



Performance Baselines and Recommendations

September 17, 2019 | Version 10.1.290.1

Table of Contents

1 Performance baselines and recommendations	4
1.1 Performance baselines and recommendations considerations	4
2 dtSearch index build and Search Terms Report	5
2.1 Data Grid and Search Terms Report results	5
2.2 dtSearch index build and STR settings	5
2.3 Performance recommendations	5
3 Structured Analytics	7
3.1 Structured Analytics results	7
3.2 Structured Analytics settings	7
3.3 Performance recommendations	8
4 Analytics	9
4.1 Analytics index build results	9
4.2 Analytics clustering results	9
4.3 Analytics categorization results	9
5 Import	11
5.1 Native/Full Text import setup	11
5.2 Native/Full Text import results	11
5.3 Import fields	11
6 Export	13
6.1 Export results summary	13
6.2 Export fields	14
7 Productions and branding	15
7.1 Staging production results	15
7.2 Branding production results	15
7.3 Production import results	15
7.4 Production export results	16
7.5 Staging and Branding settings	16
7.6 Production import settings	16
7.7 Production export settings	16

7.8 Performance recommendations	17
8 Imaging	18
8.1 Imaging results	18
8.2 Imaging settings	18
8.3 Performance recommendations	18
9 Mass conversion	19
10 Processing	20
10.1 Processing profile and data source specifications	20
10.2 Discover + Publish performance baselines	21
10.2.1 Real World data set details	21
10.3 Discover + Publish performance baselines - Data Grid	22
10.3.1 Real world data set details - Data Grid	22
10.4 Inventory performance baselines	23
10.4.1 Inventory settings	23
10.5 Performance recommendations	24
11 Data set details	25
12 Environment details	26
13 Processing environment details	28
13.1 Tier 2 environment specifications	28

1 Performance baselines and recommendations

The performance baselines and recommendations documentation is meant to be used as a reference to track overall Relativity performance from version to version. It should not be used as a benchmark of what you expect to see in a production client environment due to differences in data, infrastructure, and configuration. The results may not scale linearly. For more information, see [Performance baselines and metrics considerations](#).

Note: Published performance metrics for Relativity 10.1 are added *as they become available* and are considered final.

This documentation provides an overview of performance metrics and pertinent recommendations for Relativity 9.6. For features capable of utilizing either SQL or Data Grid, performance testing was done on both systems and noted accordingly.

See these related pages:

1.1 Performance baselines and recommendations considerations

Unless otherwise noted, the following all hold true for Relativity 10.1 Performance Metrics:

- All results are based on the EDRM Enron V2 data set. For more information about the data set details, see [Data set details on page 25](#). For more information about the environment configuration for this test, see [Environment details on page 26](#).

Note: As of Relativity 9.3, the processor model on the performance Virtual Machines was changed from **4 x Intel E5-4627 v2 3.3Ghz 8C Processors** to **4 x Intel E5-4620 v2 3.3Ghz 8C Processors**. Note that performance in 10.1 may be affected by this processor model update.

- All tests were each performed individually on the environment, without any additional tests occurring simultaneously.
- Memory on all relevant machines was cleared immediately prior to running the tests in order to eliminate any plausible effects of data caching.
- All performance results have a 5% variance, and the results may not scale linearly.

For any questions, or to receive a copy of the data sets used in this performance test, please contact support@relativity.com.

2 dtSearch index build and Search Terms Report

The performance baselines and recommendations documentation is meant to be used as a reference to track overall Relativity performance from version to version. It should not be used as a benchmark of what you expect to see in a production client environment due to differences in data, infrastructure, and configuration. The results may not scale linearly. For more information, see [Performance baselines and metrics considerations](#).

Relativity's dtSearch engine provides advanced search functionality such as proximity, stemming, and fuzzy searches. It also supports the use of Boolean operators and custom noise word lists as well as the basic searching features available in keyword searches.

Relativity's Search Terms Report (STR) simplifies the process of identifying documents that contain a specific group of keywords.

2.1 Data Grid and Search Terms Report results

For Performance testing purposes, each term within the STR was a simple keyword; no advanced searching functions were present. For example, proximity search, Boolean conditions, or wildcards. 10,000 terms, each of which matched at least 1,000 documents, were chosen as the terms used for the terms used in the STR.

Test scenario	Execution time (hh:m-m:ss)	Documents/hr	Terms/hr
dtSearch Index Build	00:06:33	11,049,900	-
dtSearch Index Build - Data Grid	00:39:21	1,699,985	-
dtSearch Search Terms Report - 10,000 terms	00:36:31	1,815,592	16,430
Data Grid Search Terms Report - 10,000 terms	00:36:55	1,841,650	16,667

2.2 dtSearch index build and STR settings

For more information about the data set details, see [Data set details on page 25](#).

For more information about the environment configuration for this test, see [Environment details on page 26](#).

2.3 Performance recommendations

Refer to the information below when deciding how large to make a dtSearch sub-index.

