



Infrastructure Planning Considerations

December 8, 2021 | Version Server 2021

Table of Contents

1 Infrastructure planning considerations	3
1.1 Servers	3
1.1.1 Secret Store	3
1.1.2 Database (SQL)	3
1.1.3 Agent (background processes)	4
1.1.4 Web (application)	4
1.1.5 Analytics (search)	4
1.1.6 File server	5
1.2 Agents	5
1.3 Analytics	5
1.4 Architecture	6
1.5 Search indexes	7
1.6 File storage	7
1.7 Integration	9
1.8 Performance	9
1.9 Processing	10
1.10 Message Broker	11
1.11 Relativity Legal Hold	11
1.12 dtSearch	11
1.13	12
1.14 Virtualizing	13
1.15 Web servers	14
1.16 Secret Store	15

1 Infrastructure planning considerations

Use this guide to learn about the recommended infrastructure for a Relativity installation. It contains general information about system requirements, server functionality, and other guidelines for infrastructure planning. It also includes a list of frequently asked questions (FAQs) grouped by category, such as security and virtualization, for easy reference.

Note: See System Requirements for a complete discussion of server recommendations.

1.1 Servers

This section briefly explains the types and functionality of the different servers used in a Relativity installation. It also includes basic guidelines for memory requirements, configurations, and other information you can use to develop an effective infrastructure plan.

1.1.1 Secret Store

The Secret Store provides secure auditable storage for Relativity secrets. A secret could be user IDs and passwords stored within the Relativity response file, a database connection string stored in an application configuration file or instance settings containing confidential credentials. Use the following guidelines when configuring Secret Store.

- **Server Requirements** - Secret Store should be deployed to a dedicated host / virtual machine that shares no other Relativity server roles.
- **Database Requirements** - Secret Store requires a database server for hosting the Secret Store database. Our recommendation would be to use the Primary Database server, however, any database server may be used provided it meets system requirements for Relativity.
- **Unseal Key** - Once configured, Secret Store will generate an unseal key used to unseal the store when retrieving secrets. If the unseal key is lost, it will not be possible to recover previously stored secrets. In this scenario, Secret Store must be reinstalled to generate a new unseal key and secrets must be reimported.

1.1.2 Database (SQL)

The SQL Server stores the databases, which contain metadata, extracted text, and SQL log files. Use the following guidelines when configuring this server.

- **Physical versus virtualized SQL** - virtualizing SQL may simplify meeting HA and DR requirements; however, it can introduce additional layers of complexity when properly configuring and troubleshooting performance and stability-related issues.
- **Disk space usage metrics** - on average, every 100,000 records consume 10 GB of database space. On average database disk space usage is one-third of the total file size (natives and images) for a workspace. These metrics have a high standard deviation and tend to scale up or down linearly. This is an average. Actual files vary in size considerably.
- **IOP and latency best practices** - no IOP or latency requirements currently exist for a Relativity installation. We recommend that you follow Microsoft best practices and work with your storage unit

provider for optimal setup and configuration. The Relativity universe contains Dell, NetApp, EMC, Hitachi, BlueArc, IBM, and other devices.

1.1.3 Agent (background processes)

Relativity uses several different types of agents to perform background processing, such as imaging, branding, and indexing. These agents run on scalable agent servers. These agents can be processor-intensive depending on the types of jobs that are running. Use the following guidelines when configuring an agent server.

- **Agent requirements** - allocate one processor core and 1 GB of RAM for each additional agent in the environment.
- **Agent configurability** - add or remove agents to maximize production, depending on the jobs that you're conducting. The number of agents you add is configurable.
- **Enhancing agent performance** - partners who use Relativity for productions and imaging often have many branding or imaging agents to speed up these jobs. If you have a larger Relativity instance, you should have several agent servers.

1.1.4 Web (application)

Relativity runs on the web (or application) server. Use the following guidelines when configuring this server.

