” basis, etc.) and improved MapReduce process that can leverage this architecture. Via these technological improvements, this approach generally enhances MapReduce technologies and widens applicability to diverse data.

62. Moreover, even if it is assumed that the claims of the '610 patent are directed to some “abstract idea,” at a minimum, each of the claims of the '610 patent includes a combination of elements that were not well-known, routine, and conventional at the time of the inventions, which I understand may preclude a finding that the claims are ineligible for patent. Specifically,

each of the independent claims of the '610 patent recite elements that, especially in combination, were not well-known, routine, and conventional at the time of filing of the patent application in October of 2006. For example, the elements recited in Claim 1 of the '610 patent were, in combination, neither well-known, routine, nor conventional at the time of filing of the application. As stated above, convention in 2006 was to not use MapReduce on heterogenous data because it was impractical. An improved MapReduce architecture that addressed this specific shortcoming was an advancement that broke convention with respect to MapReduce technologies.

63. The evidence cited above establishes that, at the time of the invention, the features and components were not well-known, routine, or conventional when taken as a combination. As such, they bear little resemblance to FedEx's characterization of the inventions as being no more than "manipulating data using a known algorithm" or an abstract idea that "encompasses human activity." The claims of the '610 patent recite inventions that constitute technical improvements of MapReduce architecture that were completely unconventional at the time of invention and helped resolve technical problems.

64. In my opinion, the claims of the '610 patent pose no risk of preempting a scientific field of endeavor, the basic tools of scientific and technical work, or any proposed abstract idea. Even if any claim was fairly characterized as an improved MapReduce architecture, there would be numerous ways of enhancing MapReduce that would not require one to perform the ordered combination of steps claimed. Dkt. 13 at 33. I understand, generally, that the principle underlying the "abstract idea" judicial exception to patent eligibility is one of preemption—the desire to prevent monopolization of *all* ways of practicing an abstract idea. The '610 patent does not do this, as no claim is directed merely to the result of "enhanced

MapReduce,” but rather only to a very specific way, using “data groups,” etc., to create a software architecture capable of performing MapReduce on diverse data.

F. Claim 1 Is Not Representative of All Claims of the '610 Patent.

65. It is further my understanding that FedEx has argued that, for the purposes of a patent eligibility analysis, Claim 1 is representative of all of the claims of the '610 patent. I disagree. For example, the claims that depend from Claim 1 recite different inventions and/or additional technical features to the innovations specified in the Claim 1.

66. For example, Claims 3 and 4 claim technical features related to specifying how intermediate data is “identifiable” to a particular data group. This constitutes a specific invention that is added to the improved MapReduce architecture of Claim 1. Similarly, Claim 5 adds the inventive concept of “employing an iterator that corresponds” to a particular “data group” in the reducing step, again directed to a specific mechanism by which processing during reduction can occur. Such limitation is meaningful, building on the selective processing feature enabled by the “data group”-centric improved MapReduce architecture. Claims 6 and 7 add several technical features, including “reducers” and method steps directed to specific mechanisms by which “intermediate data” is reduced in the reducing step of Claim 1. Claims 8 and 9 build on the inventive concepts embodied in Claim 7, themselves constituting specific inventions focused on further elucidating the partitioning step and adding methods steps, respectively. Claim 10 depends from Claim 9 and requires that the reducing step includes processing metadata. Claim 11 adds to this inventive concept, specifically limiting the processing performed in the reducing step to include an iterator corresponding to a data group, “wherein the iterator includes providing the associated metadata to the processing of the reducing step.” These are all, again, inventive concepts added by dependent claiming.

67. As additional examples, Claim 12 adds inventive concepts to Claim 5, requiring that the “intermediate data” have particular characteristics and reducers, and that the iterator introduced in Claim 5 operates in a particular manner. These are inventive concepts that constitute meaningful limitations. Claim 13 requires that processing in the reducing step includes processing data that is not intermediate data, and Claims 14 and 15 add further inventive concepts to Claim 13 by requiring that reducing steps are carried out by reducers and that the non-intermediate data processed in the reducing step is associated with one reducer or another. Claim 16 adds inventive concepts directly to Claim 1, requiring that the reducing step relates data “among the plurality of data groups.” All of these dependent claims constitute notable technical features that are, themselves, inventive—each one contains meaningful limitations that are distinctively significant in that they outline discrete elements that are directed towards different embodiments of an improved MapReduce architecture. Thus, in my opinion, Claim 1 is not representative of all of the claims of the ’610 patent. This same analysis above is generally applicable to Claims 17-32, and it is my opinion that Claim 17 is not representative of Claims 18-32 for the same reasons set forth above with respect to Claims 1-16.

68. Additionally, Claim 1 is not representative of Claims 33-39 or 40-46. A comparison of Claims 1 and 33 reveal significant and obvious differences between the two that preclude Claim 1 from being representative of Claim 33 for patent eligibility analysis. For example, different from Claim 1, Claim 33 utilizes “data sets” that belong to “data groups,” and first and second “intermediate data sets” having first and second sets of “resulting key-value pairs.” The reducing in Claim 33 is further explicitly accomplished via “iterating on at least one key from the first and second intermediate data set,” a limitation not found in Claim 1 at all. Thus, Claim 1 cannot be said to be representative of Claim 33.

69. Further, Claim 33 is not representative of Claims 34-39, because each of these dependent claims add inventive concepts of distinctive significance to Claim 33 from which they depend. For example, Claim 34 requires that “reducing is accomplished by iterating on at least one key from each of the first and second intermediate data set so that the output data set is a merging of a portion of the first and second intermediate data set.” This is a technological feature that provides a specific mechanism (and resulting benefit) with respect to the reducing step. Claim 35, similarly, adds that “a plurality of reducers” accomplishes the reducing step, and further adds two new method steps directed towards specifying a particular mechanism by which “intermediate data sets” can be reduced in the reducing step. Claim 36, depending from Claim 35, requires that “all of the key-value pairs of the first and second intermediate data set are provided to all of the reducers,” which is another inventive concept, revealing that all of the reducers receive data having, e.g., different schema. Claim 37 also builds on Claim 35, requiring that the reducers include particular tasks (such as, e.g., a “combine task” that provides specific advantages as discussed above), and further adds a method step that requires particular functionality with respect to metadata—these are meaningful limitations on top of both Claims 35 and 33. Successively building from Claim 37, Claims 38 and 39 provide limitations of distinctive significance that require that the reducing step perform particular operations with the metadata introduced in Claim 37.