Workspace document count	Recommended sub-index size
Less than 10 million	250,000 documents

Workspace document count	Recommended sub-index size
10 million to 25 million	500,000 documents
Greater than 25 million	1,000,000 documents

- **Agent count** - as with other processes in Relativity, adding more agents tends to make things faster. Unlike other agents, however, the dtSearch search agent will utilize all resources on the server it is running on. Because of this, there should only ever be one search agent running on any given agent server.
- **Agent temporary directory** - you can now configure agent servers to use specific UNC paths for dtSearch builds, system admins can segregate storage by agent server (reducing contention) or to use a NAS with faster write speeds (for building) without impacting the NAS which is optimized for reads (for searching).
 - Use caution when modifying this setting. In a virtual environment with shared storage, changing from the default may result in diminished performance due to additional data copy activities performed through the index build. Relativity recommends that you test any changes on your hardware before committing to any system-wide configuration changes.
 - If shared storage is involved or the local storage is not sufficiently isolated and fast, then you should use the default index share.
- **Agent server configuration** - for some other processes in Relativity, it is better to have many small agent servers, as opposed to a few large agent servers, or vice versa.
- **Agent configuration** - adding more CPU cores to an existing search server will make searches faster without any additional configuration.
- **Searchable content** - removing files without searchable content and using a separate index for database files and large Excel files can help optimize your index.
- **Subindex size** - this number represents the maximum size of each subindex and the number of subindexes. While less intuitive, this is perhaps the most critical setting to get right when building the index, as it impacts searching later. There are pros and cons to smaller indexes, such as the ability to make parallel at the cost of higher I/O cost. Finding a good default is a high priority.

3 Structured Analytics

The performance baselines and recommendations documentation is meant to be used as a reference to track overall Relativity performance from version to version. It should not be used as a benchmark of what you expect to see in a production client environment due to differences in data, infrastructure, and configuration. The results may not scale linearly. For more information, see [Performance baselines and metrics considerations](#).

The Structured Analytics application provides the functionality to perform convenient actions such as email threading, textual near duplicate identification, language identification, and repeated content identification on records within a workspace.

Within the area of structured analytics, Relativity 9.4 and above utilize a new Textual Near Duplicate Identification algorithm. The algorithm greatly enhances performance for both large and complex data sets. The algorithm lets users scale their Analytics server by adding CPU cores and RAM in order to achieve faster performance.

3.1 Structured Analytics results

The following table provides a breakdown of Structured Analytics performance.

Test scenario	Export duration	Structured Analytics operations duration	Import results duration	Total duration	Documents/hr
Language Identification	00:17:28	00:06:00	00:18:15	00:41:43	1,617,059
Textual Near Duplicate Identification	00:17:56	00:12:35	00:21:06	00:50:57	1,325,988
Repeated Content Identification	00:16:24	00:08:39	00:16:46	00:41:49	1,617,059
Email Threading	00:17:34	00:24:17	00:24:27	01:06:18	1,004,536
All operations	00:17:02	00:37:01	00:24:30	01:18:33	849,992

3.2 Structured Analytics settings

Structured Analytics settings	Value
Content Analytics (CA) version	4.0.5
Structured Analytics Worker Agents	8

Each Structured Analytics operation consists of three distinct phases:

1. **Export** - relevant document metadata and text is exported from Relativity to the Analytics engine.
2. **Structured Analytics operations** - the Analytics engine processes all documents, and determines the relevant information to be sent back to Relativity.
3. **Importing results** - the data that the Analytics engine extracted from the documents is sent back to Relativity through the Import API. This phase also includes job finalization and cleanup.

For more information about the data set details, see [Data set details on page 25](#).

For more information about the environment configuration for this test, see [Environment details on page 26](#).

3.3 Performance recommendations

The Relativity Analytics engine adds CPU power to your Analytics server. Adding CPU power may result in performance improvements with larger data sets. Contact support@relativity.com for more detailed recommendations on scaling the Analytics server effectively.

4 Analytics

The performance baselines and recommendations documentation is meant to be used as a reference to track overall Relativity performance from version to version. It should not be used as a benchmark of what you expect to see in a production client environment due to differences in data, infrastructure, and configuration. The results may not scale linearly. For more information, see [Performance baselines and metrics considerations](#).

Relativity Analytics helps you organize and assess the semantic content of large, diverse and/or unknown sets of documents using various Analytics features which are built upon Analytics indices.

Unlike traditional searching methods, Analytics is an entirely mathematical approach to indexing documents. An Analytics index is made up of both a training set that the index learns from and a searchable set that is searched upon.

4.1 Analytics index build results

The Analytics index build was performed using the default searchable and training sets. Start time was measured as the time the first document was sent to the Analytics server, and end time was measured as when the last document became active and searchable. Indices were built to not automatically remove English email signatures and footers.

Test scenario	Operation time (h:mm:ss)	Documents/hr
Analytics Index Build	01:17:01	860,845

4.2 Analytics clustering results

Use Relativity clustering to create groups of conceptually similar documents utilizing an existing Analytics index.