- **Number of web servers** - one web server, meeting our minimum hardware requirements of 8 processor cores and 8 GB RAM, supports 100 concurrent users. We recommend to start off with at least two web servers to help distribute user load and offer redundancy.
- **User load balancer** - Relativity includes a user load balancer to distribute user sessions evenly across all web servers included in the web farm. See Web load balancing in the Environment Optimization guide on web servers for more details. Windows NLB is also supported along with the Relativity User Load balancing option. Contact support@relativity.com if you need assistance.
- **Burstable connection** - a 10-20 Mbps dedicated, 100 Mbps burstable connection out of the data center(s) for every 50 concurrent Relativity users is sufficient. Monitor this for the most accurate bandwidth metrics.
- **Terminal services** - a Terminal Server or Citrix server is a recommended backup solution in the event that a review site has poor network connectivity or issues with installing the Relativity viewer.

1.1.5 Analytics (search)

Analytics is one of the default searching methods Relativity provides, which also includes Keyword Search (SQL Full Text Searching), and dtSearch. You can also use the analytics server for data analysis features such as email threading, near duplicate identification, clustering, and categorization. The size of the indexes and structured analytics sets that these searching analysis methods use depends on the number of columns, records, documents, and type of data being indexed on the workspace level.

- **Analytics index/set storage** - We recommend the disk location have storage speed and connectivity similar to the SQL Server, due to the complexity of the underlying database software.

- **Memory usage** - the Analytics server requires a large amount of memory, but it doesn't need a lot of processing power. The type of data dictates the RAM required. However, a good estimation for index building is 6 GB RAM for each 1MM documents. For structured analytics, estimate about 6 GB RAM for Java for each 1MM documents. Since the Analytics server uses a lot of RAM in a virtual host, use a physical server rather than virtualizing this role.

1.1.6 File server

Relativity uses the file server to store native and image files. Use the following guidelines when configuring this server.

- **File storage systems** - you can use NAS (NFS or CIFS), DAS, or SAN to fulfill this requirement.
- **UNC path requirement** - Relativity requires a UNC path to the share(s) that houses native and TIFF files. Depending on your storage unit, you may not require a Windows installation for this role. Don't install any Relativity software for this role.
- **Location** - there can be multiple file repositories across different locations. Start with one designated file repository that is a few hundred GBs in size, depending on anticipated load. You can add additional file repositories as the environment continues to grow.

1.2 Agents

The following FAQs provide information on Relativity agents and agent servers.

Which agents do the minimum requirements for an agent server support?

The minimum requirements for the agent server support one instance of all Relativity agents.

See the Agents guide for information on how many agents your version of Relativity supports and how to add them to your environment.

Can I scale the number of agents?

Yes, you can scale the number of agents. In addition, you can use more than one agent server in your environment.

For example, Relativity environments typically have multiple Branding Manager, dtSearch Index Worker, and agents. Partners who use Relativity for productions and imaging often have many Branding Manager or agents to speed up these jobs. If you have a larger Relativity instance, use several agent server.

We suggest allocating one processor core and 1 GB of RAM for each additional agent in the environment.

What are the requirements for agent machines?

Agent machines are server class. To use these machines for processing, they must communicate with the SQL Server and storage server. For specific hardware requirements on the agent servers, please reference the System Requirements document.

1.3 Analytics

The following FAQs provide information on Analytics servers, indexes, and disk space requirements.

What ports will be used for the Analytics server?

The web server must communicate with the analytics server via TCP ports 445, 8080, and 8443.

The initial communication from the analytics server to SQL is on ports TCP: 1433 and UDP:1434. By default, communication from the SQL Server to the analytics server occurs on port 8080 or 8443 (REST). However, the response from the analytics server to the SQL Server during a search query can occur on any port, thus all response (ephemeral) ports must remain open. A typical firewall rule would be to allow TCP established from analytics server port 8080 and 8443 to any port range 49152-65535.

The communication between Analytics and SQL is on port 1433.

How large will an Analytics index be on disk?

Using default settings, the average amount of disk space for the Analytics index is equal to 20% of the size of the MDF file of the workspace database. This metric has a high standard deviation and tends to scale up or down linearly. This metric indicates an average amount of disk space usage, but actual indexes may vary considerably in size. The amount of space required depends on the size of the extracted text being indexed, as well as the number of documents, unique terms, and dimensions used to build the index.

What type of file system should I use?

The Analytics Index Share houses all of your Analytics data for a particular Analytics server, and it can grow to be very large. We have found that NTFS file systems work for small environments, but if you anticipate running sets of 10 million or more documents through your Analytics Engine, you should use a file system that supports larger files such as exFAT or ReFS. We do not have a recommendation for either file system, so you must determine which is the better fit for you.