70. In my opinion, just as with the dependent claims stemming from Claims 1 and 17, all of these dependent claims stemming from Claim 33 constitute significant technical features that are themselves inventions—each one contains meaningful limitations that are distinctively significant in that they outline discrete elements that are directed towards different embodiments of an improved MapReduce architecture. Thus, in my opinion, Claim 33 is not representative of

its dependent claims, and Claim 1 is not representative of all of the claims of the '610 patent.

This same analysis above is generally applicable to Claims 40-46, and it is my opinion that Claim 1 is not representative of Claims 40-46 for the same reasons set forth above with respect to Claims 33-39.

71. In summary, it is my opinion that the claims of the '610 patent provide technical solutions centered on technical improvements (e.g., improved MapReduce architecture) to solve technical problems related to traditional MapReduce, and this is very clear to me from the specification and the claims themselves as I have discussed above. These features offer significant improvements to MapReduce technologies that existed in 2006. Thus, in my view, all of the asserted claims of the '610 patent are patent eligible and not directed to a mere "abstract idea" as posed by FedEx.

II. '157 Patent

72. In my opinion, the inventions claimed in the '157 patent encompass specific technological advancements in search engine technologies. These advancements address and overcome several problems with prior technologies known in 2009. The '157 patent claims a novel intents-driven process and accompanying programming architecture that can enable not only enhanced result ranking, but further facilitate data object display entry customization to achieve specifically-engineered entries, realizing an enhanced technological process of search engine application.

A. Technical Background and Shortcomings of Prior Technologies

73. Anyone that has ever performed a Google search is familiar with search engine technologies. Search engines receive a user query and return results that are responsive to that query. As the technology has evolved, results built by search engines have become the

differentiating factor—users want the best results for their queries, including results that are both relevant to their query and easy to sift through. Building adequate results thus distinguishes one search engine from another.

74. The specification of the '157 patent discusses shortcomings in prior search engine technologies and explains how such shortcomings are addressed via the inventions described therein. The specification explains that if, as in the case of traditional search engines, the “engine simply regards a web query as, for example, a ‘bag of words’, the search engine will search for web pages and other data objects (e.g., images, audio files, text files) that contain, or are otherwise associated with, the individual words within the query.” '157 patent at 4:1-5. In that scenario, to achieve a passable result, the search engine then “employ[s] additional techniques known in the art to more narrowly tailor search results and may additionally employ sophisticated algorithms to rank search results.” '157 patent at 4:5-8. These results can also be difficult to sift through, with each returned objects generally including a non-descript title, abstract, and link. '157 patent at 4:10-15.

75. Thus, simply treating a user query as a “bag of words,” which was the state of the art in 2009, yields results that can be far afield from the purpose of the user’s search and further be difficult to make use of.

B. The '157 Patent Describes Improvements to Search Engine Technology

76. The '157 patent recognized a better way for search engine technologies to approach this technical dilemma:

Search results could be significantly enhanced if the likely intent of the query is known. For example, search results may be ranked such that results that are more relevant to the user’s intent appear at or near the top of the search results. Perhaps more significantly, however, the user’s intent can be used to customize the display

and behavior of a search result to be narrowly targeted to a user's intent. An illustrative list of such customizations could include a customized title or abstract for the result or specialized parameters of a displayed clickable URL to provide the landing page with information regarding the user's intent or triggered by the user's intent.

'157 patent at 4:16-26.

77. The inventions of the '157 patent are explicitly presented to solve technical problems as summarized above. Specifically, the '157 patent discloses an "intent"-driven search engine process that can utilize "intent" to affect significant enhancements. The specification makes clear that "intent" correspond to "goals that users pursue in a search query," and further makes clear that "intent" constitute specific programming germane to the disclosed process executed by specific programming architecture:

The query is then classified into one or more likely intents, which can include an unclassified intent when no defined intents match the query 2300. An intent is a mapping from many combinations of keywords to a relatively small set of common goals that users pursue in a search query or session of multiple queries. Often, the intent of the query is not explicitly stated in the keywords. While the space of possible queries, is very large, the set of intents is much smaller. Examples of intents relating to product queries can be, for example: official-site, research, purchase, dealer, support, or reviews. Examples of intents relating to local/map queries: directions, reviews, phone, hours-of-operation. In one embodiment, query intent may be determined by linguistic analysis of query keywords. In one embodiment, previous queries in the user session, user profile information such as preferences, the set of all queries from all users or any subset of all users (e.g. a subset of users having specific demographics or usage patterns), and click data from previous sessions for the current user as well as the set of all users or any subset of all users are used to determine query intent.

'157 patent at 9:42-61.

78. I understand that this Court previously construed the term “intents” to have its plain and ordinary meaning. Based on my expertise, the plain and ordinary meaning of “intents” in the context of the ’157 patent is “goals that users pursue in a search query” that exist as pre-programmed items within and/or accessible to the architecture.

79. The “intents”-driven architecture and corresponding process constitute a technological advancement as of July 2009. The ’157 patent describes (such as in the excerpt above) a process where query keywords can be used in a particularly unusual manner, i.e., to call upon a plurality of pre-programmed items (the “intents”) that are configured to then affect the additional steps executed in the process by the programming architecture (discussed further below). Such feature allows the one or more query keywords to be employed for much more than just document recall, as was conventional in 2009. Indeed, the query keywords, via the “intents,” can even ultimately impact how particular data objects are constructed within a result. This provides an added benefit of enabling keywords to be utilized for more than just relevancy analysis. Additionally, because “intents” can be determined from keywords and/or combinations thereof for query after query, this architecture provides for the continued re-configuration of “intents” determination to account for actual user and search engine metrics. Thus, implementation of “intents” in the search engine process allows for adaptability based on actual, quantifiable data far beyond what is possible with the “bag of words” approach. In other words, as more and more users interact with the search engine, and as the process determines more and more “intents” from query keywords, the “intents” determination can become more and more refined, ultimately leading to more refined data delivery.³ This advantage indicates that the ’157

³ The specification contains numerous discussions regarding how quantifiable data can be implemented to modify “intents” and/or the architecture’s determination of “intents” to enhance

patent provides a significant improvement to the technological process of search engine application.

