Laserfiche Administrator II Training Manual

June 2010

Laserfiche Certified Professional Program. Not For Distribution or Resale.

Laserfiche®

Table of Contents

1.0 Introduction ... 3
2.0 Design Considerations ... 4
2.1 System Scaling ... 4
2.2 Gather Requirements for System ... 6
2.3 Planning ... 6
2.4 Network ... 12
3.0 Optimization ... 22
3.1 Folder Structure ... 22
3.2 Managing Metadata ... 22
4.0 Performance Considerations ... 35
4.1 The Client ... 35
4.2 The Administration Console ... 36
4.3 The Laserfiche Server and SQL Server ... 39
5.0 Performance Tuning ... 42
5.1 Bottlenecks ... 42
5.2 Performance Monitor ... 43
5.3 Laserfiche Search ... 56
6.0 Troubleshooting ... 63
6.1 The General Software Troubleshooting Strategy ... 64
6.2 List of Useful Tools ... 64
6.3 Error Logging ... 66
6.4 Compatibility ... 66
6.5 Repository Creation and Registration ... 66
6.6 Repository Migration and Updating ... 67
6.7 Troubleshooting Common Problems ... 68
6.8 How to Control and Use Tracing ... 74

Laserfiche Certified Professional Program. Not For Distribution or Resale.

1.0 Introduction

In this course, you will learn advanced methods for optimizing system performance, utilizing advanced metadata functionality, administering a comprehensive volume policy, implementing security for complicated scenarios, monitoring and evaluating system performance, and troubleshooting your Laserfiche implementation.

> **Note:** This course is designed for users familiar with basic functionality of the Laserfiche Client and Administration Console, as well as basic Windows tools. Be sure to watch the accompanying videos for supplementary information.

The manual covers five key areas of Laserfiche administration:

- **Design considerations:** How to plan hardware needs for a Laserfiche installation or upgrade.
- **Optimization:** Tips for getting the most out of your Laserfiche system.
- **Performance considerations:** Options in Laserfiche and MS SQL Server to improve the performance of the Laserfiche system.
- **Performance tuning:** How to analyze and maintain your operating environment to keep your Laserfiche system running efficiently.
- **Troubleshooting:** Common troubleshooting problems and suggestions for solving them, to help keep your system running smoothly.

Laserfiche Certified Professional Program - Not for Distribution or Resale.

2.0 Design Considerations

This section will walk you through the major steps you'll need to take when planning hardware needs for a Laserfiche installation or upgrade, from small-scale to "enterprise" installations. We'll then work through the requirements gathering phase, where you'll gather information about the complete picture of your needs when working with Laserfiche, and perform a comprehensive survey of existing hardware and software.

Specifying and selecting hardware is more an art than a science, because there's no single answer to any given question, and because hardware needs are never the same from company to company. Though this section contains a broad overview of the questions to ask and the decisions to make, there will be many variations that will require you to ask questions and make decisions that we've never considered.

> **Note:** The estimates and advice provided are intended as guidelines only, not as rules. We make no guarantees that the hardware we suggest here will suffice for every configuration.

2.1 System Scaling

The table below provides some guidelines for specifying how much processing power and memory might be required based on the number of concurrent users and the number of pages input to the system each day.

- **The green area** represents user needs that can typically be supported by a single server running Laserfiche, Web Access, and Microsoft SQL. We suggest at least a 2.5GHz quad-core processor with 8GB of RAM.
- **The pink area** represents needs that are best met by a dual-server configuration, with Laserfiche and Web Access on a single server, and Microsoft SQL on its own machine. For the Laserfiche/Web Access server, we recommend at least 2.5GHz quad-core processor with 4GB of RAM, and for the SQL server, the same processor with at least 8GB of RAM.
    - With a system of this size, organizations may consider virtualizing the Laserfiche infrastructure. We recommend allocating four processor cores and 8GBs of RAM to a single virtual machine hosting both the Laserfiche and Web Access servers.

    > **Note:** For performance reasons, we recommend keeping Microsoft SQL on its own physical machine, even if the other applications are virtualized.

- **The blue area** represents configurations approaching "enterprise" level, for which we recommend a three-server configuration: Laserfiche, Web Access, and Microsoft SQL each on their own machine, each with a minimum of a 2.5GHz dual-core processor and 8GB of RAM.

Laserfiche Certified Professional Program. Not For Distribution or Resale.

- For virtualization, we recommend a similar configuration where each server (except for MSSQL) resides on a dedicated virtual machine. The virtual machines should be allocated two processor cores and 4GBs of RAM.

- **Optimum storage configurations** also depend on the size of the system. For smaller Laserfiche implementations, such as the green area below, we recommend a single RAID 5 array. For the other estimated loads (both pink and blue areas), the ideal configuration is to set up three separate RAID arrays, one each for the Laserfiche volumes (RAID 5), Microsoft SQL Server data file (RAID 1 or RAID 1+0), and Microsoft SQL Server transaction logs (RAID 1 or RAID 1+0).

		Active Concurrent Users				
		10	25	50	100	200
PAGES/DAY 50kb ea	200	15GB*	30GB	45GB	60GB	75GB
	500	30GB	45GB	60GB	75GB	90GB
	1,000	60GB	75GB	90GB	105GB	120GB
	2,000	105GB	120GB	135GB	150GB	165GB
	3,000	150GB	165GB	180GB	210GB	240GB
	5,000	210GB	240GB	270GB	300GB	360GB
	10,000	450GB	495GB	540GB	585GB	660GB

*Please note that all storage estimates listed are for three years, which is the estimated life of current off-the-shelf hardware.

When you're dealing with systems that are truly enterprise, there are some best practices to follow. Primary among them is putting all Laserfiche components on separate servers. At Laserfiche, we've found that doing so is well worth the extra cost to avoid potential resource-sharing problems. Another consideration is additional environments—you're planning for Laserfiche in production, but you may also need to configure multiple additional support environments. Development, testing, training, and production environments may all be required for seamless performance, depending on your or your customer's requirements.

Lastly, as you create a hardware plan, realize that Laserfiche supports multiple cores per processor, as well as multiple sockets per server. Planning your hardware system to include multiple cores and sockets will allow you to take advantage of multi-threaded processing and help your operations run more efficiently and quickly.

2.2 Gather Requirements for System

The first step when planning the implementation/upgrade of Laserfiche is to determine what is expected from Laserfiche.

The most important question involves user base. When considering your user base, it is important to remember that there is a difference between the number of users and the number of users that might be estimated to be accessing the system—and therefore using your hardware resources—at any given time. We suggest an estimate of one-in-four total users accessing the system at any given time. Therefore, if you're looking at 200 user licenses for an enterprise implementation, you'll want to plan for your system to accommodate 50 simultaneous users at any given time.

> **Note:** Throughout this section, we will use the term "users" to refer to simultaneous, concurrent user load on the system. This term used here should not be confused with the number of user licenses or the previous Laserfiche licensing schema of "concurrent user licenses."

A related question is whether you will be deploying Laserfiche over the Web (e.g., Web Access or WebLink), and if so, to what portion of the user base, in order to calculate the amount of traffic that will be hitting your Web servers as well.

The next important question is the expected scanning volume: how many documents will be going into the system on a daily and monthly basis? The best approach is to try to predict this number over a three to four year timeframe (considering three to four years to be the typical lifespan of computer hardware). This figure should be in pages—later we will discuss converting this number to GB and its implications for storage needs. As a corollary to this, you'll need to determine if image capture will primarily be done through an installed Laserfiche Scanning client on the same local area network (LAN) as the Laserfiche Server or via Web Scanning. Again, this information will help you calculate loads required for your Web servers as well as any effect to network bandwidth.

Workflow and Quick Fields are harder applications to specify hardware for, but for the initial requirements gathering, it's enough to know how intense the use might be, and what, if any, Quick Fields add-ons will be employed.

2.3 Planning

Hardware needs for each application component will be considered separately.

Laserfiche Certified Professional Organization. For Distribution Not for Resale.

2.3.1 Windows Server 2008

For the Windows Server 2008 Standard edition, the maximum server resources supported are four processor sockets and 32GB of memory. For the Enterprise edition, the maximum server resources supported are eight processor sockets and 2TB of RAM.

2.3.2 Laserfiche Server

Generally speaking, having more cores is better for performance than having faster processors, and this is particularly true for the Laserfiche Server, which is a multi-threaded application and can make use of multiple processors or cores.

Concurrent users	Cores	RAM
200-400 users	2-4	4-8GB
400-1,000 users	4-12	8-16GB
>1,000 users	Contact reseller or Regional Manager	

Note: Currently, Laserfiche requires that the Laserfiche Server and the Laserfiche Full-Text Search be installed on the same physical machine. In the future, with Laserfiche 8.2, the Laserfiche Full-Text Search (LFFTS) will be able to be located on a separate server machine from the Laserfiche Server. When that becomes possible, you may want to plan for a separate machine to house the LFFTS, to optimize the resources available to each service.

2.3.3 MSSQL Server

When specifying hardware for your Microsoft SQL Server, we recommend allocating a separate physical machine. The first reason for this is to protect other component resources and confine MSSQL to only the resources you designate. The second is that internal Laserfiche tests have shown that MSSQL performs better on a physical box than on a virtual machine.

Physical machine requirements will vary depending on the Laserfiche load estimates, as well as the version of MSSQL Server. The MSSQL Standard version supports four sockets for processors and the maximum memory supported by the operating system. The Enterprise version—which you may want to use due to its support of failover clustering, a potentially important requirement for larger Laserfiche installations—supports more sockets and memory, up to the maximum number supported by the operating system.