Can I have multiple index/set locations for a given analytics server?

No. An analytics server may only reference one disk location for the server's analytics indexes and structured analytics sets.

Can I have multiple analytics servers in my environment?

Yes. Analytics servers are scalable and are available as Servers in Relativity. You can allocate the Analytics servers to different workspaces via Resource Pools. Additionally, you can dedicate Analytics servers to Analytics Indexing, Structured Analytics, or both feature sets. For more information on configuring these options, see Adding an Analytics server in the Admin guide.

1.4 Architecture

The following FAQs provide information on scaling, performance, memory, and other issues related to the architecture of a Relativity installation.

What architectural model does Relativity use?

Relativity is a web-based review platform that rests on a distributed, grid-based computational model.

Can I connect servers used by a Relativity installation over the Ethernet?

You can use an Ethernet connection for the servers in your Relativity installation. Minimal connectivity speeds between SQL and the storage location are 4Gbps. Refer to the System Requirements document for specific network connectivity requirements.

Can I install the agent and the web servers on the same machine?

We advise against this server configuration. When users are heavily utilizing the agents, this configuration causes performance problems. It also makes troubleshooting performance issues difficult.

Can Relativity run over a WAN with a cache buffer?

Yes, Relativity can run over a WAN if it's set up for pass-through authentication. All authentications must pass through to Relativity.

Do you support changing SQL communication from port 1433 to another port?

For a default named SQL Server, switching port 1433 won't work.

For a named instance of SQL Server, you can use a static port for SQL communication. By opening up port 1434 and setting the SQL Server to use another port (for example, 1435), SQL will begin to communicate using port 1435.

However, for Relativity to work and communicate effectively over the chosen static port, UDP port 1434 also needs to remain open. Relativity uses this port for the initial connection to the SQL Server and then begins to use port 1435. Once SQL has been configured to use a static port such as 1435, if port 1434 isn't open, certain actions in Relativity might break. Some of the places where it breaks include upgrading and the Relativity viewer.

1.5 Search indexes

The following FAQs provide information on storing and managing search indexes in Relativity.

Where are the indexes for Relativity's various searching methods stored?

Indexes for the three searching methods offered in Relativity are stored in the following locations:

- **Keyword Search** - Relativity stores indexes for these searches in the SQL database.
- **dtSearch** - Relativity stores dtSearch indexes in the file repository.
- **Analytics** - You should store indexes and structured analytics sets on a drive letter attached to the Analytics server. This should have similar storage speed and connectivity of SQL Server.

Can I use multiple file repositories and dtSearch indexes in a Relativity installation?

Yes, you can configure your Relativity environment with multiple repositories and dtSearch indexes as necessary.

1.6 File storage

The following FAQs provide information on various file storage solutions for Relativity.

Is a Windows installation required for a file server?

No, a Windows installation isn't required, but you must provide Relativity with the path to the location where the files reside.

Can I use thin virtual provisioning with Relativity?

We don't recommend using thin virtual provisioning in VMware for the following reasons:

- **Host system on SAN** - SAN vendors differ in their support of thin virtual provisioning. If a device throttles connectivity speeds when certain thresholds are met, monitor these thresholds and configure alerts to notify system admins of these events. Provisioning can cause Relativity to reach capacity and a disk grow (zeroing) event to occur. If provisioning uses zeroing to grow the disks, performance is throttled at as little as 50 percent of expected disk IO.
- **Exhaustion of shared storage pools** - Relativity can exhaust storage pools when LUNs are over-provisioned and data writes to them. To avoid this situation, you must put procedures, warnings,

and protocols in place. Use thick provisioning and zero disks to prevent a rogue host from writing large amounts of data that infringes on other storage users.

- **Mismatched block sizes** - each implementation of thin provisioning uses a different block size. For example, HDS uses 42 MB on the USP, EMC uses 768 KB on the DMX, IBM uses variable sizes from 32 KB to 256 KB on the SVC, and 3Par uses blocks of 16 KB. While the reasons for these practices vary, file systems created on thin provisioned LUNs don't typically have a matching block size structure. This makes thin provisioning a bad choice for SQL data.
- **Mismatched Block Sizes Effects SQL Efficiency** - by default, Windows NTFS uses a maximum block size of 4 KB for large disks—unless it has been configured for a specific size. The mismatch between the block sizes for thin provisioning and the file system causes SQL performance problems on these systems when data is created, amended, and deleted over time. Consequently, thin provisioning negatively affects the efficiency of the SQL database when performing these operations.