Analytics clustering was performed in a new cluster on all documents in the workspace. Start time was measured as the time the mass operation was initiated, and end time was measured as the time the job was removed from the relevant queue in the database.

Test scenario	Operation time (hh:mm:ss)	Documents/hr
Analytics Clustering	00:20:56	3,314,970

4.3 Analytics categorization results

Use Relativity categorization to create a set of example documents that Analytics uses as the basis for identifying and grouping conceptually similar documents.

Each Analytics categorization was run using a new categorization set consisting of all documents within the set. Samples were taken by coding the first 1,000 or 5,000 documents with extracted text, sorted by ArtifactID ascending, and coding them to one of the categories. Times were recorded from audits of the categorization set.

The test scenarios in the following table use a batch size of `AnalyticsCategorizationBatchSize = 100K`.

Test scenario	Examples	Categories	Operation time (hh:mm:ss)	Documents/hr
Scenario 1	1,000	2	00:44:59	1,473,866
Scenario 2	5,000	2	01:11:55	921,892
Scenario 3	5,000	5	01:11:54	922,106

The test scenarios in the following table use a batch size of AnalyticsCategorizationBatchSize = 50K.

Test scenario	Examples	Categories	Operation time (hh:mm:ss)	Documents/hr
Scenario 1	1,000	2	00:44:14	1,498,856
Scenario 2	5,000	2	01:11:59	921,038
Scenario 3	5,000	5	01:13:05	907,175

For more information about the data set details, see [Data set details on page 25](#).

For more information about the environment configuration for this test, see [Environment details on page 26](#).

5 Import

To test the import performance of Relativity 10.1, the Relativity Desktop Client (RDC) was used to import a folder and its sub-folders from Relativity. The test is conducted with a single load file where native files are copied to the location and text is written to its final location.

5.1 Native/Full Text import setup

Native + Full Text Imports setup	
# of documents	1,104,990
File Transfer Mode	Direct Mode
Bulk Share Configuration	Direct Mode
Default Repository Access	Direct Mode
Overwrite	Append Only
Parent Info	Folder Information Column = Folder Path
Extracted Text	Cell contains file location
Encoding for Undetectable Files	Western European

5.2 Native/Full Text import results

Test scenario	Import time	Records/hr	Avg Metadata transfer rate (MB/sec)	Avg SQL process rate (docs/sec)
Native/Full Text (Copy to Repository)	02:02:55	543,438	155.99 KB/sec	693

5.3 Import fields

- 46 fields were imported:
 - 12 long text fields (including Extracted Text)
 - 18 fixed-length text fields
 - 2 whole number fields
 - 7 date fields
 - 1 yes-no field
 - 2 single-choice fields
 - 2 decimal fields

- 1 multi-choice field
- 1 single-object field

6 Export

The performance baselines and recommendations documentation is meant to be used as a reference to track overall Relativity performance from version to version. It should not be used as a benchmark of what you expect to see in a production client environment due to differences in data, infrastructure, and configuration. The results may not scale linearly. For more information, see [Performance baselines and metrics considerations](#).

To test the export Performance of Relativity 10.1, the Relativity Desktop Client (RDC) was used to export a folder and its sub-folders from Relativity.

All exports were completed in direct mode, copying files directly from the repository. Export was the only action executing on the machine. Images were exported as single-page tiffs. Of the 1,104,990 documents tested, 114,783 had images. Extracted Text was exported into separate files. All metrics are calculated using the total number of documents. The average file size of all documents exported was 0.11 MB.

6.1 Export results summary

Export setup	
# of documents	1,104,990
File Transfer Mode	Direct Mode
# of images	383,849
# of documents with images	114,783
Default Repository Access	Direct Mode
Text Precedence	Extracted text
Text File Encoding	Western European

The following tables provide a breakdown of export performance:

Test scenario	Records/hr	MB/hr	.dat size (MB)	Avg File Transfer Rate (MB/sec)	Metadata transfer rate includes SQL Processing (MB/sec)	Export time (hh:mm:ss)	File size, natives + images (MB)
Natives + Full Text + Images(Copy to Repository)	432,293	49,293	1040	8.72	4.42	02:33:22	126,000

6.2 Export fields

- Natives + metadata
 - 46 fields were exported:
 - 11 long text fields (including Extracted Text)
 - 18 fixed-length text fields
 - 2 whole number fields
 - 7 date fields
 - 1 yes-no field
 - 2 single-choice fields
 - 2 decimal fields
 - 1 multi-choice field
 - 1 single-object field

7 Productions and branding

The performance baselines and recommendations documentation is meant to be used as a reference to track overall Relativity performance from version to version. It should not be used as a benchmark of what you expect to see in a production client environment due to differences in data, infrastructure, and configuration. The results may not scale linearly. For more information, see [Performance baselines and metrics considerations](#).

Relativity can image and produce your documents when review completes. Production is a Relativity Dynamic Object (RDO), and you can use different data sources instead of a mass operation to select the documents to produce. This means that you can set up a production set with a data source during any stage of document review so that you can click a button to stage and run the production when ready.