The numbers required for the Laserfiche MSSQL database fall within the limits of the Standard and Enterprise MSSQL versions. For a load on MSSQL of 200 to 1,000 concurrent Laserfiche users (remember, this is actual users hitting the system, not

Laserfiche Certified Professional Content - Not For Distribution or Resale.

the number of licenses), we recommend a dual-processor, quad- or six-core configuration with 4GB to 16GB of RAM. For more than 1,000 concurrent users, we recommend multi-processor systems and more than 16GB of RAM. For more details or guidance on systems of that size, talk to your Laserfiche reseller or regional manager.

A best practice when planning and purchasing your system is to invest in multiple processors. Tests have been performed in-house in two environments: one with a SQL Server with a single Intel Xeon x5365 processor, and another exactly the same but with two processors. Search durations are similar in both systems if there's only one concurrent search. But if more users log in to the system to search at the same time, the system with more processors is faster.

Another best practice is to buy some of the fastest processors available, because searching, indexing, and full-table scans all benefit greatly from increased processor speed. In a test performed with an in-house system versus a production system configured by Laserfiche, we observed the difference in search result speed a single gigahertz of processor speed can make. Both systems had identical specifications, except that the production site had a 2GHz processor, and our test bed had a 3GHz processor. As the graph below indicates, the in-house test bed (red bars) returned the same search results in half the time of the production system (blue bars).

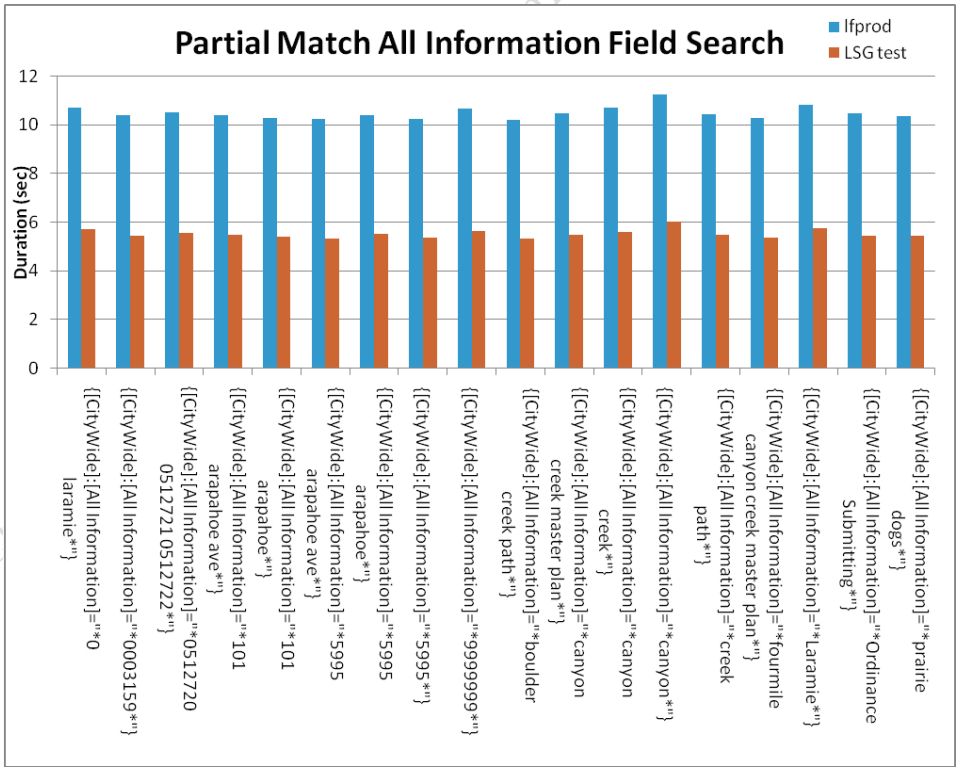

Another consideration with regard to your MSSQL server is how to handle multiple databases. In the Laserfiche suite, Workflow, Audit Trail, and the Laserfiche Client/Server will each require a separate database, though they don't all have to be

located on separate boxes. However, the Audit Trail database in particular will benefit from being separated onto its own machine, due to its reporting function, which is very resource-intensive. Some organizations with Laserfiche even replicate all Laserfiche audit data into a separate reporting environment to minimize interference from other system demands and to maximize performance.

> **Tip:** Through customer site tests and our own Laserfiche in-house tests, we have determined that buying the second-fastest processor available may actually be the best practice. The second-fastest processors can cost as much as 20-30% less than the fastest processors on the market, and the speed differential is typically not substantial enough to warrant the premium.

Some customers choose to utilize an existing MSSQL server installation for Laserfiche MSSQL needs. Up to a certain number of concurrent users and/or system load, this can be a cost-effective solution. We recommend a dedicated Laserfiche MSSQL installation for more than 50 concurrent users. Beyond that, we advise consistent monitoring of MSSQL usage and loads, to ensure enough resources are allocated for optimum Laserfiche performance.

2.3.4 Storage

The type of storage system connected to a Laserfiche implementation doesn't matter; whether it's a storage area network (SAN), direct attached storage (DAS), or network attached storage (NAS), Laserfiche can be configured to store volumes on any drive able to be referenced by a UNC storage path. Most enterprise systems, however, establish a SAN dedicated to storage needs.

The key metric for estimating storage capacity is how many images the system will house. As shown in the table below, Laserfiche recommends estimating 75GB per one million images, where each image is a single-page, 8.5x11, 300dpi, in black and white. The amount of storage space required will increase in ways not easy to predict when working with oversize and/or color images. Storage required for SQL data and log files can be estimated at a percentage of total volume storage.

Recommended configurations for storage arrays are also detailed in the table below. For volume storage, we recommend at least a RAID 5 array, specifically for its redundant data storage (parity), which employs at least three disks to protect against total failure of any one. Note that for RAID 5 arrays larger than 2TB, we recommend RAID 6, with two parity disks, due to increased possibility of a failed disk. For SQL data and log files, we recommend either a RAID 1 array, for its mirroring function, or a RAID 1+0 array, which distributes data storage and mirrors data across four or more disks.

Laserfiche Certified Professional Program. Not For Distribution or Resale.

	Sizing calculations	Storage arrays
Laserfiche volume storage	75GB per 1 million images (8.5x11, 300dpi, black/white)	RAID 5 Serial Attached SCSI recommended Serial ATA OK
SQL data file	10-15% of total volume storage	RAID 1 or RAID 1+0 Serial Attached SCSI recommended with fast drives
SQL log file	5-10% of total volume storage	RAID 1 or RAID 1+0 Serial Attached SCSI recommended with fast drives
tempdb	5% of total volume storage (for larger Laserfiche installations, usually >200 users)	Shared with SQL log file

It is worth mentioning that the ratio between the SQL data file and the transaction log can vary greatly between "simple" and "full" recovery modes. In simple mode, the transaction log will stay quite small, and as per the table above, we would correspondingly recommend nearer to 5% of total volume storage for the SQL log file. In full mode, however, the transaction log can become very large, depending on how often you perform backups, and we would recommend closer to 10% of estimated total volume storage.

2.3.5 Backup

Most organizations will have backup procedures in place, so planning for Laserfiche system backup has more to do with determining the potential amount of data that will need to be backed up, in addition to current load. If a backup system isn't robust enough for the additional load or isn't completely set up, Laserfiche recommends a disk-to-disk-to-tape process: backing up to disk first—which can usually be the SAN, NAS, or DAS device—and then backing up the disk to a rotating set of tapes. If storing on a DAS, we recommend a large RAID 5 array with a number of SATA disks, or at least enough capacity for one month's worth of grandfather-father-son rotation. Larger systems should use LTO-standard backup tapes; the most recent version is LTO-4, with 800GB uncompressed capacity. Smaller systems with less data to back up can use smaller-capacity DAT drives.

Tip: If backups are being performed at night, the backup server (if sufficiently large) can be shared with another low-resource server that is only used during the day, such as a VMWare Infrastructure Management server.

Laserfiche Certified Professional Program Not for Distribution or Resale.

2.3.6 Other Component Servers

2.3.6.1 Web Server

More cores can mean better performance from your Laserfiche Web Server because of the way ASP.NET (the .NET host for Web-based .NET applications) handles input/output threads in IIS. RAM usually isn't as important or as relevant as the number of cores, with the exception of support for certain processes, such as PDF generation (for printing or exporting).

Laserfiche Web configuration offers unique flexibility in the Laserfiche system. With the Laserfiche Web server, you can scale out and add more Web servers to handle more load, even after initial implementation, as the load on your Web server increases. To set up Web-based load balancing, you will need an application that will dynamically allocate requests across the Web, such as Microsoft Network Load Balancing.

With that said, we can make some basic estimations of hardware required for user load. For 200 concurrent users, we recommend 2-4 cores and 8GB of RAM; for 300-400 concurrent users, 4-12 cores and 16GB of RAM.

2.3.6.2 Workflow Server

Laserfiche Workflow is a difficult product to specify hardware for, because the resource use of Workflow isn't driven by users. Rather, it is driven by the number of actual workflows running, and the speed at which they run is limited by the number of threads that Windows Workflow Foundation can run.

Part of the load on the Laserfiche Workflow system certainly is predictable—we can reliably estimate how many users are configuring workflows and starting rules in the Workflow Designer. But this isn't the issue. Where Workflow impacts the system most from a planning perspective is when system actions are performed by workflow instances being triggered, which impacts the Laserfiche Server and SQL server. It's hard to know just how resource intensive an individual workflow might be, even when creating it, much less when planning for a future system.

That said, there are some guidelines to follow and some tips for maintaining Workflow performance. The first recommendation is virtualization, which would allow you to add (or remove) cores as necessary to your Workflow Server as you monitor production performance. The next recommendation is a faster processor with many cores, to take advantage of as many threads as the Windows Workflow Foundation can handle. We suggest starting with a 4-8 core system with 4-8GB of RAM.