What is the required speed or IOPS for a Relativity storage unit?

We do not have any specific storage speed or IOPS requirements for Relativity. However, you can use the following guidelines when selecting storage devices:

- **Fast storage** - select the fastest available storage for SQL data.
- **Minimal latency** - per Microsoft best practices for database SAN LUNs, SQL Server disk read/write latency should generally be 20 ms or less. The SQL Server disk read/write latency for transaction log SAN LUNs should 5 ms or less following Microsoft best practices.
- **Optimal performance configuration** - follow the guidelines from your storage device provider to ensure that it's configured for optimal performance. In addition, work with your device provider or follow Microsoft best practices to determine where different SQL data should live. Ensure that the link between the SQL Server and storage device is configured correctly for optimal performance by consulting your device provider. This setup becomes more complicated when virtualizing SQL.
- **Consider iSCSI** - while fiber has traditionally been considered a better option for enterprise class storage, iSCSI protocols have changed that to some extent. Even though most fiber cards now provide 4 GB of connectivity while Ethernet is limited to 1 GB, the speed differences aren't very pronounced. Fiber tends to cost more, but you want a dedicated Ethernet switch for iSCSI traffic. We have multiple clients using iSCSI without any problems.
- **Storage maintenance** - when you purchase and evaluate storage, be sure to discuss a disaster recovery plan for it, the qualifications of the staff required to manage it, and how it will support continued growth of the environment.
- **Space requirements** - on average, every 250,000 records in a Relativity workspace consumes about 20 GB of MDF disk space. As a rough metric, the total SQL data for a workspace is about one-third the total file size (natives and images) of that workspace.
- **Auto-tiering** - many new SAN devices can automatically shift data (auto-tier) across different types of available storage based on the activity levels of the stored data. You can use these devices to maximize performance with minimal manual intervention, reducing IT-related costs and maximizing use of available storage.
We use Compellent storage, while many of our partners have recently purchased EMC VNX devices. Other devices, such as NetAPP and IBM, also use some type of auto-tiering.

1.7 Integration

The following FAQs summarize information about integrating Relativity with different systems and platforms.

What other systems or platforms can be integrated with Relativity?

System or Platform	Integrates with Relativity?	Comments
DB2 database platforms	No	
SQL database platforms	Only Microsoft SQL	
Lotus Domino email system	Loosely integrated	The Native Imaging platform can read email messages stored in Lotus Notes databases.
Microsoft Outlook Exchange email system	Yes	
EMC SourceOne email archiving system	No	
IBM CM8 (z/OS platform)	No	
Network file systems	No	
IBM Lotus Quickr and IBM TeamRoom	No	
LDAP/Active Directory servers	Yes	
ImagePlus system	No	

What middleware servers does Relativity support?

- Windows Server 2012 R2
- Windows Server 2016
- Windows Server 2019 (Message Broker and File Server)

1.8 Performance

These FAQs provide information about optimal performance configurations for Relativity installations as well as latency detection.

How does the Bandwidth Tester that is built into Relativity measure latency?

In Relativity, the Bandwidth Tester utility uses an HTTP request to measure latency. The viewer records the sent time for a server request from a particular web page and the time at which it receives the page from the server. It determines the latency by subtracting these time values. These time values are high when network latency exists or when the web server is busy.

For more information, see the Admin Guide.

What is Virtual Machine Communication Interface (VMCI)?

The Virtual Machine Communication Interface (VMCI) is an infrastructure that provides fast and efficient communication between a virtual machine and the host operating system. It also provides fast and efficient communication between two or more virtual machines on the same host.

You can enable your application to use VMCI programmatically. VMCI provides greater inter-guest communication when virtual machine guests share a host. We don't offer built-in support for this feature. However, you may want to research existing tools for porting specific TCP endpoints to VMCI sockets.

1.9 Processing

These FAQs provide information about Processing in a Relativity environment.