C. The Claims of the '157 Patent Are Directed Towards Improved Search Engine Technology.

80. The claims of the '157 patent describe in-depth the technological features implemented to realize the improved search engine process described above. For example, Claim 1 outlines a technical solution to address deficiencies in then-existing search engine technologies. The claim is reproduced below for convenience, with particular emphasis on exemplary technical aspects:

A method comprising the steps of:

receiving, over a network, a query from a user, the query comprising at least one query token;

analyzing the query, using at least one computing device, to *identify at least one query keyword*;

determining, at least the one computing device, *a plurality of intents from the at least one keyword, each of the plurality of intents indicates a type of information regarding the query keyword that is likely to be desired by a user submitting the query*;

classifying the query, using the at least one computing device, *into at least one of the plurality of intents*;

identifying, using the at least one computing device, a plurality of data objects available over the network that match the at least one query keyword;

assigning, using the at least one computing device, *at least one of the plurality of intents to at least some of the plurality of data objects*;

ranking, using the at least one computing device, the plurality of data objects;

the effectiveness of the intents-driven process. *See, e.g.*, '157 patent at 4:27-35; 9:47-61; 11:53-63.

building a result, using the at least one computing device, using the ranked plurality of data objects, the result comprises a plurality of display entries, *at least one display entry customized to a respective assigned intent is constructed for each of the ranked plurality of data objects*; and transmitting the result, over the network, to the user.
(emphasis added).

81. In my opinion, Claim 1 itself outlines a novel process executed by a specialized programming architecture that constitutes a significant improvement in computer functionality. Each of the technical features emphasized above operates cooperatively to enhance the technological process of search engine application, and these advances define a novel improvement in computer capabilities.

82. For example, the step of “analyzing the query, using at least one computing device, to identify at least one query keyword” is a first step in realizing an improved search engine programming architecture that avoids the “bag of words” approach disparaged in the specification. As seen later in Claim 1, the at least one query keyword plays an important role in guiding the architecture by facilitating intent determination. The ’157 patent discloses specific embodiments of how keywords are determined, e.g., through specialized modules and/or computing device(s) capable of parsing and/or tokenizing queries, and these modules and/or computing device(s) are configured to identify “common entities, phrases, spelling errors, synonyms, variants of stems, and noisewords (eg ‘the’, ‘a’, ‘in’).” ’157 patent at 11:47-52; 9:36-41. Via such exemplary capabilities, the modules and/or computing device(s) can identify at least one keyword within a given query.

83. The next step of “determining, at least the one computing device, a plurality of intents from the at least one keyword, each of the plurality of intents indicates a type of information regarding the query keyword that is likely to be desired by a user submitting the

query” constitutes an important technical feature that builds on the keyword determination in the preceding step. Instead of simply including the keywords in a “bag of words” as would generally be done in traditional search engine applications, the ’157 patent employs the at least one keyword to explicate a plurality of pre-programmed features, i.e., “intents,” from a discrete collection of “intents” available to the architecture. As discussed above, the programming architecture of the ’157 patent is, in one embodiment, capable of “mapping” to one or more “intents” “from many combinations of keywords.” Claim 1 conveys to me, then, that these “intents” reside as particular programming accessible to the architecture, and the one or more keywords elucidated from the query serves to guide the architecture to the proper set of “intents.” *See also* ’157 patent at 9:44-47.⁴

84. As mentioned above, the “intents” serve to characterize the additional steps in the process of Claim 1. For example, Claim 1 reveals that “at least one of the plurality of intents” is assigned to “at least some of the plurality of data objects” identified by the process that match the at least one query keyword. Subsequently, the process builds a result wherein “at least one display entry customized to a respective assigned intent is constructed for each of the ranked plurality of data objects.” That is to say, the “intents,” determined from one or more query keywords (and/or combinations thereof) for a given query, “can be used to customize the display and behavior of a search result to be narrowly targeted to a user’s intent.” ’157 patent at 4:19-22. The specification further provides that “[a]n illustrative list of such customizations could include a customized title or abstract for the result or specialized parameters of a displayed clickable URL to provide the landing page with information regarding the user’s intent or triggered by the

⁴ The specification contains extensive discussion of “intents,” including specific examples of how one or more “intents” can be implemented in specific programming arrangements. *See, e.g.*, ’157 patent at 4:41-55; 5:56-64; 6:15 – 8:30; 8:40-54; 10:25-37; 13:26-63; 14:5-39; 15:1-23.

user's intent.” ’157 patent at 4:22-26. This provides a technological advantage that is an improvement upon search engine technologies existing in 2009.

85. For example, the state of the art at that time of the ’157 patent was to enhance search engine functionality via ranking algorithms, and such algorithms are discussed in the ’157 specification, with Claim 1 further including the step of “ranking...the plurality of data objects.” See ’157 patent at 8:62-64. Unconventionally, however, Claim 1 of the ’157 patent focuses not just on enhancing ranking, but on improving search engine functionality in an entirely original way—Claim 1 uses “intents,” determined from query keywords, to affect data object display entry customization in result generation, regardless of and/or in addition to data object ranking. Thus, data objects of the result may not only be ranked, but display entries thereof are further engineered to correspond in function and/or presentation to a particular “intent” in a myriad of different ways. This practice, in my opinion, greatly improved upon the existing technological processes of search engine application known at the time, magnifying the power of simple keywords to impact more than just object recall.