Lastly, for Laserfiche Workflow in particular, we recommend making sure that your test systems are similar to your production systems, so that you have the best data possible for system load and resource allocation.

Laserfiche Certified Professional Program. Not For Distribution or Resale.

2.3.6.3 Quick Fields and Quick Fields Agent Servers

Though Quick Fields itself is a client-side application, it should be addressed in hardware plans because it will work the hardest of any of the Laserfiche applications, and it is very resource intensive. Faster processors will improve speed, particularly of popular processes like OCR that tend to be processor-intensive.

Quick Fields uses up to four threads, so we recommend a quad-core processor, and 4GB of RAM. For Quick Fields Agent, however, the baseline system you should start with is a dual quad-core with similar RAM, to take advantage of the fact that Quick Fields sessions can take advantage of parallel processing. Since the number of processors or cores can be adjusted as needed, you might want to consider virtualizing the Quick Fields Agent environment. For storage sizing, you'll want to start by allocating a minimum of 15GB in the VM for the operating system and applications, plus 50KB (single-page, 8.5x11, 300dpi, in black and white) times the number of images to be scanned across all sessions.

2.3.6.4 Audit Trail

Audit Trail is also tricky to specify hardware for, because the impacts to the system depend so much on which events are selected to be logged. When planning a Laserfiche implementation that will include Audit Trail, you need to know two key points about tracking and reporting on system events: one, Audit Trail event tracking can generate huge amounts of data, which can, in turn, compete for storage space with other applications, and two, Audit Trail reporting can be an extremely resource-intensive process.

As a best practice, if you know that Audit Trail will be an important part of your system, we recommend replicating audit logs into their own reporting system in a completely separate environment. Many customers have found that this is the most efficient means of maintaining performance across their Laserfiche installation.

2.4 Network

Network connections from SQL Server to the Laserfiche Server and from the Laserfiche Client to the Laserfiche Server are vitally important. The Client requests information from the Server, such as a particular document or a search. The Laserfiche Server then communicates with the SQL Server and requests the actual data. Once the Laserfiche Server has the data, it returns that information to the Client. All of this information exchange occurs on a network. If the network was a highway, the Laserfiche Server would be the traffic cop, shunting networking data from one location to another.

2.4.1 Port and Communication Considerations

The Laserfiche Server communicates with client applications using the HTTP/WebDAV protocol. In order to allow the Server and clients to interact with

Laserfiche Certified Professional Program – For Distribution Outside.

one another, you will need to ensure that certain TCP ports are open for these applications to allow communication to take place. By default, Laserfiche Server 8 listens on TCP port 80 or 5050. Laserfiche defaults to port 80, except when installed on the 32-bit edition of Windows XP, where Laserfiche defaults to port 5050. The Laserfiche Server broadcasts notifications on port 5051. Port 5053 also needs to be open on the Server machine for search engine functionality. If there is a firewall between your Laserfiche Server 8 instance and your client application, please make sure that ports 80 or 5050, and port 5051 are open on the firewall. You can use the Server Settings node of the Laserfiche Administration Console to modify the port settings.

2.4.2 Speed Considerations

Laserfiche has to run on a network that supports the TCP/IP suite of protocols. Connections between the Laserfiche Server and the SQL Server must be fast; a 100 Mbps switched connection is required at a minimum. Most networks these days have a 1Gbps switched connection on their backbone. For those installations running any of the Laserfiche Web products (Laserfiche WebLink and Laserfiche Web Access), the same network bandwidth requirements apply.

Scanning is one of the most network intensive processes in Laserfiche. In order to support scanning with devices at their rated speed, a 100Mpbs switched connection is recommended from the Laserfiche Client to the Server. Scanning across a WAN or over the Web using Web Access is supported, but is very dependent on the broadband access speed and the reliability of the connection. On commercially available Internet connections, scanning speeds of 4-20ppm are usually seen. Any sort of VPN or encrypted connection between source and destination will also result in a slight decrease in scanning speed because of the overhead to encrypt the connection.

2.4.3 Legacy Considerations

If you are using Laserfiche Server 8 with a Laserfiche 6 or 7 client application, you will need to make additional configurations.

Laserfiche 7

Laserfiche 7 used DCOM for communication between the Client and Server. If you are using a Laserfiche 7 client application with a Laserfiche 8 Server, the client applications will still be using DCOM. You must therefore either use the DCOM Configuration Utility to modify the default DCOM configuration, or use the Firewall Compatibility Component (FCC) to communicate across a firewall. If all of your Laserfiche components are within your firewall, you can install and run the DCOM Configuration Utility on the Server computer. If your security policy does not allow this type of configuration, or if you are using a firewall between your Server and your client applications (for example, between the Server and your WebLink Web server), you can install and configure FCC instead.

Laserfiche Certified Professional Program. No Portfolio distribution or reuse.

To access either the DCOM Configuration Utility or the Firewall Compatibility Component, install the Laserfiche Server 8 as a Custom installation. Choose either Laserfiche FW_DCOM (for the DCOM Configuration Utility) or Laserfiche FCC (for the Firewall Compatibility Component).

This configuration is not necessary if the client applications and the Server are all version 8.

Laserfiche 6

By default, Laserfiche client programs (version 6) use port 1888 for communication with the Laserfiche Server. If there is a firewall between your Laserfiche Server 8 instance and your version 6 client program, please make sure that port 1888 is open on the firewall.

2.4.4 Virtual Environments

Virtualization means simulating the interface of a server, a desktop computer, a storage device, an operating system, or even an application. Virtualizations are run on physical servers, computers, and storage devices, but the effect of virtualization is to hide the physical characteristics of the host machine.

There are many benefits to the practice of virtualization, but the fundamental difference between running an application in a physical environment (e.g., on a server machine) and running in a virtual environment is that in a virtual environment, you can have multiple virtual resources on one physical resource. In addition, you can also pool multiple physical resources to make them appear as a single virtual resource.

In practice, this means that if you estimate you require 20 servers to handle the needs of the organization, instead of purchasing 20 separate, physical server machines, you could purchase two or three larger servers and set up 20 or more virtual machines. What the virtual approach also gives you is high availability, scalability, portability, automatic failover, and increased flexibility for hardware upgrades, desktop management, and prioritizing resource allocation.

Some of the software products available for virtualization are VMware (used and recommended by Laserfiche), Microsoft Virtual PC, Microsoft Hyper-V, and XEN hypervisor (opensource).

2.4.4.1 Benefits

- Failover: virtualization management software can move virtual machines from one physical server to another ("hot standby") in the event of failure
- Resource efficiency: resources tend to be underutilized by up to 80%; with a VM you can be sure you're using the full capacity
- Easy hardware upgrades: move the VM to new physical machine with no reinstall of software needed (depends on implementation)

Laserfiche Certified Professional Program. Not For Distribution or Resale.

- Disaster recovery: in the event of complete physical server loss, with VM backup offsite, all you need to do is replace the physical server, reinstall the VM software, and restore the VM from backup; no need to reinstall multiple applications and configure them
- Increased security management: have users work on VMs on top of a secure base, which gives you the ability to manage potential security breaches, such as viruses or malware, by "wiping" and rebuilding the VM for the user in minutes, without having to rebuild all the software on the physical machine
- Dynamic allocation of resources: allocate more memory on physical machine to VM that will take more load
- Licensing savings: Windows Server 2008 Enterprise license for one physical machine or four virtual servers; Microsoft Data Center Edition server license for unlimited virtual servers on the same machine or machine cluster.
- Less expensive hardware and savings of physical space: three large machines instead of as many as 20 smaller ones
- High availability with failover and dynamic resource allocation
- Portability: move VM from one physical machine to another with **no interruption of service**
- Scalability: resize hard disks on the fly and use more of the SAN resources
- Save energy: save power by not keeping low-use physical machines running all the time
- Easier to manage: only working from one physical machine with the management software
- Run Windows-only software on a Macintosh through a virtual machine
- Software not competing for resources

2.4.4.2 Drawbacks

- Performance: in theory, software running on a VM will be slower than if installed directly on a physical machine—though it should be noted that in practice, products typically share space (and therefore compete for resources) on physical machines more often than they do on VMs
- Cost: need to invest in infrastructure, extra hardware, license fees, and VM software
- Consolidation and centralization: all your eggs in one basket if all servers fail

Laserfiche Certified Professional Program. Not For Distribution or Resale

2.4.4.3 Best Practices

While the needs and resources available at every installation will vary, Laserfiche offers a few general recommendations as the best practices for running Laserfiche in a virtualized environment.

First, we recommend running SQL on its own physical machine. The reasons are twofold: to confine SQL to only the resources you designate and because internal Laserfiche tests have shown that a higher level of SQL performance on a physical box than on a virtual machine.

Next, we recommend starting with hardware clustering of server machines with multi-core processors and lots of RAM. On those we recommend running VMware software, which has proven capable through extensive use both internally and in customer configurations. In addition, a SAN (Storage Area Network) is essential for the ability to move VMs while they are running. The SAN serves as a hard drive for everything attached to it and allows you to centrally manage disk usage.

> **Case Study:** The Laserfiche Information Technology department uses five servers and a SAN, each with dual quad-core processors and 48GB RAM, running VMware ESX server virtualization software. On that setup we are running some 170 virtual machines, everything from low usage test VMs to key infrastructure components, such as the enterprise mail server and domain services servers.

In addition, we recommend a modification to the standard file storage architecture of the Laserfiche Server. In Laserfiche 8, supporting files—such as audit logs—can be stored separately from the Laserfiche Server. When working in a virtual environment, we recommend storing data files outside of the virtual server on a network storage device. Specifically, these files are

- Audit logs (Laserfiche Server)
- Volume files (Laserfiche Server)

Separating the files from the system components will maximize the processing resources available to the software and improve performance, such as a configuration with the virtual servers installed on one physical server, and all supporting files stored on the storage area network.