If native imaging and processing jobs are scheduled for the same time, how will the workers on the worker agent handle this situation?

By default, imaging jobs take higher priority over processing jobs. However, you can access the management to change the priority of running jobs. There is never a need to manually configure things. If a processing worker has nothing to do, it asks the Queue Manager to give it something to do. Any processing worker can work on any job, regardless of job type. It's uncommon for even a single worker to have multiple threads working on both imaging and processing simultaneously.

Is it necessary to back up the processing databases? My understanding is that these hold temporary data for the most part, but should they be backed up like all other Relativity workspace databases?

Yes, you should back up processing databases.

What is the processing 250 GB space requirement used for on the worker server?

We recommend a decent amount of HD space on the worker machine for processing, because up to 16 threads can perform processing on each worker. Some of the files each worker processes could be quite large. 250 GB is a reasonable amount of space to perform temporary work before files get transferred to the network. The space is strictly used for temporary storage of data while workers are processing.

What is going on in the Invariant processing databases, is it very I/O, CPU, or memory intensive?

Use Processing to ingest raw data directly into your workspace for eventual search and review without the need for an external tool.

Does Relativity support virtualized processing SQL Servers?

Yes, Relativity supports virtualized processing SQL Servers as long as the server meets the minimum requirements.

Can a worker server be used for just specific types of jobs?

Yes, you can specify the type of jobs a worker can perform by editing the Worker designated work field on the worker server. You can configure a worker for one or more or all of the following jobs:

- Processing
- Save as PDF
- Native imaging

1.10 Message Broker

Relativity utilizes RabbitMQ as a message delivery service that communicates information about agent jobs to different application components. This infrastructure feature supports this communication by routing messages between application components. For example, Relativity uses RabbitMQ for submitting conversion jobs to agents and returning converted documents.

1.11 Relativity Legal Hold

The following FAQs provide information about Relativity Legal Hold (RLH).

What versions of Relativity does RLH work with?

Legal Hold compatibility matrix for version support details.

Can you upgrade from an older version to 3.0?

You can only upgrade from version 2.1.127.5 and above to 3.0. Any version beneath that, including 1.0, does not have a direct upgrade path. If a client really needs to maintain their data during an upgrade from an older version, reach out to PM.

Where can I confirm version information?

You can see it all in one place from the Configuration tab or look at the installed Applications.

Can RLH be installed to any workspace?

You can install to any workspace that doesn't have Processing already installed.

How many of the RLH agents should I have per environment?

We recommend starting out with a single agent because it doesn't use a lot of system resources. You can add more agents, if necessary.

How does the Custodian Portal authenticate?

There is no authentication done when accessing the Portal, but the Portal link is unique to every custodian. To prevent unauthorized use (via email forwarding of the link), you can specify the maximum number of times the link can be used or specify a link expiration date.

Does the Custodian Portal work with RSA authentication?

No, it doesn't.

1.12 dtSearch

The following FAQs provide information about dtSearch in a Relativity environment.

Can I use multiple file repositories and dtSearch indexes in a Relativity installation?

Yes, you can configure your Relativity environment with multiple repositories and dtSearch indexes as necessary.

Can my dtSearch Search agent reside in the agent server along with the other Relativity agents?

Each Relativity dtSearch search agent should have its own server or virtual machine, similar to existing dtSearch WebAPI web servers. System admins have the ability to add the search agent to servers with

other Relativity agents. However, this is not advisable as one dtSearch search agent will be able to utilize all available processor cores on a server. Therefore, each Relativity agent server that is designated to be a dtSearch search agent server should only have one dtSearch search agent and nothing else.

How can I estimate the size of my dtSearch indexes to allocate adequate storage space?

The dtSearch indexes grow to approximately 25% of the total database size. For example, if you had 10 TB of total files, the database would be 3-4 TB, and you would need about 1 TB for the index storage.

What is the default size of the dtSearch sub indexes that are created?

The default size of the sub-index is controlled through an instance setting. It can also be adjusted per each index creation through a setting on the index creation configuration page. The default value will be 250,000 documents. Larger data sets may benefit from having the sub-index size increased to 500,000 documents.

1.13

The following FAQs provide information on managing a SQL Server in a Relativity environment.