86. In short, the ’157 patent, via its novel approach to search engine processes, created an entirely new metric, above and beyond simple ranking, by which to judge search engine efficacy: the existence and/or implementation of non-uniform data object display entries capable of indicating particular information, such as, e.g., user intent, object type, etc. Further, Claim 1 outlines a technical implementation of this exact improvement, setting forth specific process steps explaining *how* a process executed via a specialized programming architecture goes

from receiving a query to engineer data object display entries to accomplish intent-driven customization.⁵

D. The Inventions of the '157 Patent Offer Unconventional Technological Features that Improve the Functioning of a Computer.

87. As discussed above, the inventions claimed in the '157 patent improve the speed, efficiency, effectiveness, and functionality of computer systems. Moreover, the inventions provide an improvement in computer functionality rather than an improvement in some other task for which a computer is used in its ordinary capacity. For example, the focus of the '157 patent is to avoid the “bag of words” (followed by extensive ranking) approach and ultimately achieve better, more-usable computer-generated results as compared to technologies that existed in 2009. Such a problem is necessarily a technical one, only brought about due to the existence and usage of search engines, and the '157 patent thus claims improvements to search engine technologies. My understanding is that claims that improve an existing technological process constitute patent eligible subject matter, and it is therefore my opinion that the claims of the '157 patent are patent eligible.

88. It is my understanding that FedEx has argued that the '157 patent is directed towards the abstract idea of “[p]roviding search results based on a user’s intent,” and that FedEx has characterized the inventions in the claims of the '157 patent as being “what a sales clerk does every day.” Dkt. 13 at 25-26. However, this position and characterization greatly oversimplify the inventions of the claims and ignore the technological features discussed above.

⁵ Specific embodiment regarding display entry customizations, such as via customizable presentation rules, are thoroughly discussed throughout the specification and addressed in dependent claims of the '157 patent. *See, e.g.*, '157 patent at 4:36 – 8:67; 11:1-23; 12:23 – 13:3.

89. The '157 patent claims much more than “providing search results based on a user’s intent.” In fact, Claim 1 does *not* claim that results are returned based on “intents;” rather, the process “identif[ies]...a plurality of data objects...*that match the at least one query keyword,*” meaning that results are returned based on query keyword, not “intents.” Instead, and as discussed extensively above, “intents” in Claim 1 are programmed items, elucidated according to query keywords, that drive query classification and how data objects are displayed in the results. Dependent claims add further features facilitated by intents, such as ranking. This key technical feature, necessarily rooted in computer functionality by its very nature, is ignored by FedEx, and such improvement goes far beyond the activities of a sales clerk in an auto parts store. *See* Dkt. 13 at 26. Additionally, FedEx’s position fails to consider the context of the '157 patent that is echoed in the claims—the claims are solving the technical problem presented by existing search engines that stem from the “bag of words” approach and typical result generation. The claims are not solving an everyday problem encountered by humans wherein a computer is used merely as a tool.

90. Further, in my opinion, the process claimed by the '157 patent was far from conventional at the time of invention. Indeed, in 2009, these features and components were unconventional when taken alone and were even more unconventional when taken in combination. The pre-programmed items (“intents”), determined via query keyword and employable to classify queries and facilitate data object display entry customization, as well as the improved search engine process tailored to deliver specially-engineered results realizable because of such “intents,” are not “abstract.” They are instead technical in nature, and in combination, cannot be characterized as being well-known, routine, and conventional in at least

July of 2009. I am unaware of any software or other service that put these elements into practice at that time. Thus, these elements cannot be ignored in determining the nature of the claims.

91. Importantly, and as outlined above, Claim 1 of the '157 patent explicitly provides for *how* the data object display entry manipulation is accomplished, reciting process steps that delineate how a received query ultimately culminates in customized entries on a per-“intent” basis. In my opinion, these limitations disclose a specialized process executed by a novel programming architecture that improves the way a computer can store, retrieve, generate, rank, and/or present data. This goes far beyond simply using conventional methods to “provid[e] search results based on a user’s intent” via a computer. Dkt. 13 at 25-26. In my opinion, these claims recite a specific improvement over prior systems, resulting in an improved search engine technology. It is my understanding that such improvements are patent-eligible.

92. It is further my understanding and opinion, given my study and research relating to search engines, that the claimed improved, intent-driven search engine process constitutes a specific structure inextricably intertwined with a computing device itself, particularly as it relates to the function of search engine application. For any search engine to be well-designed, it is critical that the search engine deliver the best possible content to users based on the least detailed search queries. To this end, the '157 patent claims a specific approach for transforming how and what search engines generate in response to queries, designing both the programming architecture (which includes the “intents”) and improved search engine process that can leverage this architecture—via these technological improvements, this approach generally enhances search engine technologies, such as by improving the overall generated results of search engines with specially-customized display entries.

93. Furthermore, even if it is assumed that the claims of the '157 patent are directed to some “abstract idea,” at a minimum, each of the claims of the '157 patent includes a combination of elements that were not well-known, routine, and conventional at the time of the inventions, which I understand may preclude a finding that the claims are ineligible for patent. Specifically, each of the independent claims of the '157 patent recites elements that, especially in combination, were not well-known, routine, and conventional at the time of filing of the patent application in July of 2009. For example, the elements recited in Claim 1 of the '157 patent were, in combination, neither well-known, routine, nor conventional at the time of filing of the application. As stated above, convention in 2009 was to utilize a “bag of words” approach and subsequently rank results, and implementing an intents-driven process and accompanying programming architecture to realize specially-engineered display entries to indicate information regarding, e.g., user intents, was an advancement that broke convention with respect to search engine technologies.

94. The evidence cited above establishes that, at the time of the invention, the features and components were not well-known, routine, or conventional when taken as a combination. As such, they bear little resemblance to FedEx’s simplistic abstract idea of “providing search results based on a user’s intent.” The claims of the '157 patent recite inventions that constitute technical improvements of search engine technology that were completely unconventional at the time of invention and helped resolve technical problems.