All times were measured using the timing information from the audits on the production set object.

7.1 Staging production results

Test scenario	# of documents	# of data sources	Total Stage Time (mm:ss)	Docs/secs
1	99,311	1	00:02:46	598

7.2 Branding production results

Test scenario	Average family size	Production time (hh:mm:ss)	Branding time (hh:mm:ss)	Images/hr
New Bates	N/A	00:00:32	01:11:56	223,513
New Bates with Relational Attachment	1.567	00:00:29	01:11:43	224,256
Original Bates	N/A	00:02:55	01:13:36	219,588
Original Bates with Relational Attachment	1.567	00:02:57	01:14:36	216,721

7.3 Production import results

Test scenario	Import time (hh:mm:ss)	Images/hr	Avg metadata transfer rate (KB/sec)	Avg SQL process rate (Docs/sec)
Images only	00:45:10	366,935	137.34 KB/sec	557

7.4 Production export results

Test scenario	Export time (hh:mm:ss)	Documents/hr	Avg metadata transfer rate (KB/sec)	Avg file transfer rate (MB/sec)
Natives only	00:05:27	3,375,804	949.37 KB/sec	7.56 MB/sec

7.5 Staging and Branding settings

Staging and Branding settings	Value
# of documents in production set	99,311
Images	281,317
Branding Agents	8
Files to Produce	Images Only
Attachment Relational Field	Group Identifier
Production Numbering	Page Level Number
Left Footer	Production Bates Number
Right Footer	Control Number

7.6 Production import settings

Production imports settings	Value
Copy option	No Copy
.opt file	32.0
File Transfer Mode	Direct Mode
# of images	281,317

7.7 Production export settings

Production export settings	Value
Copy option	Copy
File Transfer Mode	Direct Mode
# of exported documents	99,311

For more information about the data set details, see [Data set details on page 25](#).

For more information about the environment configuration for this test, see [Environment details on page 26](#).

7.8 Performance recommendations

Because productions rely on saved searches, when constructing the searches used as processing data sources, follow best practices for searching in Relativity.

8 Imaging

The performance baselines and recommendations documentation is meant to be used as a reference to track overall Relativity performance from version to version. It should not be used as a benchmark of what you expect to see in a production client environment due to differences in data, infrastructure, and configuration. The results may not scale linearly. For more information, see [Performance baselines and metrics considerations](#).

Imaging uses Invariant workers to complete both basic and native imaging. Testing was performed for native imaging only. All imaging tests were submitted through an imaging set and times recorded from audits on the imaging set.

8.1 Imaging results

Test scenario	Images created	Imaging time (hh:mm:ss)	Images/hr
Native imaging – 2 worker servers	280,618	02:04:48	134,912

8.2 Imaging settings

Imaging settings	
Document Pool	99,275
Average images/document	2.9
Image Type	TIFF
Imaging Set Basic image output Quality (DPI)	300
Basic Image Size	Original Setting

For more information about the data set details, see [Data set details on page 25](#).

For more information about the environment configuration for this test, see [Environment details on page 26](#).

8.3 Performance recommendations

Imaging does not currently have any performance recommendations.

Note: Running the exact same imaging set twice may result in a different number of documents. This is a known issue with the vendor software.

9 Mass conversion

The following table provides mass conversion performance via the service bus:

Release	Relativity 9.6
Platform	Service Bus
# of Docs converted	1,103,413
# of Workers/Agents	2
Total Conversion Duration (hh:mm:ss)	02:57:53
Documents/Hr	367,586

10 Processing

The performance baselines and recommendations documentation is meant to be used as a reference to track overall Relativity performance from version to version. It should not be used as a benchmark of what you expect to see in a production client environment due to differences in data, infrastructure, and configuration. The results may not scale linearly. For more information, see [Performance baselines and metrics considerations](#).

Relativity's processing feature ingests raw data directly into your workspace without the need for an external tool. The following details apply to the data on this page:

- The performance speeds included below are based on tests conducted against a specific environment setup. For the details of this setup, see [Processing environment details on page 28](#).
- We recommend that you conduct performance tests against your current environment setup, compare your speeds to those included on this page, and then adjust your environment accordingly, if necessary.
- Performance baseline calculations use the discovered file size, not the source data file size (e.g., File Size–after Discovery / [Discovery Time + Publish Time]).

To obtain the folders containing the EDRM and Real World test data sets, email support@relativity.com.

10.1 Processing profile and data source specifications

The processing profile used in the job that generated these performance baselines had the following specifications. The only field that was edited to something other than its default value was the Auto-publish set field, which was set to Yes.