On the SQL Server, can an SSD be used to improve performance of the temp and log drives?

Yes, you can use a solid-state drive (SSD) on the SQL Server to improve performance.

Where should dtSearch indexes be stored?

We don't recommend storing dtSearch indexes on the SQL Server. Other than this location, you can store the indexes in the file repository location.

What is the recommended ratio of tempdbs to cores on an SQL Server?

Multiple files provide better I/O performance and less contention on the global allocation structures. For higher end systems with more than 16 cores, Microsoft suggests beginning with eight database files and adding more if there is still contention.

There are scripts available to measure contention. Most Relativity environments work well with eight data files. The Environment optimization guide describes several issues that you may want to consider when configuring the tempdb.

Is SQL aliasing or mirroring supported?

No, SQL aliasing isn't standard configuration supported by Relativity. Our Customer Support will work with you to determine whether or not aliasing can be successfully used in your environment.

We support Relativity in a mirrored environment, but mirroring is known to fail. If you use this configuration, and experience a failover, follow the instructions in the Configuring Relativity after a Disaster Recovery Failover document in the Relativity Community. Relativity won't be usable until you run the scripts included in that document.

Can I use a Gigabit Ethernet connection for inner server communication?

Yes, the focus on connection speed should be between the SQL and the storage system. A SQL Server on 1 GB iSCSI isn't sufficient. It constrains you during backups, index rebuilds, table scans, etc. These large sequential operations can easily saturate a 1 GB pipe, and storage quickly becomes a bottleneck. The minimum requirement we have is 4GBPS, but if possible, we recommend moving this up to 8-10GBPS.

What is the benefit of using multiple active SQL Servers in a Relativity installation?

With the use of multiple active SQL Servers, you can distribute your workspace databases across more than one SQL instance.

After configuring an active/passive cluster with one instance, can I modify it?

Yes, another instance can be added to this configuration and then moved to a second node or converted to an active/active cluster.

If there is a need to change the name of the Master Relativity SQL instance during the reconfiguration, Relativity Client Services will assist you with updating references to it in Relativity. (The Master Relativity SQL instance houses the EDDS database.)

What are the recommended allocations for disk space on the initial SQL instance?

For information about SQL partition sizes, see the System Requirements guide. That document includes sample partition sizes for the databases, logs, tempdb, and other features.

Should I activate the Priority Boost setting on the SQL Server?

No, you should never activate the Priority Boost setting.

Are other options available for database encryption besides "transit"?

SSL transit is the most commonly used database encryption method in practice. Some clients have been known to use Transparent Data Encryption, which is available in the enterprise edition of SQL Server.

Are there plans to support SQL Server 2019 with Relativity?

Relativity Server 2021 only supports SQL Server 2017 and SQL Server 2019.

Is Microsoft Windows Enterprise required for the SQL Server?

No, it is not. But you will want to make sure that the version of SQL server you plan on using supports the amount of ram you plan on allocating.

Is Microsoft SQL Server Enterprise Edition required for the SQL Server?

This depends mostly upon your memory requirements for the SQL Server. SQL Server 2012 and below support up to 64 GB RAM for Standard Edition. SQL Server 2014 supports up to 128 GB RAM for Standard Edition. For more information, see the System Requirements guide.

How large is the database and what kind of growth can I expect?

The database is small at the outset. It grows as it ages (from audits) and as data is loaded into it. The amount of SQL data created during a load can be anywhere from 1-2 times the size of the load file being loaded. This can possibly be even more if a lot of ANSI data is loaded into Unicode fields.

Does the application just connect to the database via a service account?

We require both a Windows Relativity service account and a SQL Server account (EDDSDBO), but the EDDSDBO server account is the only one that connects to the database itself.

The Windows Relativity service account requires permissions to the shared BCP Path folder on the SQL Server so that it can drop data there for loading.

1.14 Virtualizing

In this section, the FAQs provide recommendations for virtualizing servers and other related information.

What are the recommendations for virtualizing servers in a Relativity installation?

We recommend virtualizing all Relativity roles. Any hypervisor is supported.

What are your recommendations for using VMware?

We do not have any special recommendations for using VMware with Relativity. For general information, see the VMWare white papers.

1.15 Web servers

The following FAQs provide information web servers and other related information.