95. In my opinion, the claims of the '157 patent pose no risk of preempting a scientific field of endeavor, the basic tools of scientific and technical work, or any proposed abstract idea. Even if any claim was fairly characterized as “providing search results based on a user’s intent,” there would be numerous ways of “providing search results based on a user’s

intent” that would not require one to perform the ordered combination of steps claimed. I understand, generally, that the principle underlying the “abstract idea” judicial exception to patent eligibility is one of preemption—the desire to prevent monopolization of *all* ways of practicing an abstract idea. The ’157 patent does not do this, as no claim is directed merely to the result of “providing search results based on a user’s intent,” but rather only to a very specific way, using specific intents-driven process and/or software architecture capable of facilitating such process to accomplish the improved search engine technology.

E. Claim 1 Is not Representative of all Claims of the ’157 Patent

96. It is further my understanding that FedEx has argued that for the purposes of a patent eligibility analysis, Claim 1 is representative of all of the asserted claims of the ’157 patent. I disagree. At least the dependent claims recite different inventions and/or add additional technical features to the innovations specified in the independent claims.

97. For example, Claim 4 builds upon the analysis step of Claim 1, specifying that the “analysis step identifies query keywords by parsing the tokens and identifying common entities, phrases, spelling errors, synonyms, variants of stems, and noisewords.” Such limitation is distinctively significant because it provides a specific invention wherein the analysis step is accomplished via a particular mechanism. Similarly, in Claims 5-10, inventive concepts are added to the classifying step of Claim 1 by specifying particular mechanisms included in the classifying step. Claim 11 also adds inventive concepts to the ranking step of Claim 1, specifying that “the plurality of data objects are ranked by intent.”

98. Claim 12 similarly adds significant technical features to Claim 1, specifying that “in the building step, each of the ranked plurality of data objects to which at least one intent has been assigned are matched against at least one presentation rule associated with the respective

data object and matching presentation rules are used to construct the at least one display entry for the respective data object.” Such limitations are distinctively significant, providing specific mechanisms for display entry customization that constitute stand-alone inventive concepts with respect to the building step in Claim 1. Claims 13-18 add significant inventive concepts to dependent Claim 12, requiring particular characteristics for one or more presentation rules. Claim 19, depending from dependent Claim 15, adds even more technical features providing more limitations with respect to the determined “intents,” data objects, and presentation rules. Claims 20-21 add similar limitations and inventive concepts to Claim 16. All of these limitations found in the dependent claims constitute meaningful limitations, setting apart the dependent claims from the independent claims as distinct inventions.

99. With respect to Claims 2 and 22-23, the same discussion as laid out above with respect to Claims 1 and 3-21 applies.

100. In summary, it is my opinion that the claims of the ’157 patent provide technical solutions centered on technical improvements (e.g., intents-driven architecture and an improved search engine process capable of leverage said architecture) to solve technical problems related to search engines, and this is abundantly clear to me from the specification and the claims themselves as I have discussed above. These features offer significant improvements to search engine technologies that existed in 2009. Thus, in my view, all of the asserted claims of the ’157 patent are patent eligible and not directed to a mere “abstract idea” as posed by FedEx.

III. ’329 Patent

101. Upon review of the claims of the ’329 patent and accompanying disclosure, it is my opinion that the ’329 patent describes specific technical advances in search engine technologies. The technologies claimed in the ’329 patent overcome several problems with prior

technologies by challenging well-established paradigms that defined (and inherently limited) search engine and ranking technologies known in 2007. Such paradigms and resulting drawbacks are discussed at length in the specification of the '329 patent.

A. Technical Background and Shortcomings of Prior Technologies

102. Like the '157 patent discussed above, the '329 patent describes advances in search engine technologies that address failings in the art. For example, with respect to traditional search engines, the specification discusses that these search engines generally “rank the documents according to relevance to the search query. Relevance is a measure of how closely the subject matter of a document matches a search query’s terms.” '329 patent at 1:67 – 2:3. Such reliance of search engines on relevance in ranking was conventional in 2007, and as the '329 patent explains, this functionality was problematic—ranking dictated by relevance meant that search engines could be tricked, such as by nefarious parties, “into recalling documents and inflating their ranking” using techniques known as “search engine spamming.” '329 patent at 2:6-8.

103. “Search engine spamming” consists essentially of cramming documents full of particular terms as a way to artificially inflate relevance and ultimately “trick search engine ranking algorithms into recalling and highly ranking documents that contain, for example, sponsored links to a web merchant.” '329 patent at 2:8-11. This technological issue greatly hindered search engine performance, causing a returned result for a given query to include irrelevant content that the querier did not desire, and perhaps even subjecting unwary users to malicious behavior. '329 patent at 2:14-17. The specification offers an illustration of these issues via an example of a user shopping for a camera:

A typical example of search engine spam is when a user tries to search for the terms “digital camera reviews” and expects to find pages which review various

models of digital cameras, detailing performance specifications, sample images and reviewer pros and cons list. Having this expectation when the user clicks on a link for one of the results, the user is instead led to a page that contains nothing but a plethora of keywords and links to other stores where he can buy the camera.

'329 patent at 2:18-27.

B. The '329 Patent Describes Improvements to Search Engine Technology.

104. Because of these issues with traditional search engine applications discussed above, the '329 patent recognized the need to depart from the standard protocol of having relevancy dictate search engine ranking while also providing barriers to search engine spammers.

See, e.g., '329 patent at 2:34-37. To that end, the '329 patent provided a novel approach:

As a crawler examines an individual document, one of the attributes that can be considered is section structure. In examining the various sections, the crawler identifies sections to ignore, that is, to not index in search engine indexes and or otherwise use for recalling the document. Such sections are referred to herein as “no-recall sections.” Those portions that are indexed for recalling are referred to as recall sections. In an embodiment, a crawler ignores no-recall sections demarcated by, for example, a tag. In another embodiment a no-recall section may be identified by analyzing section content rather than examining only delimiters. The terms inside no-recall sections do not contribute to the document term frequency counts and are not used for recalling the documents in response to search engine queries. However the no-recall sections are included as input to forms of analysis of the document that affect, for example, the document's ranking. Links inside the no-recall sections as well as the rest of the document may be followed in order to discover new content. The document may be analyzed for the amount of advertisements or other features in its entirety. Therefore, terms inside the no-recall sections can affect document ranking. Such analysis prevents hiding of search engine spam. A page containing copious

amounts of advertisements, or low quality links will be readily identified and ranked accordingly.