- **Default time zone** - UTC Coordinated Universal Time
- **Defaults OCR languages** - English
- **DeNIST** - Yes
- **DeNIST mode** - DeNIST all files
- **Extract children** - Yes
- **When extracting children, do not extract** - <blank>
- **Excel header/footer extraction** - Do not extract
- **OCR** - Enabled
- **OCR accuracy** - Medium (average speed)
- **OCR text separator** - Enabled
- **Auto-publish set** - Yes
- **Deduplication method** - None
- **Default document numbering prefix** - REL

- **Default destination folder** - <name of workspace>
- **Do you want to use source folder structure** - Yes

Note: As of Relativity 9.5, the number of mapped fields has been increased to 125 to provide a more representative metric. Due to this change, all end-to-end Processing metrics take 10% longer.

In addition to the above settings, all of the available optional fields were mapped prior to running the job. For more information on the processing profile fields, see the Relativity Processing User Guide.

The test data used to generate these performance baselines (the Real World data set) was selected for the Source Path field on the processing data source. Note the following details of the Real World data set:

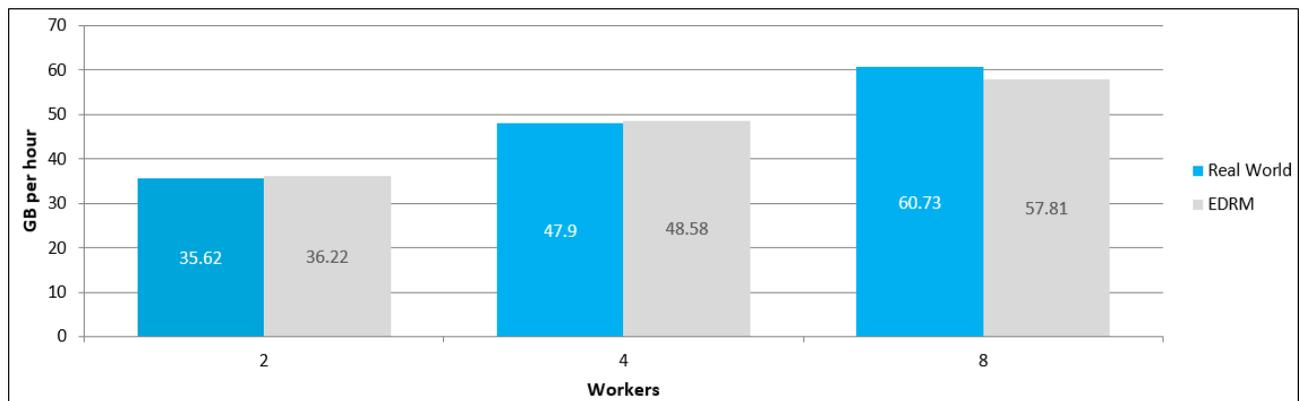
- File size - 60 GB
- File count - 100,000 assorted files
 - 70% emails - 42 GB of various PSTs and loose emails
 - 15% Office files - 10 GB of Excels, Word, PPTs,
 - 10% PDF - 6GB PDFs
 - 5% other (CAD files, Text files) - ~3GB

For information, see the Relativity Processing User Guide and [Processing environment details on page 28](#).

10.2 Discover + Publish performance baselines

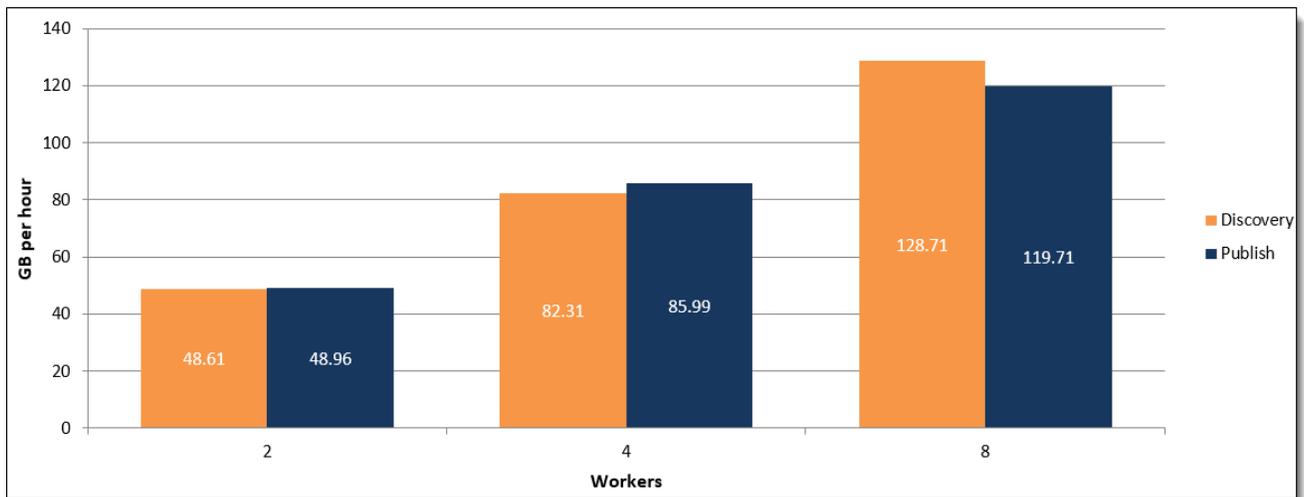
All time entries in the following table appear in the (hh:mm) format.