How many web servers are recommended for a Relativity installation?

Use the following guidelines to determine the number of web servers required for your Relativity installation:

- One web server dedicated for importing and exporting with Relativity Desktop Client (RDC).
- The Web (Application) role supports 100 simultaneously logged in users for every 8 processor cores and 8 GB RAM.

How do I set up Relativity in the DMZ so that users outside my network have access?

To configure the servers for DMZ access, you must open the following ports between the web server in the DMZ and all other servers in your environment:

- 445 (Microsoft-DS Active Directory, Windows shares)
- 1433 (MSSQL Server TCP)
- 1434 (MSSQL Monitor UDP)
- 8080 (Analytics)
- 8443 (Analytics)
- 443 (SSL) open to the outside
- 25 (SMTP)
- TCP 9090 (Secret Store)
- TCP 10000 (Secret Store)

In addition to the web server, what other servers should be in the DMZ?

Only the web servers should be in the DMZ.

What web load balancing methods does Relativity support?

- **Windows load balancing(single affinity)** - after a user's session begins on a single web server, it must remain on that server, or errors occur. In Relativity, a session state is stored in proc.
- **RULB (Relativity.User.Load.Balancing)** - RULB requires each included web server to have its own externally facing URL. To enable RULB you would go into the Servers tab of the admin page and edit a web server you want to include in the RULB. You then enter the URL for that web server you are editing.

Note: Once you click Save, Relativity checks the URL to see if it accessible so make sure the path is set up beforehand.

RULB works well with NLB and needs no extra configuration. All you have to do is turn on RULB, and Relativity does the rest. You probably won't see the NLB URL in IE change, but users should land on different web servers.

RULB also works with SSL certificates.

- **Round robin** - this is the default load balancing method. Round robin mode passes each new connection request to the next server in line, eventually distributing connections evenly across the array of machines being load balanced. Round robin mode works well in most configurations, especially if the equipment that you're load balancing is roughly equal in processing speed and memory.
- **Least connections (member) and least connections (node)** - the least connections methods are relatively simple. The LTM system passes a new connection to the node that has the fewest current connections. The least connections methods work best in environments where the servers or other equipment being load balanced have similar capabilities. These are dynamic load balancing methods. They distribute connections based on aspects of real-time server performance analysis, such as the current number of connections per node or the fastest node response time. Load balancing calculations may be localized to each pool (member-based calculation) or they may apply to all pools of which a server is a member (node-based calculation).
- **Priority-based member activation** - you can load balance traffic across all members of a pool or across only members that are currently activated according to priority number. In priority-based member activation, each member in a pool is assigned a priority number that places it in a priority group designated by that number. With all pool members available (meaning they are enabled, marked up, and have not exceeded their connection limit), the LTM system distributes connections to all members in the highest priority group only. That is, the group designated by the highest priority number. The Priority Group Activation value determines the minimum number of members that must remain available for traffic to be confined to that group. If the number of available members in the highest priority group drops below the minimum number, the LTM system also distributes traffic to the next highest priority group, and so on.

Which load balancers are commonly used?

RULB (Relativity.User.Load.Balancing) or NLB are the most commonly used load balancers.

What load balancing algorithms are available as Predictor options?

The following load balancing algorithms are available as Predictor options:

- **Round robin** - distributes connections evenly among the servers (default setting).
- **Weighted round robin** - enables priorities to be assigned to certain servers. The higher a server's weight, the more connections the load balancer sends to it.
- **Least Connections** - directs a service request to the server with the fewest open connections.

1.16 Secret Store

The following FAQs provide information on Secret Store and related information.

What ports are used for communication with the Relativity Secret Store?

The default ports utilized for secret store are TCP 9090 (the service port) and TCP 10000 (the nonce port). All Relativity Servers must be able to communicate with Secret Store over these ports and these ports are configurable.

Is it possible to recover secrets from the store if we have lost the unseal key?

No, it is not possible to recover secrets from the store without the unseal key. In this scenario, it is recommended that the existing secret store database be backed up, deleted and that Secret Store be

reinstalled to create a new database and generate a new unseal key. Secrets should be reimported into the store at this point.

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.

©2021. Relativity ODA LLC. All rights reserved. Relativity® is a registered trademark of Relativity ODA LLC.