2.4.5 High Availability Systems

What is it?

High Availability is a systems design paradigm that seeks to minimize system down-time in the event of component failure. It incorporates the concepts of fault tolerance (systems and software redundancy), failover clustering, and disaster recovery, creating a system that can quickly and automatically recover from component failure and minimizing downtime that arises from system maintenance/upgrades.

Laserfiche Certified Professional Program. Not For Distribution or Resale.

True High Availability systems can fail over and correct themselves within seconds of failure, before users even realize that there was an issue. The most common metric for High Availability is uptime as a percentage of total time per year. The standard definition of an HA system is one that can deliver five 9s service, meaning 99.999% uptime, and only 5.25 minutes per year of downtime. What High Availability really does is take the concept of redundant systems that you typically see on the hardware level, and extend to the system, application, and informational levels.

Pros and cons of it

High Availability systems have costs and benefits associated with them, just as with any significant system change. It is necessary for a requirements analysis to determine if a high availability solution is a good fit for your situation.

High Availability can provide reliability, little or no downtime, and better protection of critical information and mission critical applications. There is a considerable cost involved, complex system design, and long term maintenance involved in a High Availability solution as well. You will need to evaluate whether the application is mission critical, if the cost of downtime outweighs the cost of the extra hardware and software, and what is a realistic acceptable level of High Availability service.

How we do it?

Fault tolerance, virtualization, clustering, and offsite replication are how Laserfiche implements a High Availability solution.

High Availability and fault tolerance are all about redundancy. Redundant systems allow you to recover quickly from a single point of failure. One example is redundant power supplies in your servers. If one power supply goes down your server doesn't go down with it, since it's got a backup. Redundant systems can go beyond just simple fail-over, however. With redundant power supplies, you can also hook up the power supplies to separate UPCs or breakers or grids, etc. The point is that you now have limited n-way tolerance in the system, so that if there is a power failure that incapacitates one of your power supplies, you still have one connected to a separate UPC to keep your server operational.

There are two schools of thought when it comes to High Availability systems. The first depends on custom hardware and components that have been specifically designed to be highly available, and controlling hardware/software to manage these components. On the other side of the spectrum you find organizations using off-the-shelf commercial grade components to build redundant hardware and then using software to extend this to the system and informational levels. The most popular ways of implementing High Availability are Clustering, Virtualization, and Offsite replication.

Laserfiche Certified Professional Program. Not For Distribution or Resale.

2.4.5.1 Clustering

You can use Laserfiche software and the data stored in a Laserfiche repository in conjunction with failover clustering to maintain high availability. High availability means that users can still access data stored in a Laserfiche repository after the server experiences a hard disk, network, or software failure. You can achieve this by creating a server cluster that acts like one or more backup servers. If a failure is detected on the active server, the active server will be taken offline and another server will take over.

An administrator can decide to simulate a failure by manually taking a server off-line to either test failover clustering or to perform planned maintenance. Once testing, planned maintenance, or failure recovery is completed, the administrator can then put the server back online. A server cluster can be used to ensure that a Laserfiche repository remains available during planned or unplanned server downtime.

2.4.5.2 SQL Server

SQL Server provides support for failover clustering through virtual servers. These virtual servers can be installed and configured when installing SQL Server on a node of your failover cluster. You should install SQL Server after you have created and configured all of the desired nodes. For detailed information on planning and installing SQL Server support for failover clustering, please refer to SQL Server documentation.

After you have implemented SQL Server support for failover clustering, you will need to make the Laserfiche Server resource dependent on the cluster resources created for SQL Server. Detailed instructions for creating this dependency are available in the Laserfiche Server section below.

Note that this section is geared towards failover and high availability, but you can also set up MS SQL Server in a performance cluster.

2.4.5.3 Laserfiche Repository Location

The Laserfiche repository location is the default root folder for Laserfiche volumes and contains the search index file directory and audit logs. The location of this folder was specified when the Laserfiche repository was created. If a physical disk resource does not already exist for the hard disk on which the Laserfiche repository resides, you will need to create one. You will also need to make the Laserfiche Server resource dependent on the physical disk resource corresponding to the hard disk on which the Laserfiche repository resides. For information on how to make the Laserfiche Server dependent on another resource, please refer to the Laserfiche Server topic below.

Laserfiche Certified Professional Property - Not For Distribution or Resale.

2.4.5.4 Laserfiche Volumes

Each Laserfiche volume consists of images, text, and electronic files, as well as thumbnails and word location data. If a volume is not available, users who attempt to open documents in that volume will not be able to see the images, text and electronic files that are part of those documents, or thumbnails in the Thumbnail Pane. However, other information associated with the document, such as field information, tags, and versions, will still be available. A Laserfiche installation will continue to function after a Laserfiche volume experiences a hardware failure, although the contents of documents will not be accessible. It is therefore not strictly necessary to include the hard disks on which volume data resides as a part of a failover cluster. However, you can make the volumes available by including the hard disks on which they reside in your cluster.

If you decide to do this, we recommend that you do not make the Laserfiche Server dependent on it, unless the volume resides on the same storage device as the Laserfiche repository or SQL database files.

Note that, as stated previously, we recommend at least a RAID 5 array for volume storage, specifically for its redundant data storage (parity), which employs at least three disks to protect against total failure of any one. For RAID 5 arrays larger than 2TB, we recommend RAID 6, with two parity disks, due to increased possibility of a failed disk.

2.4.5.5 Laserfiche Server

The Laserfiche Server allows users to access and modify the data stored in one or more Laserfiche repositories; it is the only way to access the information in the repositories. It is therefore important to include the Laserfiche Server in your server cluster.

The Laserfiche Server software is cluster-unaware, and thus its behavior does not change when it is installed in a clustered environment. This does not mean that the Laserfiche Server cannot take advantage of clusters. If the Server fails over, users accessing it using remote clients may see a brief interruption of service. However, they will not need to log back in, and will not lose their modifications to open documents.

When implemented with failover clustering, the Laserfiche Server operates as an active/passive cluster, which means that only one instance of the Laserfiche Server can run at any given time. This means that only the Laserfiche Server on the node that is in control of the quorum will actively broadcast a Laserfiche repository. Instances of the Laserfiche Server installed on other cluster nodes will be inactive, or passive. These other instances of the Laserfiche Server will remain inactive until the currently active node fails over. During failover, the Laserfiche Server on the node that takes over the quorum will become active and the Laserfiche Server on the node that was taken offline will become passive.

Laserfiche Certified Professional Program - Not For Distribution or Resale.

Configuring your Laserfiche Server

After you have created and configured your server cluster, you will need to install the Laserfiche Server and LFFTS on each node. After which, you will need to configure both the Laserfiche Server and LFFTS to run in it. This can be performed from the configuration wizard in the Cluster Administrator on the active node. It is important that you choose or set the following settings:

- **Virtual server**: A virtual server is used by client workstations to access an application hosted on a server cluster. When configuring the Laserfiche Server, you should choose to use the existing virtual server called "Cluster Group."
- **Application cluster resource**: An application cluster resource is used to manage the settings of the Laserfiche Server. A "Generic Service" resource should be created for the Laserfiche Server and the Full-Text Search Engine.
- **Service name**: The name of the service used by the Laserfiche Server 8.1 is "LFS." The service used by the Laserfiche Full-Text Indexing and Search Service is "LfFTSrv." You can verify the service name by viewing the properties of the Laserfiche Server service.
- **Registry replication key**: The registry branch that keeps track of Laserfiche repository settings must be replicated on each node. This branch is the following: HKEY_LOCAL_MACHINE\SOFTWARE\LASERFICHE\Engine.
- **Resource dependencies**: It is important that the Laserfiche Server be dependent on the resources it needs to make a Laserfiche repository available across your network. Establishing a dependency relationship ensures the Laserfiche Server will always have access to the resources that it requires for proper operation. The Laserfiche Server should be dependent on the following resources: SQL Server, physical disk resource(s) for the hard disks on which the repository and SQL database files reside, cluster IP, and cluster name. Additionally, the Laserfiche Full-Text Search Engine must be dependent on the Laserfiche Server.
- **Restart policy**: The restart policy for the Laserfiche Server determines the number of times that the Cluster service will attempt to restart the Laserfiche Server when it detects that it is no longer started. If it is not able to start after the specified number of times, it will fail over the node. It is important to limit the number of restart attempts, since each attempt will increase the amount of downtime for a Laserfiche repository.

Laserfiche Certified Professional Program. Not for Distribution or Resale.

2.4.5.6 Creating/Registering Laserfiche Repositories

Once you have properly configured your server cluster, SQL Server, and Laserfiche, you are ready to create a repository or to register an existing one. When creating or registering a Laserfiche repository, keep the following in mind:

- The name of the server cluster should be used to identify the computer hosting the Laserfiche Server. This name was specified when creating the server cluster.
- The SQL Server instance is the name of the virtual SQL Server, followed by a backslash, and then the instance name. If an instance name was not specified, you may simply use the name of the virtual SQL Server.
- The Laserfiche repository location should be set to a folder on the cluster storage device.

Note: Once you have created or registered a Laserfiche repository, it will be immediately broadcast across your network. If the Scan Network for Repositories option has been turned off, or if for some other reason the Client cannot see the Laserfiche repository, it will need to be attached. When attaching a Laserfiche repository, you should make sure to use the name of the server cluster when setting the Laserfiche Server option. This will ensure the client will be able to access the Laserfiche repository after a failover occurs.

Clustering needs to automatically detect and handle faults. A two-node cluster has an active node and a passive node for application failover, and it is the only cluster that Laserfiche will recognize. The primary goal of clustering is availability, not performance. Laserfiche supports Microsoft Windows Server Enterprise.