The following graphic depicts how many GB per hour are both discovered and published, per data set, per workstation setup.



10.2.1 Real World data set details

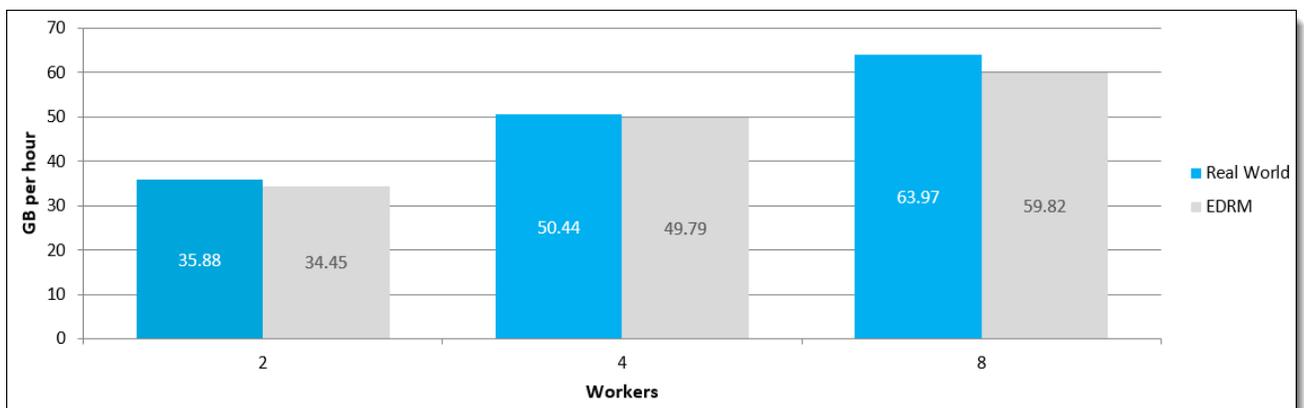
The following graphic represents the Real World Data set only; it depicts the breakdown of total time between discover and publish times, per workstation setup for the Real World Data Set only.



10.3 Discover + Publish performance baselines - Data Grid

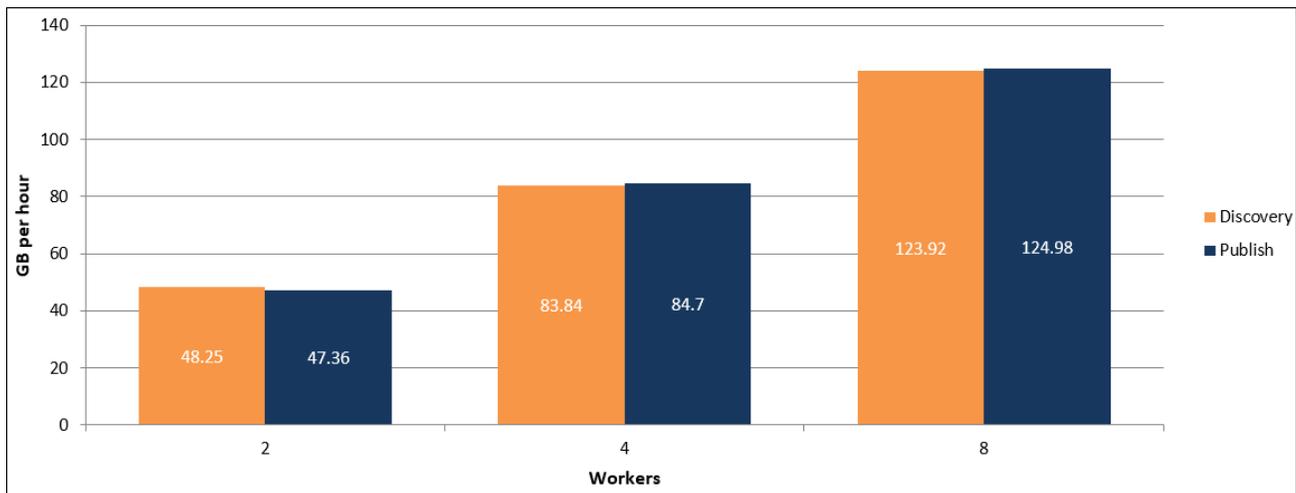
All time entries in the following table appear in the (hh:mm) format.

The following graphic depicts how many GB per hour are both discovered and published, per data set, per tier of workstation setup.



10.3.1 Real world data set details - Data Grid

The following graphic represents the Real World Data set only; it depicts the breakdown of total time between discover and publish times, per tier of workstation setup for the Real World Data Set only.



10.4 Inventory performance baselines

Inventory narrows down your files before discovering them by eliminating irrelevant raw data from the discovery process through a variety of preliminary filters. With inventory you can exclude certain file types, file locations, file sizes, NIST files, date ranges, and sender domains. Doing this gives you a less-cluttered data set when you begin to discover your files.