'329 patent at 3:7-31.

105. This approach outlined above culminates in two overarching technological improvements: 1) an improved data structure that is capable of facilitating both search engine recall and improved ranking via the attributes of recall and no-recall sections, and 2) an improved ranking process rooted in a specialized computing device and/or software capable of delineating between and selectively employing recall and no-recall sections found in a plurality of the aforementioned improved data structures. These two technological advancements, working in tandem, realize a discrete process and/or system that greatly improves upon search engine technology that existed in 2007.

C. The Claims of the '329 Patent Are Directed Towards Improved Search Engine Technology.

106. Significantly, and as an example, Claim 1 of the '329 patent embodies both of these technological features elucidated above. Claim 1, in its entirety, is reproduced below. The portions of the claim directed to the first technological improvement recited above are emphasized in blue, and the portions of the claim directed to the second technological improvement recited above are emphasized in green:

A method, comprising:

ranking a plurality of documents recalled by a search engine for a query;

wherein the plurality of documents contain certain documents, *each document of said certain documents containing at least one section that is not used by said search engine for recall* and one or more sections that are used by said search engine for recall;

wherein ranking a plurality of documents includes ranking said plurality of documents *based, at least in part, on the at least one section of said certain documents not used by said search engine to recall documents*; and; wherein the method is performed by one or more computing devices. (emphasis added).

107. The emphasized portions above are directed explicitly to the technological improvements set forth in the specification and explained above. Further, these claimed technological innovations solve the exact problems identified in the '329 specification—nefarious attempts to game Web query rankings are thwarted via implementation of 1) novel data structures (e.g., “documents” that include both recall and no-recall sections), and 2) an improved ranking process that circumvents sole reliance on relevancy by using such no-recall sections in ranking documents. '329 patent at 1:67–2:17. Further, by providing data structures with no-recall sections that are utilized by an improved ranking process, barriers are posed to would-be spammers—subjecting no-recall sections to ranking analysis prevents documents from being “designed so that content that increases recall and/or ranking potential is placed in the recall section and content that diminishes high ranking potential is hidden in a no-recall section.” '329 patent at 4:1-9. This allows “[a]ll the attributes in all of the sections of a document such as ‘links’, frequency of terms, coloring, font, etc.” to be considered in the spam and relevancy analyses. '329 patent at 4:13-16. A result is that a search engine can “affect the recall and ranking of documents to more accurately reflect relevance of the documents to search engine queries.” '329 patent at 3:1-3.

108. This inherent technical nature of the claims is apparent from both the claim construction rulings previously promulgated by this court for the '329 patent and the claim language itself. For example, the first limitation of Claim 1 requires “ranking a plurality of

documents recalled by a search engine for a query,” and this limitation alone divorces the claims from the “abstract idea” characterization FedEx has suggested and illustrated via the lawyer/librarian example. I understand that this Court previously construed “document” to mean “any unit of information that may be indexed by search engine indexes” to conform with the definition found in the ’329 specification, and further construed “section” to mean a “defined portion within the structure of a document.” *R2 Sols. LLC*, No. 4:21-cv-00091, Dkt. 54 at 54, 49. This Court also previously construed “recall” to mean “generating results for a search engine query.” *R2 Sols. LLC*, No. 4:21-cv-00091, Dkt. 54 at 49. Clearly, then, in my opinion, Claim 1 is expressly directed to improving a technical process as opposed to an abstract idea, because the language explicitly refers to and defines an improved data structure and an improved ranking process that only arise or otherwise have any meaning in the context of search engine applications, which is apparent from the Court’s construction of these terms. Thus, Claim 1 must be directed to something that is not an “everyday human activity,” because it is expressly limited to enhancing computer functionality. Dkt. 13 at 18.

109. By improving upon the functionality of existing search engine technology to address problems arising from implementation of such technology, the claimed invention is necessarily rooted in the improvement of computer functionality, as opposed to, e.g., enhancing the economy of a task usually performed by hand. Indeed, the entire focus of the ’329 patent is on solving a problem that exists *only* for computing devices that are implementing search engines—i.e., the claims of ’329 patent are specifically asserted improvements to computer capabilities in the field of search engine applications and ranking systems. My understanding is that claims that improve an existing technological process constitute patent eligible subject

matter, and it is therefore my opinion that the asserted claims of the '329 patent are patent eligible.

D. The Inventions of the '329 Patent Offer Unconventional Technological Features That Improve the Functioning of a Computer.

110. It is my understanding that FedEx has argued that the '329 patent is directed towards the abstract idea of “ranking search results,” and that FedEx has characterized the inventions in the claims of the '329 patent as being simply “abstract human activity that, at best, invokes computers merely as a tool.” Dkt. 13 at 17. For example, I understand that FedEx has argued that the invention of the '329 patent “ensnares everyday human activity,” such as a lawyer performing legal research, or a librarian using an electronic index. Dkt. 13 at 18. However, FedEx’s position is incorrect for a number of reasons.

111. Importantly, FedEx’s position fails to consider that, as mentioned above, the entire problem that the '329 patent seeks to solve arises *only* in the realm of computer-implemented search engines. “Search engine spamming” clearly is not an “abstract human activity” or “everyday human activity” that a lawyer or librarian encounters, and it clearly is not a problem that can be solved by any sort of “human activity.” Dkt. 13 at 18. Thus, the claims of the '329 patent, which are (per the patent itself) provided to solve this technical problem, must themselves be technical in nature. In other words, the claims of the '329 patent cannot be considered only an abstract idea for which computers are invoked merely as a tool, because the claims themselves only offer utility in the form of improving upon a computer process itself—that is to say, the claims do not implement a human activity via a computer, because the claims themselves describe a process that can only be performed on a computer to solve a problem that only exists within a computer.