2.4.5.7 Other Considerations

Virtualization utilizes purpose-specific software that allows for a simulated computer environment. It allows failover images to be spread over different resource pools, reduces scheduled down-time, and offers better scalability. Virtualization also allows for clustering and for the failover of entire VMs from one node to another.

Offsite replication is information and system redundancy. There is an automatic failover to an offsite location in the event of a catastrophe, like an active/passive cluster, except the offsite location is a separate and distant location to the main system. It requires double the hardware and software requirements, a reliable and robust network connection, replications of the necessary Laserfiche databases, volumes, and specific offsite failover software that will be able to take over in the case of a system failure.

Laserfiche Certified Professional Program. Not For Distribution or Resale.

3.0 Optimization

To ensure you get the most out of your Laserfiche system, there are a number of tips for making it easy to work in an efficient manner.

> **Note:** Be sure to watch the accompanying video on Optimizing Laserfiche.

3.1 Folder Structure

The goal of Laserfiche users is to find and use content in the Laserfiche repository. Typically users will do some combination of searching and browsing for information, but browsing can be difficult for users if the way information is organized is non-intuitive or haphazard. Disorganized folder structures and inconsistent or confusing naming conventions—for both folders and documents—can get in the way of employees using the Laserfiche system in the most efficient manner.

You can take advantage of the repository's folder tree and the ability to create a hierarchy of folders and subfolders. Breaking up your documents into multiple folders prevents the performance problems that can occur if you place thousands of documents in a single folder. Creating an intuitive folder tree also allows users to locate documents more easily.

> **Example:** If you have several types of document in a single folder—for instance, invoices, purchase orders, and receipts—you could create a subfolder for each document type. You could also create subfolders based on date or alphabetical by document name.

3.2 Managing Metadata

A key component of a user-friendly system is metadata that is easy to understand and work with. Attributes that are self-explanatory and that make it easier to find information ensure that users are comfortable working with the system.

3.2.1 What is Metadata?

In technical terms, metadata is information about information. In that sense, many elements of the Laserfiche repository can be referred to as metadata: the names of documents, the folder structure, the annotations, the security, and so on. However, when we speak of Metadata in the context of a Laserfiche repository, we mean a specific set of information. Metadata in that sense refers to user-defined information that is applied to the entry as a whole. By that definition, Laserfiche has five types of metadata: templates, fields, document relationships (also called links), tags, and versions. In this manual, "metadata" will be used to mean those five items as a group.

Laserfiche Certified Professional Program | Not for Distribution or Resale.

3.2.2 What is Metadata For?

Metadata stores any information you want to have connected with an entry. You can use metadata to help you organize and locate entries, to make entries more easily searchable, to provide additional information that would be useful to have alongside an entry, or to jump directly to related entries. You should use metadata to store information that will be useful to you in the future, either to find the entry, to find a related document, or to provide additional information when you have a document open.

3.2.3 General Principles of Metadata Use

3.2.3.1 Streamline Metadata Types

In general, it is a good idea to use as few metadata types as you can while still storing all the information you need. The reason for this is twofold. First, the more metadata elements you have in your repository, the more performance will be impacted. When you open a list of available fields, tags, or document relationships, the list will open much more quickly if it contains 30 entries than if it contains 300.

Second, and perhaps more important, users generally find it easier to select the appropriate element from a shorter list than a longer one. If a user is presented with 10 informational tags, they can scan the list quickly and choose the appropriate tag within seconds. If a user is presented with 50 informational tags, however, the process is more difficult. If the user knows what tag they want, they must still scroll through dozens of irrelevant tags to get to it, which is unwieldy and often frustrating. Worse, if the user doesn't know what tag to apply, they must skim all options and decide which of many tags seems most appropriate. In addition, if a user wants to search for items by tag, they will have the same decision-making process when confronting the list of 50 tags, whereas a list of five or 10 presents much less of a hurdle to searching.

In many cases, a single metadata element can serve more than one purpose, simplifying the process of choosing and applying them. For example, it's not necessary to have five different "Customer" fields just because you have five templates. If all of those fields will contain a client name, you can create a single "Customer" field and use it in all of the templates. In another case, you might have three different document relationships to indicate that one document is a supplement to another—for example, "Report" and "Supplement," "Agenda" and "Packet," and "Presentation" and "Handout." In some cases, it may be important for the way your organization handles searching to keep these separate, but in many cases they can be combined into a single document relationship: "Document" and "Supplement." This allows users to choose the relationship more easily, as they are confronted with fewer choices. It also makes searching simpler.

Laserfiche Certified Professional Program Not for Distribution or Resale

3.2.3.2 Use Appropriate Metadata Types

Certain kinds of information can be stored in more than one way, by more than one metadata type. In some cases, more than one metadata type is equally appropriate, but in most cases one type is better-suited to the task. In general, you should consider how the information will be set and used, and choose the type most suited to the task.

The following scenarios illustrate when you might want to use different types of metadata, including different types of fields and templates:

- Use an informational tag to mark if a document has been filed. The document could be tagged "To Be Filed" and then users could remove the tag once the document has been filed, or search for that tag to find all documents that need to be put in the appropriate folder.
- Use a list field to contain approval status of a form. Thus, when one of your users sets the document to Pending, Approved, or Denied, they simply have to select one of the options from the list—and the options will always be consistent and easy to search. You could also set up field constraints; for instance, you might want to ensure that all zip codes stored in the Zip Code field consist of five numbers. You can also specify that a field should be able to contain multiple values; for instance, your Phone Number field might be configured to permit multiple values so that you can store an office number and a cell phone number.
- Create a template with sets of fields you use regularly. For instance, you might want to create a character field to contain the address of the client associated with the form, a date field to contain the day you received the form, and a list field to track the approval status of the form. You could create a template for each of these forms—one for Invoices, one for New Accounts, and so on. You could then add fields relevant to those templates, or create new fields from within the template dialog. Some fields might be useful for both the Invoices and New Accounts templates, such as Date Received. If you add the field to both templates, you could search for a particular date and return documents belonging to both templates, or you could narrow down your search to only documents from one template or the other. Other fields might only belong to one of the templates—for instance, an Invoice Number field might only be applied to the Invoice template.

As a reminder, when creating a field, consider carefully which field type would be most suited to what you are trying to do.

3.2.4 Field and Template Design Considerations

Since fields and templates are one of the major ways many users locate documents in the repository, you should carefully consider your fields and templates before creating them.

Laserfiche Certified Professional Program. Not For Distribution or Resale.

3.2.4.1 Field and Template Creation Considerations

There are two important questions you should ask yourself when creating a field or template: what information users of this repository will use to search for documents, and how best to categorize documents within your repository. The first question is most useful when deciding on specific fields, as field values are an excellent way to search for entries. The second question is most useful when choosing what templates to include, as templates are a good way of categorizing documents and folders.

Field Creation Considerations

A field has two purposes: to provide searchable information to help users locate a document, and to provide supplemental information that the user might need once they have found it. Therefore, when deciding what fields to create (either independently or within a template), bear in mind what kind of information users will be looking for and what kind of information they will need once they have the document or folder open.

For instance, if you know that users will be searching for documents by author, it is a good idea to have an Author field that they can use for this search. If you know they will be referencing invoices by an invoice number, you could add an Invoice ID field of some kind. Once the user has the entry open, they might need different information about it: the same user who used an invoice ID to locate an invoice might then want to have the customer's phone number available so they can follow up. Including a Client Phone field would be useful for this purpose.

Keep in mind that fields can be applied to documents as part of a template or as an independent field. If only a few documents will need a particular field, you can create that field independently and simply apply it to the documents that need it.

Template Creation Considerations

Once you know how much and the type of information that should be stored as field data, you will need to decide on what type of template strategy your organization will use. Templates allow you to quickly add several related fields to a document; they are also useful for categorizing documents and allowing you to find all documents of a particular type. It is therefore common for organizations to create one template for each major document type, one template per department, or a combination of the two.

Whatever template strategy you choose, it is a good idea to minimize the number of templates as much as possible. Presenting a user with a long list of templates to choose from can be daunting; if some of those templates are only used by a few documents or folders, you may be better served by combining them into a single template. Remember that you can use independent fields to customize a more general template, which allows you to add exactly the fields you want to an entry without creating extra templates.

Laserfiche Certified Professional Program. Not For Distribution or Resale.

When choosing a template strategy, consider the following:

- What types or categories of document are in my repository?
- What information is important for each document type?
- Will users be able to tell which template should be applied to individual entries? (If users don't know which template to choose, they may have trouble filing or searching for entries.)
- Should the template apply to multiple types or categories of document? (If you will process many documents of a particular type, it may be a good idea to have a single template dedicated to that type of document; however, if you have only a few documents of a particular type—or if you prefer to organize your templates by department rather than by document—it may be more efficient to design one template that can be used for several purposes.)
- How many templates will I use? (Too many templates may be confusing for users, but using only one or two templates doesn't allow you to take advantage of templates for categorization.)

3.2.4.2 Field and Template Optimization

Once you know the type of information that should be stored in each field, and how you want to organize them into templates, you will want to optimize the manner in which users work with each field and template. This means making it more efficient to specify, view, and search for field data.

Field Names

The name assigned to a field is the first thing that a user will notice when working with field data. You should name your fields so that users can immediately tell what information should be stored in them. For instance, you might wish to create a field to store the shipping date for an invoice. If you name the field "Date," users may not know whether you mean the order date, filing date, or shipping date. Naming it "Shipping Date" would reduce confusion. However, if you wish to apply the same field to more than one template, you may wish to make the name general enough to apply to all the templates. If you want to keep track of the name of the customer for your invoices and applications, creating one field called "Customer Name" and applying it to both templates may be best.

Field Type

Another important setting is the field type. This setting determines the type of information that can be assigned to a field. An analysis of when you should use each field type is provided below.