10.4.1 Inventory settings

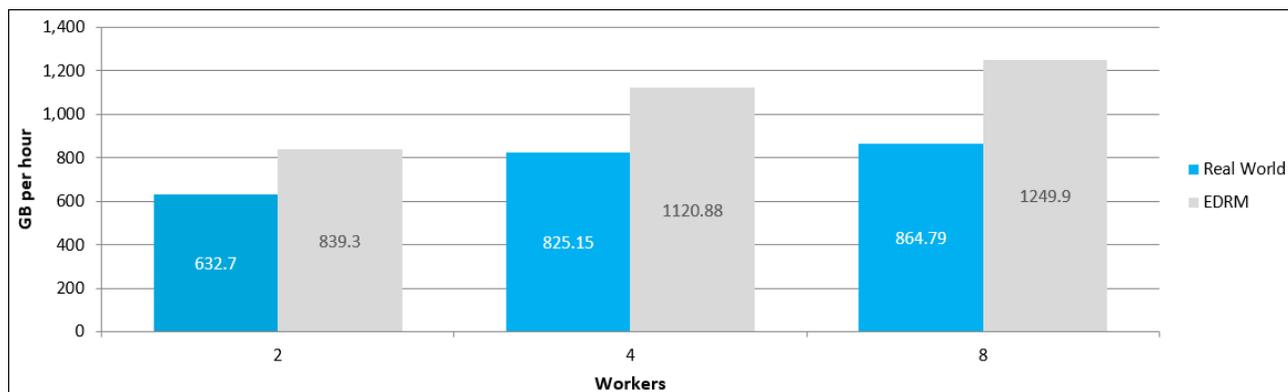
Inventory set up	
Data Set	Real World Data Set
Source file Size	60 GB
Source File Count	100k Assorted Files
Inventoried Documents	688,623

The following list contains the file breakdown of the Source File Size:

- 70% emails - 42 GB of various PSTs and loose emails
- 15% Office files - 10 GB of Excels, Word, PPTs,
- 10% PDF - 6GB PDFs
- 5% other (CAD files, Text files) - ~3GB

Note: There is almost no contention on the file share during the publish phase, but there is some SQL contention as data is inserted into the workspace database. SQL processor utilization will increase during the publishing phase.

The following graphic depicts how many GB per hour are processed by how many workers present in each workstation setup.



10.5 Performance recommendations

- **Worker Threads** - by default and if enough memory is available (750 MB/thread), Processing will run two worker threads per CPU core on each worker server up to 16 worker threads per worker machine. Further improvements could be realized by scaling horizontally with multiple worker machines.
- **Worker Machines** - when scaling Relativity Processing Worker machines horizontally, it is equally important to scale the file server if you want to continue to see linear improvements in performance. It will do no good to put 20+ worker machines online if your file server does not have the available IOPS to support them.
- **Invariant Queue** - during discovery, the Invariant Queue Manager, workers, and default file share for the workspace are heavily utilized. CPU, network, and disk resources should be closely monitored, to detect and prevent any issues which may impact performance for other processes in the system, particularly review and imaging.
- **Deleting a processing set** - this task runs as a background process, and will only proceed while other activity is not ongoing for the data store. This process may take days or weeks for large data sets or busy servers. If the data deletion needs to proceed more quickly, we recommend creating the following index in the relevant data store:
 - CREATE INDEX [IX_Matter_StorageId] ON [dbo].[Matter] ([StorageId])

Note: This index must be dropped or disabled prior to further processing work using that data store.

- **Scaling single sets** - in a workspace, you can horizontally run up to eight workers on a single processing set. If you run heavy imaging and processing loads, set up dedicated imaging workers. The maximum of eight workers applies only to a single processing set.

11 Data set details

The following are the details of the data set on which the performance tests were run:

Data set name	File count	Compressed Size (GB)	Expanded Size (GB)
EDRM Enron V2	1,111,439	41.62	97.14
Real World	1,045,510	60.17	99.51

12 Environment details

The following table provides Relativity environment server specifications:

Server	Memory (GB)	Cores	Host details
SQL*	256	32	<ul style="list-style-type: none">■ Dell PowerEdge R720xd Rack Server■ 2 x Intel E5-2650 v2 2.6Ghz 8C Processors■ 256GB 1866Mhz RDIMM (16 x 16GB)■ 18 x 600GB 10K SAS Hard Drives, 4 x 1.2TB 10K SAS Drives, 4 x 400GB SATA MLC Solid State Drives■ Broadcom 57800 2x10Gb BT / 2x1Gb BT Network Adapter■ PERC H710P RAID Controller