112. FedEx’s position also seems to rely on a pronounced oversimplification of the technology that minimizes the benefits enabled by the inventions. Although the generic idea of “ranking search results” certainly existed prior to the invention, the asserted claims of the ’329 patent are again directed to a particular manner of structuring data into discrete units with particular characteristics and components, and further leveraging those characteristics and/or components to accomplish an improved ranking process that maintains the benefits of recalling documents using relevance, while also avoiding inherent drawbacks of traditional search engine ranking. In my opinion, FedEx’s position that the asserted claims are directed only to the generic idea of “ranking search results” reflects an extreme distillation of the inventions into an underpinning abstract idea that is merely involved with the claims, as opposed to being the focus of the claims; this abstraction is not what the claims are directed to. Indeed, in my view, FedEx’s position ignores the key technical features that represent the heart of the innovation, i.e., the specific and novel type of data structure designed to symbiotically facilitate application of an improved ranking process in connection with data recalled by search engines.

113. Further, in my opinion, the process of “ranking a plurality of documents recalled by a search engine for a query,” “wherein the plurality of documents contain certain documents, each document of said certain documents containing at least one section that is not used by said search engine for recall and one or more sections that are used by said search engine for recall,” “wherein ranking a plurality of documents includes ranking said plurality of documents based, at least in part, on the at least one section of said certain documents not used by said search engine to recall documents,” and “wherein the method is performed by one or more computing devices” as claimed by the ’329 patent was far from conventional at the time of invention. Indeed, in 2007, these features and components were unconventional when taken alone and were even more

unconventional when taken combination. The particular data structure enabling improved ranking, as well as the improved ranking process capable of using such data structures, are not “abstract” in any sense of the word—they are technical in nature, and in combination, cannot be characterized as being well-known, routine, and conventional in at least January of 2007. I am unaware of any software or other service that put these elements into practice at that time. Thus, these elements cannot be ignored in determining the nature of the claims.

114. For instance, Claim 1 of the ’329 patent requires that certain documents returned by a search engine “contain[ed] at least one section that is not used by said search engine for recall and one or more sections that are used by said search engine for recall.” Such limitation, especially in light of the claim construction previously decreed by this court, explicitly provides for *how* the improved ranking process is accomplished, i.e., via the novel technical aspects of a plurality of these specific types of data structures that have these specific attributes.⁶ Claim 1 further requires that “ranking a plurality of documents includes ranking said plurality of documents based, at least in part, on the at least one section of said certain documents not used by said search engine to recall documents”—this limitation also explicitly claims *how* the improved ranking occurs, i.e., via use of a specific technical feature of the novel data structure—the no-recall section.⁷ In my opinion, these limitations disclose a specific data structure that

⁶ The specification discusses multiple specific embodiments of a no-recall section being identified within a document, such as demarcation via one or more tags, section attributes, content attributes, link attributes, etc. ’329 patent at 3:14-31; 4:17-60.

⁷ The specification is replete with exemplary embodiments of the improved ranking process facilitated by no-recall sections. FIG. 3, for example, provides a flow chart of a process for improved ranking employing delineation and exploitation of no-recall sections. The portion of the specification discussing FIG. 3 additionally describes ranking analyses and elements thereof that are possible, including the creation of document models via “theoretic, algebraic or probabilistic approaches;” determination and usage of various specific attributes throughout sections of a document; link characteristics such as destination, text, quantity, and quality; section ratings; frequency of change in a particular section as compared to other sections;

improves the way a computer can store, retrieve, and/or rank data, and further disclose an improved ranking process capable of leveraging such data structures. This goes far beyond simply using conventional methods to “rank[] search results” on a computer. Dkt. 13 at 17. In my opinion, these claims recite a specific improvement over prior systems, resulting in an improved search engine technology. It is my understanding that such improvements are patent-eligible.

115. It is my understanding and opinion, given my study and research relating to search engines, that the claimed improved ranking process implementing discrete and novel data structures constitutes a specific structure inextricably intertwined with a computing device itself, particularly as it relates to the function of search engine application. Indeed, in January 2007, search engines were prominent, so prominent, in fact, that unique problems engendered by the existence of search engines themselves had already begun to pose challenges to their implementation (e.g., “search engine spamming” discussed above). As discussed in the specification, search engines at the time included “a combination of integrated software components” such as crawlers, indexes, ranking functionality, etc. ’329 patent at 1:18-19. As I discussed with respect to the ’157 patent, for any search engine to be well-designed, it is critical that the search engine deliver the best possible content to users based on the least detailed search queries. To this end, the ’329 patent claims a specific approach for designing both the data structure that defines the content accessible by the search engine, and the ranking process that interacts with the data structures, and via these technological improvements, this approach

demarcation; etc. ’329 patent at 4:17-60. In fact, multiple of these specific embodiments of improved ranking processes are covered by dependent claims of the ’329 patent, such as claims 6-7 and 13-14.

generally enhances search engine technologies, such as by addressing the issue of “search engine spamming” set forth in the specification.

116. Furthermore, even if it is assumed that the claims of the ’329 patent are directed to some “abstract idea,” at a minimum, each of the claims of the ’329 patent include a combination of elements that were not well-known, routine, and conventional at the time of the inventions, which I understand may preclude a finding that the claims are ineligible for patent. Specifically, each of the independent claims of the ’329 patent recite elements that, especially in combination, were not well-known, routine, and conventional at the time of filing of the patent application in January of 2007. For example, the elements recited in Claim 1 of the ’329 patent were, in combination, neither well-known, routine, nor conventional at the time of filing of the application. As stated above, convention in 2007 was to rank results based on relevance to the search query, and implementing a data structure with a no-recall section that could be exploited by a ranking process was an advancement that broke convention with respect to search engine technologies.