- **Char**: If a field should be able to store any combination of numbers, characters, and symbols, you should use a character field, which is the most flexible type. You can set custom constraints on character fields as well—for

Laserfiche Certified Professional Program. Not For Distribution or Resale

instance, you could force users to input social security numbers in the format 123-45-6789. You will also need to set a width when creating a character field, and you should choose one that can contain the information users will store in that field. Note that you can increase the width later without problems, but decreasing the width later may result in data loss.

- **List:** If a field should only store values from a predefined list, you should use list fields, which enforce consistency. For instance, you could create a list field to store the state portion of an address and add all the U.S. states to that list. You would then know that you could search for "California" in that field, without having to worry that the user had input the state as "CA."
- **Number:** If users should to be able to input both decimal and integer values, but not other types of characters, use this field type. You can set custom constraints on numeric fields to ensure the values stored fall within a certain range.
- **Date:** If a field will only store dates that comply with short date format, use a date field. This field type allows users to take advantage of date range searches and the ability to search for a date from a single field.
  - **Date/Time**: If a field will only store dates that comply with short date format and you would also like to store the time, use a date/time field.
  - **Time:** If a field will only store a time, without the associated date, use a time field.
  - **Integer**: If a field will only store numbers that fall within 0 - 64,999, use an integer field type. This field should be used only if you do not wish to allow decimal numbers; for general-purpose numbers, use a number field. You can set custom constraints on numeric fields to ensure that the values stored fall within a certain range.
  - **Long Integer**: If a field will only store numbers that fall within 0 - 3,999,999,999, use an integer field type. This field should only be used if you do not wish to allow decimal numbers. For general-purpose numbers, use a number field. You can set custom constraints on numeric fields to ensure that the values stored fall within a certain range.
  - **Data:** The data field allows you to associate binary files with a field. This field type is most often used with integrations and other customizations.

Multi-Value Fields

You can configure fields to allow multiple values for a single entry. For instance, you might want to create a "Document Type" field to store the type of document for a field. If a document arrives that matches multiple types, you could create the field as usual and specify that it should be a multi-value field. List fields can also be configured to accept multiple values; in that case, users will be able to select values

Laserfiche Certified Professional Program. Not For Distribution or Resale

from more than one drop-down list. The document would then be returned in a search for either type. When creating your fields, consider if it would be useful for users to be able to set several values within the field.

Default Field Values

If a user does not have sufficient information to assign a value to a field, he or she may leave it blank. If blank fields are undesirable, you can configure a default value for the field in question. This causes a default value to be automatically assigned to that field, unless the user sets a different value. You can use tokens to create dynamic default fields—for instance, you might want to use the date token to automatically fill a "Filing Date" field with the current date.

Required Fields

To make sure a value is always specified for a particular field, you can indicate that it is required, which means a value is mandatory for that field when a user is setting field data. A required field is only required when it is applied as part of a template. The required setting has no effect on independent fields.

Template Field Order

In general, it is a good idea to place more commonly-used fields toward the top of the template, and less commonly-used fields toward the bottom. This allows users who are scanning and filing documents to quickly type values into the most relevant fields without having to tab or scroll down. It also presents the user who searches for or opens the document with the most relevant fields first. You should also bear in mind how users are used to finding information: if they have previously worked with Laserfiche templates, or if they are very familiar with forms that present information in a particular order, it may be more intuitive to mimic that familiar order.

Security

Security can be enforced on fields and templates. This feature is meant to prevent unauthorized users from viewing or setting field data, but it can also be used to make templates more relevant. If a user does not have rights to read the values of a field, they won't see the field at all. Thus, you can use field security to make even large templates manageable, by showing users only the fields that are relevant to them.

Optimizing Field Searches

Searching by the keywords that have been stored in a field can either be a very efficient way to find a document or it can lead to frustration. There are two major factors that determine the effectiveness of field searches. The first is the extent to which that field has been populated. The second is if the field values assigned to documents use consistent terminology and formatting.

Laserfiche Certified Professional Program - Not For Distribution or Resale

A field that has been applied to an entry is only useful if it has had a value assigned to it. If no value has been input, the entry will not be returned in field searches. It is therefore a good idea to ensure that fields are filled whenever possible. For instance, you may wish to set a default value on a field; this ensures that there is some value in the field at all times, while allowing users to input a different value if necessary. You can also set fields as required; when the field is applied as part of a template, users will need to fill in the field before they can save the document.

It is also important to ensure that terminology and formatting are consistent when filling in field values. Otherwise, it can be difficult for users to find documents even if the fields have been filled. For example, a user may describe the District of Columbia as any of the following: Washington D.C., D.C., DC, District of Columbia, etc. If a user needs to find all documents associated with Washington D.C., he or she may perform a search using any of those terms. If the user who filled the template field set it to "District of Columbia," but the user performing the search searches for "DC," the entry will not be returned by the search.

You can avoid this problem either by creating terminology policies for your repository, or by creating list fields to ensure consistency in the way a field is filled. If the field will contain a limited number of values that can be applied to multiple documents (e.g., a list of states, departments, personnel, etc.), use a predefined list. However, if the field will store unique values (e.g., addresses, phone numbers, etc.), don't use a predefined list, but create a set of criteria that will guide a user when determining which term or format should be used.

> **Note:** An administrator does not need to predict all possible values that can be assigned to a list field when setting it up. List fields can be modified in the future; see "Setting a List of Values That Can Be Applied to a Field" in the Laserfiche Administration help files for more information.

If the values will be unique in content but of a consistent format—for instance, phone numbers—you can use field constraints to ensure consistent formatting. Users may try to store phone numbers in a variety of formats: 1 (562) 988-1688, (562) 988-1688, 562-988-1688, 5629881688, 9881688, etc. If a user performs a field search using the wrong formatting convention, then the desired document will not be found. Field constraints define a pattern that a value must match before it can be assigned to a field. In the example above, you may require that users specify phone numbers using the following pattern: (###) ###-####. By requiring that field data match a particular pattern, users can be confident as to how to specify a value when searching that field.

3.2.5 Consolidating Templates and Fields

After your system has been in use for a while, you may find that there are two templates that can logically do the job of one or several fields that all hold very

Laserfiche Certified Professional Program. Not for Distribution or Resale.

similar information. This may become a problem for usability, where a user has to search for the same information more than once.

If you have multiple fields that can logically hold the same information, they can be merged together in the Administration Console. If you have two similar templates and you want to combine them, you can do so through Laserfiche Workflow, Quick Fields, or a Laserfiche SDK script.

3.2.5.1 Merging Fields

You can merge multiple fields into a single field after they have been created. This is useful if you have created fields in more than one template, realize that they store the same or similar information, and want to streamline them. When you merge fields, Laserfiche will create a new field and migrate all values from the merged fields into the new field. Laserfiche will then delete the original fields, leaving the repository with a single new field containing all original field values, allowing users to search and select fields more efficiently.

For instance, you might create two templates, one with a Client Phone Number field and one with a Customer Phone Number field. You might also independently create a Phone Number field and later decide to streamline your repository by merging these fields into a single Phone Number field. You could do so with field merging. The process will preserve the original values in the three fields, but consolidate them into a single field, making it easier to search and to select fields from a list.

You will need the **Modify Templates and Fields** privilege to merge fields.

To merge fields

1. Start the Laserfiche Administration Console.
2. In the console tree, expand the desired Laserfiche Server.
3. Select the desired Laserfiche repository.
4. If security has been enabled on that repository, log in as any user who has been granted the **Modify Templates and Fields** privilege.
5. Expand the Metadata Management node.
6. Select the Fields node.
7. Select the fields to merge.
8. From the **Action** menu, select **All Tasks** and **Merge Fields** to open the **Merge Fields** dialog.
9. Under **Name**, type the name of the new, merged field.
10. Under **Type**, select the type of the field. The type you select should be compatible with all of the fields to be merged. For instance, if you are merging

Laserfiche Certified Professional Program. Not for Distribution or Resale.

a Number and Long Integer field, you should select Number, as it will be able to handle both types of fields.

- If you have selected the List type, you can click the **Edit List** button to modify the list options. The new list will contain all list items from the previous fields.

11. Optionally, you can set a default value under **Default value**. Any existing default values for the selected fields will be automatically listed.
12. Select **Allow multiple values** to create the new field as a multi-value field. This will also allow you to store more than one existing value in the repository in the new, merged field.
13. Select **Required** to create the field as a required field.
14. Click the **Constraint** button to set a field constraint. See "Restricting Field Data to a Particular Format" in the Laserfiche Administration help files for more information.
15. If a single entry has more than one of the fields to be merged, choose how to handle them. Under **Conflict Resolution**, choose how to handle this conflict.
    - To save all existing values as a multi-value field, ensure that the **Allow multiple values** option has been selected and choose **Retain both values**.
    - To cancel the merge process if a conflict is encountered, select **Cancel merge and display error**.
    - To use one field's value and discard the other, select **Use field *Field* value**, where *Field* is the name of the field whose value you want to retain.

3.2.6 Making a Copy of a Metadata Object

You can create a copy of any metadata object—templates, fields, tags, document relationships, or public stamps—in the same repository or in a different repository, either by copying and pasting the object or by exporting the object and importing it at a later time. The copy of the metadata object will retain the same settings as the original user or group with the following exceptions:

- Only the metadata definition will be exported. The list of entries to which it was applied will not be included. For example, a field will retain all of its properties, but will not retain a list of documents to which it was applied, or the values stored in it for those documents.
- Security will not be included. For example, an exported template's template security will not be retained.

All other properties and settings will be included.

Laserfiche Certified Professional Program. Not For Distribution or Resale.

3.2.6.1 Copying and Pasting Metadata Objects

You can create a copy of any metadata object and paste it in the same repository or in a different repository.