Server	Memory (GB)	Cores	Host details
Cache	4	4	<ul style="list-style-type: none"> ■ Dell Poweredge R820 Rack Server ■ 4 x Intel E5-4620 v2 2.6Ghz 8C Processors ■ 128GB 1866Mhz RDIMM (16 x 8GB) ■ 16 x 600GB 10K SAS Hard Drives ■ Broadcom 57800 2x10Gb BT / 2x1Gb BT Network Adapter ■ PERC H710P RAID Controller
Virtual File	4	4	
Web 1	8	8	
Search	64	8	
Data Grid - Data Note 1	16	8	
Data Grid - Data Node 2	16	8	
Data Grid - Data Node 3	16	8	
Agent 1	8	8	
Agent 2	8	8	
Invariant SQL	32	8	
Web2	8	8	
Data Grid - Master Node	4	2	
Data Grid - Client Node 1	4	2	
Data Grid - Client Node 2	4	2	
Data Grid - Data Node 4	16	8	
Data Grid - Data Node 5	16	8	
Data Grid - Data Node 6	16	8	
Worker*	16	8 physical, 16 logical	<ul style="list-style-type: none"> ■ Dell Poweredge R420 Rack Server ■ 2 x Intel E5-2420 v2 2.2Ghz 6C Processors ■ 16GB 1600Mhz RDIMM (4 x 4GB) ■ 1 x 400GB SATA MLC Solid State Drives ■ Onboard Dual Gigabit and Broadcom 57800 2x10Gb BT Network Adapter ■ PERC H310 RAID Controller

* The Relativity SQL Server and worker server are physical machines. All others are VM's.

13 Processing environment details

The tables below provide details on the tier 2 environments on which the processing performance tests were conducted. These environments falls into the environments outlined on the processing system requirements page. See Processing system requirements in the System requirements guide.

13.1 Tier 2 environment specifications

Server	Memory (GB)	Cores	Host details
SQL *	256	32	<ul style="list-style-type: none"> ■ Dell PowerEdge R720xd ■ 2 x Intel E5-2650 2.6Ghz 8 Core Processor ■ 16 x 16GB (256GB) 1866MHz RDIMM Memory ■ Broadcom 57800 2x10Gb BT / 2x1Gb BT Network Adapter ■ PERC H710P RAID Controller
Web	8	8	Split between two of the following: <ul style="list-style-type: none"> ■ Dell PowerEdge R820 ■ 4 x Intel E5-4620 v2 3.3Ghz 8C Processors ■ 16 x 8GB (128GB) 1866MHz RDIMM Memory ■ Broadcom 57800 2x10Gb BT / 2x1Gb BT Network Adapter ■ PERC H710P RAID Controller
Agent	8	8	
File	4	4	
Search/CA/Loading	64	8	
Invariant SQL	32	8	
Worker *	16	<ul style="list-style-type: none"> ■ 8 physical ■ 16 logical 	<ul style="list-style-type: none"> ■ Dell PowerEdge R420 ■ 2 x Intel E5-2420 v2 2.2Ghz 6C Processor ■ 4 x 4GB (16GB) 1600Mhz RDIMM ■ Onboard Dual Gigabit and Broadcom 57800 2x10Gb BT Network Adapter

* The Relativity SQL Server and worker servers are physical machines. All others are VM's.

Note the following additional details about Relativity's performance testing:

- Network topology
 - Relativity runs performance testing is run on multiple identical environments. Each environment consists of a physical SQL server and two ESXi hosts. The physical SQL is used for Relativity SQL, the ESXi hosts are used for hosting virtual machines which make up the other components of the Relativity environment. Shared components include the physical file servers and physical Invariant workers. All physical servers communicate over a 10Gbps network.
- Hard disk drive speed
 - The storage used is either 10K SAS HDD or SSD. We recommend accounting for RAID type and number of disks in each RAID group. Relativity was deliberate in creating RAID groups, in that we optimized each RAID group for its purpose.
- Bandwidth between physical servers
 - All physical servers in the environment are connected to a 10Gbps switch.
- Performance between physical and virtual servers
 - Performance difference between physical and virtual can vary from none to about 10%. CPU-intensive workloads such as SQL perform better on a network in which a virtual machine is installed directly on hardware rather than within the host operating system (OS). Note that Relativity SQL and Invariant workers are not virtualized and that virtual CPU and memory are not over-allocated in the environment, which eliminates CPU and memory contention on the ESXi hosts.

For more details on Relativity's performance testing environments, contact [Support](#).

Proprietary Rights

This documentation (“**Documentation**”) and the software to which it relates (“**Software**”) belongs to Relativity ODA LLC and/or Relativity’s third party software vendors. Relativity grants written license agreements which contain restrictions. All parties accessing the Documentation or Software must: respect proprietary rights of Relativity and third parties; comply with your organization’s license agreement, including but not limited to license restrictions on use, copying, modifications, reverse engineering, and derivative products; and refrain from any misuse or misappropriation of this Documentation or Software in whole or in part. The Software and Documentation is protected by the **Copyright Act of 1976**, as amended, and the Software code is protected by the **Illinois Trade Secrets Act**. Violations can involve substantial civil liabilities, exemplary damages, and criminal penalties, including fines and possible imprisonment.

©2019. Relativity ODA LLC. All rights reserved. Relativity® is a registered trademark of Relativity ODA LLC.