117. In my opinion, it is important to remember the shortcomings of the prior art systems discussed above. While it was well-known that sections of a document may be excluded from a search index, there was no service, product, or technology in existence prior to the inventions of the patents-in-suit that utilized improved data structures with particular sections recognizable by a ranking processes as to be used in ranking but not recall, much less any service, product, or technology that was so long in use and so commonly used that it would have gone *beyond* prior art at the time of invention to have become “well-known, routine, and conventional.” The ’329 patent resolved the underlying technical issues via the technological improvements discussed above. Indeed, this technical feature was the basis for the examiner’s

allowance of the '329 patent. In an interview held before allowance of the '329 patent, the examiner indicated that “the use of data within noindex metatags for recall or ranking is contrary to industry standards and is generally frowned upon.” Dkt. 13-3 at 57. And the Notice of Allowance stated that “[t]he reason for allowance lies specifically in the concept of using an ignored section only for purposes of ranking and not indexing.” Dkt. 13-3 at 28. In other words, the examiner recognized that ignoring a section for purposes of indexing but using it for ranking was a novel (and counterintuitive) method for ranking documents recalled by a search engine.

118. The evidence cited above establishes that, at the time of the invention, the features and components were not well-known, routine, or conventional when taken as a combination. As such, they bear little resemblance to FedEx’s simplistic abstract idea of “ranking search results.” The claims of the '329 patent recite inventions that constitute technical improvements of search engine technology that were completely unconventional at the time of invention and helped resolve technical problems.

119. In my opinion, the asserted claims of the patents-in-suit pose no risk of preempting a scientific field of endeavor, the basic tools of scientific and technical work, or any proposed abstract idea. Even if any claim was fairly characterized as “ranking search results,” there would be numerous ways of “ranking search results” that would not require one to perform the ordered combination of steps claimed. I understand, generally, that the principle underlying the “abstract idea” judicial exception to patent eligibility is one of preemption—the desire to prevent monopolization of *all* ways of practicing an abstract idea. The '329 patent does not do this, as no claim is directed merely to the result of “ranking search results,” but rather only to a very specific way, using specific data structures and a computing device and/or software capable of interacting with such data structures to accomplish the improved ranking process.

E. Claim 1 Is not Representative of all Claims of the '329 Patent.

120. It is further my understanding that FedEx has argued that for the purposes of a patent eligibility analysis, Claim 1 of the '329 patent is representative of all of the asserted claims of the '329 patent. I disagree. At least the dependent claims recite different inventions and/or add additional technical features to the innovations specified in the independent claims.

121. For example, Claim 2 includes a technical feature building upon the improved data structure set forth in Claim 1, specifying that the “at least one section of said certain documents not used by said search engine to recall documents are demarcated by a tag of an element having a name from a set of one or more names.” Such limitation is distinctively significant because it provides a specific invention wherein a no-recall section is indicated via a particular mechanism (i.e., a tag) with particular characteristics. Similarly, Claim 3 recites that “at least one document of said certain documents does not comprise data that contains element tags”—such limitation adds inventive concepts to the improved ranking process claimed in Claim 1, because specifying that “at least one document of said certain documents does not comprise data that contains element tags” necessarily requires that the ranking process used some other mechanism besides “element tags” to accomplish the “ranking [of the] plurality of documents.”

122. Claim 4, like Claim 2, adds further technical features to the invention in the form of another method step, requiring the “search engine [to] generat[e] data representing results of said query”; “wherein said data includes an abstract describing each document of said plurality of documents; and ” “for each respective abstract of each document of said certain documents, said abstract excludes terms from the respective at least one section not used by said search engine to recall said each document.” These limitations lend extensive and new inventive

concepts to Claim 1. Most obviously, the method of Claim 4 has an entirely new step involving the search engine from Claim 1—the search engine must generate a completely new data object (“abstract”), and the data object must not include terms from the no-recall section of the document to which the data object relates. Thus, the search engine in Claim 4 must be capable of distinguishing between, as well as selectively exploiting, both recall and no-recall sections to accomplish both document recall and the generation of the new data object(s) as to a particular one or more recalled documents. This is, again, a technical improvement and an additional inventive concept as compared to Claim 1 alone.

123. Claim 5, like Claim 4, also adds a completely new method step requiring that the “search engine us[e] a search engine index to recall said plurality of documents;” “wherein said search engine index does not index any term in said certain documents that is found only in one or more sections of said certain documents established as a no-recall section.” There is distinctive significance to these limitations as they add specific technologic enhancements to the search engine from Claim 1. In addition to specifying that the search engine uses a search engine index, it is clarified that the “search engine index does not index any term” that is found only in a no-recall section of one or more documents. Thus, the search engine index is itself a novel data structure, differing from traditional search engine indexes by specifically excluding those terms originating only in no-recall sections of documents. This narrows the method of Claim 1 by providing a specific mechanism by which the search engine can recall documents while simultaneously allowing for particular sections of the same documents to be passed over—e.g., the novel data structure of the claimed search engine index is configured to exclude such sections to enable the search engine to properly perform the requisite recall function.

124. Claims 6 and 7 also provide additional inventive concepts such that Claim 1 is not representative of these claims for a patent eligibility analysis. These claims provide specific mechanisms utilized by the improved ranking process of Claim 1, such as a “vector space model associated with each of said documents” and a “language model associated with each of said documents.” Like the previously discussed independent claims, Claims 6 and 7 are therefore distinctively significant in that they add inventive concepts to the independent claim they stem from.

125. With respect to Claims 8-14, the same discussion as laid out above with respect to Claims 1-7 applies.

126. In summary, it is my opinion that the claims of the ’329 patent provide technical solutions centered on technical improvements (e.g., novel data structures enabling and working with improved ranking processes) to solve technical problems, e.g., “search engine spamming,” and this is apparent to me from the specification and the claims themselves as I have discussed above. These features offer significant improvements to search engine technologies that existed in 2007. Thus, in my view, all of the asserted claims of the ’329 patent are patent eligible and not directed to a mere “abstract idea” as posed by FedEx.

CONCLUSION

127. In sum, I believe all asserted claims of the patents-in-suit are directed to patent-eligible subject matter. Factually and technically, no claim is directed to any abstract idea and all asserted claims recite an inventive concept, not merely that which was “well-known, routine, and conventional” at the time of invention. Each claim is presumed valid, and I am aware of no evidence, much less clear and convincing evidence, that would disturb that presumption.

I declare under penalty of perjury under the laws of the United States that the foregoing is true and correct.

Executed this 14 day of April, 2022

By: Bill Davis
Bill Davis