Before copying a metadata object, make sure to log in to the Administration Console as a user that has sufficient permissions. For fields and templates, this is **Manage Templates and Fields**; for document relationships, it's **Manage Links**; for tags, it's **Manage Tags**, and for public stamps, it's **Manage Stamps**.

To copy and paste a metadata object

1. Start the Laserfiche Administration Console.
2. In the console tree, expand the desired Laserfiche Server item.
3. Select the desired Laserfiche repository.
4. If security has been enabled on that repository, log in as any user who has been granted the appropriate privileges for the specified repository.
5. Click the Metadata Management node.
6. Select the node for the type of object to copy.
7. Select the object or objects to copy.
8. From the **Action** menu, select **Copy**, or drag and drop the metadata object.
9. Do one of the following:
    - To copy the metadata object within the same repository, select **Paste** from the **Action** menu.
    - To copy the object to a different repository, repeat steps 2 - 8 for the appropriate repository before selecting **Paste** from the **Action** menu.

3.2.6.2 Importing and Exporting Metadata Objects

You can export and later import a copy of any metadata object in the same repository or in a different repository. The exported metadata objects are saved as an XML file. This allows you to copy metadata objects, to transfer metadata objects to a different repository, or to back up metadata objects. (For example, if you want to reorganize the templates in a field but want to be able to revert to your original template settings if you dislike the new ones, you could export your template before you begin. If you don't like your new template organization, you can delete the new template and then import the exported file.)

Before importing or exporting a metadata object, make sure to log in to the Administration Console as a user that has sufficient permissions. For fields and templates, this is **Manage Templates and Fields**; for document relationships, it's **Manage Links**; for tags, it's **Manage Tags**, and for public stamps, it's **Manage Stamps**.

Laserfiche Certified Professional Program. Not For Distribution or Resale.

To export a metadata object

1. Start the Laserfiche Administration Console.
2. In the console tree, expand the desired Laserfiche Server item.
3. Select the desired Laserfiche repository.
4. If security has been enabled on that repository, log in as any user who has been granted the appropriate privileges for the specified repository.
5. Expand Metadata Management.
6. Select the metadata object or objects to export:
   - To copy one or more individual metadata objects, open the relevant node and select the objects to export. (Hold down CTRL as you select to select more than one object.) Right-click, point to **All Tasks**, and select **Export as XML...**. Alternately, once you have selected the trustees, open the Action menu, point to **All Tasks**, and select **Export as XML...**. To copy objects of different types (for example, fields and tags, or document relationships and templates) you will need to perform this step once for each type of trustee.
   - To copy all of the objects of a particular type (e.g., all fields or all tags), select the relevant node, right-click, point to **All Tasks**, and select **Export as XML...**. Alternately, once you have selected the node, open the Action menu, point to **All Tasks**, and select **Export as XML...**.
7. You will be prompted for a location to save the exported XML file. Browse to a location and click **Save**.

To import a metadata object

1. Start the Laserfiche Administration Console.
2. In the console tree, expand the desired Laserfiche Server item.
3. Select the desired Laserfiche repository.
4. If security has been enabled on that repository, log in as any user who has been granted the appropriate privileges for the specified repository.
5. Right-click the repository name, point to **All Tasks**, and select **Import Repository Objects**.
6. Browse to the exported XML file.
7. In the Conflict Resolution option, select one of the following:

> **Note**: A metadata object is considered to match another metadata object if they share a same name and their properties match.

Laserfiche Certified Professional Training - Not for Distribution or Resale.

- **Create new objects only if matching objects do not exist:** This option will create a new metadata item if no matching item exists in the repository. If a matching item does exist, it will be left alone, and no new duplicate item will be created.
- **Always create new object**: This option will create a new metadata item for every item listed in the exported file. If a matching item already exists, another item will be created, and a number will be appended to its name. (For example, if two duplicates of the tag "Confidential" were imported, they would be created with the names "Confidential (2)" and "Confidential (3)".)
- **Only update existing objects (do not create new objects)**: This option will replace any matching items in the repository with the version in the exported file. It will not create any new items.
- **Update existing objects or create new objects**: This option will replace matching items with the version in the exported file, and will create new items if no matching items exist.

8. Click **Open** to import the metadata objects.

Laserfiche Certified Professional Program. Not for Distribution or Resale.

4.0 Performance Considerations

In addition to general usability considerations, the speed of the Laserfiche system will have an impact on the overall Laserfiche experience. This section highlights options in Laserfiche and MS SQL Server that will improve the performance of the Laserfiche system. The next section in the manual will cover diagnosing and correcting performance issues at a deeper level.

4.1 The Client

This section has tips for speeding up the user experience within the Laserfiche Client.

4.1.1 Folders and the Folder Tree

You can take advantage of the repository's folder tree and the ability to create a hierarchy of folders and subfolders. Breaking up your documents into multiple folders prevents the performance problems that can occur if you place thousands of documents in a single folder. Creating an intuitive folder tree also allows users to locate documents more easily.

It is strongly recommended that you avoid placing thousands of entries in a single folder. Retrieving, loading, and displaying all of those items can be very slow, delaying the user's ability to access the repository. You can avoid this problem by breaking up large folders into smaller folders. For instance, if you have several types of documents in a single folder—for instance, invoices, purchase orders, and receipts—you could create a subfolder for each document type. You could also create subfolders based on date or portions of the alphabet, by document names. You can use certain Laserfiche tools to automate folder creation. Laserfiche Scanning, Quick Fields, Import Agent, and Workflow can all file documents based on a token value—for instance, a date token—or even dynamically create a new folder. This allows you to break up documents without needing to manually create many new folders.

If you cannot break up a large folder into smaller folders, or if you need a workaround while you decide on and implement a plan for splitting up a large problem folder, you can use folder shortcuts to allow users to tunnel down to a relevant folder and bypass large folders higher in the folder tree. In this case, only users who need something in that large folder will actually need to open it; users who need something from a subfolder can jump directly to it.

4.1.2 Column Display

Certain columns in Laserfiche take more time to load and display in the Folder Browser than others. This is particularly noticeable in folders that contain many documents. Tailoring your column display to exclude information you don't need can cause these folders to open more quickly, especially if you can exclude those

Laserfiche Certified Professional Program. Not For Distribution or Resale.

columns that take a long time to load. You can modify the columns that are displayed on a user-by-user basis, through the Client; this allows each user to display just what they need to see. You can also set up default values and assign them to the Everyone group—via Attributes—so that there is a consistent starting point for all new users.

Most columns that you can choose to display load relatively quickly because the values are drawn directly from the SQL database. However, certain values are not stored and must be calculated whenever they are needed. Calculating these values can slow down the column display. Some of the columns with the highest potential for slowing performance are listed below:

- **Version**. This will slow as the number of versions in your repository increases.
- **Security.** The final results of a user's rights on a document once all security controls have been calculated.
- **Path**. This column is only available for search results and displays the folder path to the document. This must be calculated.

Certain Records Management columns can also slow display. For a complete list, please see Knowledge Base article 1011219 INFO: Certain Browser Columns Can Require Significant Loading Times.

4.1.3 Remembering Folder Columns

You can configure the default columns that are displayed in the Folder Browser for folders and documents. Thus, if you need information from potentially-slow folder columns like OCRed Pages in some folders but not others, you can display those slower columns only where they are needed. In the Options menu, select **Browser** and **General**. Under **Column display should be…**, select **The same for all folders** if you want a consistent column display, or **Remembered for each folder** if you want the ability to customize column display for each folder.

4.2 The Administration Console

4.2.1 Columns

As in the Laserfiche Client, you can choose which columns to display in the Laserfiche Administration Console. This functionality is particularly useful in the Volumes node. When you view the Volumes node in Details format, it will attempt to perform the volume size calculation for each volume. If you have very large volumes, or hundreds of volumes, this can take a long time and result in significant delay in opening this node. Changing the display time to List—or simply removing the problematic columns—resolves this problem.

Laserfiche Certified Professional Program - Not For Distribution or Resale.

4.2.2 Metadata

While metadata can make it easy for users to quickly locate documents they need, too many different metadata elements can both slow performance and make it difficult for users to determine which metadata item they should use. Whenever a list of metadata items is displayed—for instance, in the template selection list—the values must be retrieved from the server. If the number is very large (e.g., if there are hundreds of templates) this can be a very slow process. A smaller number of fields, templates, tags, and links will make the metadata dialogs load faster, and will also make searching faster and more efficient.

4.2.2.1 Character Field Width

The *character* field type lets you specify the width—in characters—that the field can be. In Laserfiche 8.1.1, an optimization was introduced to make searching in string fields significantly faster for fields of fewer than 100 characters. Therefore, the best practice is to define string fields to be shorter than 100 characters whenever possible.

4.2.3 Volumes

It is a good idea to find a balance between size and quantity of volumes, rather than having one or two very large volumes to store all of your files. Large volumes can cause problems when backing up and archiving, as they are time-consuming to export, attach, and copy and may not fit on standard backup media. Additionally, moving large amounts of data around can become unwieldy.

However, be careful not to overcompensate by creating hundreds of small volumes. These create more upkeep, make the volume selection list unwieldy, and make exporting and attaching volumes difficult.

A good compromise is to use logical, or "rollover," volumes. See the section below for details.

4.2.3.1 Logical Volumes

You can use logical volumes in Laserfiche to keep volumes from growing unmanageably large. Logical volumes appear to the user as one volume, but divide files between multiple sub-volumes, each of which has a size limit. You can configure the size limit to ensure that your sub-volumes remain small enough to avoid performance issues and fit onto your backup media. Whenever one of the sub-volumes reaches its size limit, another sub-volume is automatically created by the Server; however, users can continue to scan, import and search as if there were only one volume.

In effect, the logical volume acts as a container for the sub-volumes, which are invisible to the user through the Laserfiche interface but contain the physical files in their locations on the hard disk. The work of keeping the sub-volumes at a

Laserfiche Certified Professional Program - Not For Distribution or Resale

manageable size is done behind-the-scenes by the server. Using logical volumes in this way allows you to keep a relatively simple volume structure without allowing your volumes to grow too large.

When you configure your size limit for logical volumes, be aware that the volumes can grow after they have reached their size limit. This happens because Laserfiche does not split pages from a single document between multiple volumes. If an additional page is added to a document in a volume that has already "rolled over," an electronic document is printed using Snapshot, a document that was not OCRed prior to rollover is OCRed, or an electronic document has text extracted the size of the volume will increase a bit. Additionally, exported volumes contain XML files that have folder, metadata, and annotation information. Because of this, it is a good idea to start by setting the size limit of your logical volume to 20-50% of your final desired size. If you determine after several months of use that this is too conservative, you can always increase the maximum size later.

Laserfiche Certified Professional Program. Not For Distribution Outside.

4.2.3.2 Modifying Logical Volume Settings

The settings assigned to a logical volume can be modified after it has been created. Modifying logical volume settings does not affect existing physical volumes, but it will affect default name, paths, security, and/or file size limit for all future physical volumes created for this logical volume.

> **Note:** A physical volume can also be modified (e.g., location, size limit, or read-only status), but unless settings are modified on that physical volume's logical container, the settings will not be changed for future physical volumes.

The read-only option is a unique setting for a logical volume in that it prevents new physical volumes from being created. It is able to prevent volume creation by disallowing data from being assigned to the logical volume. When this option has been set, a user will not be allowed to add content to the logical volume. However, if a user has been granted the **Create Documents** or **Append Data** right on a physical volume associated with a read-only logical volume and the physical volume has not been specifically assigned read-only status, he or she can add data directly to the physical volume.

For the most part, physical volumes created for a particular logical volume are administered the same way as other volumes, though there are key distinctions to be aware of. For example, a user does not need the **Create Documents** right on the physical volume to migrate or assign a logical volume to a document, because Laserfiche assigns the appropriate physical volume to the document.

Another distinction is that new physical volumes are assigned the settings configured for the logical volume. Although these settings are assigned to the physical volume by default, they can always be changed without having a direct effect on the logical volume. However, changing them can affect a user's ability to

import to a logical volume. For example, if you restrict a user's rights to the current physical volume being used by a logical volume, that user will not be allowed to import to the logical volume until a new physical volume is created. Likewise, if the current physical volume is assigned read-only status, no users will be allowed to import to the logical volume.

4.2.3.3 Splitting Up an Existing Large Volume Into Logical Volumes

If you have been using the same volume or volumes for some time, especially if there is a large amount of scanning, your volumes may have become very large. When volumes become too large, you may experience difficulty exporting them; you may also begin to run out of hard drive space. This is why we recommend using logical volumes that will automatically and invisibly "roll over," creating a new volume behind the scenes but allowing users to continue to scan and import as usual. However, you cannot directly convert an existing non-logical volume into a logical volume.

If you have a volume that is too large, or you wish to preemptively switch to a logical volume to avoid future issues, you can instead migrate from an existing volume into a new, logical volume. The best practice recommendation for doing so is to perform a search for all documents in the current volume, sort them by creation date, and then migrate them in that order to the new logical volume. This will generate the logical volumes as if they had been there from the beginning, with the oldest documents in the first physical volume and so on.

Note that the process of migrating the documents may take a long time, especially if the original volume is quite large. You may wish to begin the migration at the end of the working day or on a Friday. This is also a good reason to perform this migration before you begin to see problems, as the migration is faster when the volume is smaller.

For more information, please see Laserfiche Volumes Introduction and Best Practices.

4.3 The Laserfiche Server and SQL Server

4.3.1 'Heavy Load' Actions

Certain actions produce a higher load on the Laserfiche Server than others. In many cases, it is possible to schedule these operations for off-peak hours, either by starting the processes at the beginning of an off-peak time (for instance, at the end of the day, or the beginning of the weekend), or by writing and scheduling Workflow or Toolkit scripts to handle them for you.

Laserfiche Certified Professional Program. Not for Distribution or Resale.

The following processes are potentially high-load, and, where possible, should be either broken up into small batches or run at off-peak hours.

- High volume scanning
- Indexing
- Migrating documents
- OCRing documents
- Briefcase import
- Volume attach
- Briefcase export
- Volume export
- Publishing with Plus

4.3.2 SQL Server and Laserfiche Server

The Laserfiche Server and your DBMS (Oracle or Microsoft SQL Server) can both be resource-intensive processes. This is particularly true if your Laserfiche Server installation has a high level of activity—for instance, many users connected and working at the same time, a high scan volume, frequent briefcase import or volume attaching, frequent audit reports, et cetera. If your Laserfiche Server and DBMS are on the same computer, the two services will compete for memory. This can in turn result in performance problems.

It is therefore recommended that you install your Laserfiche Server and DBMS on different computers. If you have already installed your Laserfiche Server and DBMS on the same computer, and wish to make this change—for instance, if you are beginning to run into performance issues—you can move either your Laserfiche Server or your DBMS. This will require you to take your repository offline—unregistered from Laserfiche and the database detached from MS SQL—and should be done only after verifying that a good backup of the system exists.

4.3.3 SQL Server

Your SQL Server performance is partially dependent on how many databases you have and how many operations are being performed in SQL Server at a given time on all databases. Thus, if you have a Laserfiche database that is not exposed to heavy load, but shares a SQL Server with a database that does have a heavy operation load, Laserfiche may run slower.

4.3.4 SQL Memory Usage

By default, Microsoft SQL Server will attempt to use all available memory. This can result in serious performance issues if Laserfiche is hosted on the same computer. In that situation it is strongly recommend that you cap MS SQL's memory usage to prevent it from consuming all available memory.

Laserfiche Certified Professional Program. Not Distribution or Resale.

4.3.5 Microsoft SQL Server: Express Edition

MS SQL Server Express is the free edition of MS SQL Server. As such, there are limits built into Express that are not present in the Core Editions of the product that will impact performance and storage capacity.

With regards to performance, SQL Server Express is limited to one CPU/core and 1 GB RAM. These limits mean that a Laserfiche system hosted by SQL Server Express should be limited to 10-15 users accessing the system simultaneously, so as not to experience performance degradation.

The capacity limits for the database size (.MDF file) are 10 GB for SQL Express 2008 R2 and 4 GB for earlier versions of Express. This will limit the amount of content that can be stored in the repository. Files are stored on the file server, not the DBMS, but the standard estimate is 1 million pages for the 4 GB database size and 2.5 million pages for the 10 GB database size limit.

4.3.6 SQL Indexes

Inserting or removing large amounts of data into a SQL database can fragment the SQL indexes. This can, in turn, slow SQL retrieval times significantly. To avoid this, you should develop a database maintenance plan that includes provisions for rebuilding your SQL indexes on a regular basis. See your SQL server documentation for more information on database maintenance.

4.3.7 Windows Defragmenter

Keeping the hard drive or drives that host your volume and database files defragmented can help prevent performance issues. You should regularly defragment the computers in question; third-party defragmentation scheduling software can help with this. You may need to stop your SQL Server before defragmenting to ensure that database files will defragment.

Laserfiche Certified Professional Program Not For Distribution or Resal

5.0 Performance Tuning

Performance tuning is the process of analyzing and maintaining your operating environment in order to keep the system running efficiently. This includes the hardware and software for the different components of your Laserfiche system (and other systems).

The most important element of performance tuning is striking a balance with regard to information you store and display and with regard to the process that you run on your repository. Careful consideration of what you need for your particular installation will help you strike this balance, which ensures that you can find the information you need as efficiently as possible.

While relevant and important information should be stored and displayed, you should eliminate any information that is irrelevant or unlikely to be used. A large quantity of extraneous information serves only to slow down your repository, especially if your repository is very large; limiting data only to what is actually needed ensures that users have what they need and encounter fewer performance issues. It is particularly important to note that, in some cases, even data a particular user can't see can affect performance. This section will help you determine how to cut down on extraneous information in your repository without losing important data.

You can also reduce performance issues by breaking up large collections of data, or lengthy processes, into smaller batches or steps. For instance, if a folder containing thousands of entries is slow to load, you might want to break it up into several smaller folders. Similarly, if you need to perform a slow process—for instance, migrating documents from one large volume to another—you can often break the process into smaller steps and take advantage of off-peak hours.

5.1 Bottlenecks

If you are encountering performance issues using your repository, it is useful to be able to figure out where the bottleneck in performance is occurring, which allows you to make changes to fix the problem itself, rather than guessing. You can use Laserfiche Auditing, Windows Performance Monitor, Laserfiche Server performance counters, and SQL performance counters to locate the source of the problem.

Once you have determined the source of a bottleneck—for instance, you see from Auditing that performance slows in the afternoon because several scanner operators are all scanning and storing large batches of images at that time—you can take steps to do something about it. You can distribute the load into multiple smaller batches and spread them out or automate the process at an off-peak hour. You can use other Laserfiche applications such as Quick Fields, Import Agent, or the SDK/Toolkit to help you automate these processes. Or you can upgrade your system by adding memory or disk space. Locating a bottleneck and resolving it can significantly improve the performance of your Laserfiche system.

Laserfiche Certified Professional Program. Not For Distribution or Resale

5.2 Performance Monitor

Note: Be sure to watch the accompanying video on Performance Monitoring in Laserfiche.

5.2.1 Getting Started with Microsoft Performance Monitor

Performance Monitor is a Microsoft utility that includes counters for monitoring different aspects of system performance. You can use these counters to measure the performance of various system features, across both hardware and software, and create performance logs. Performance logs can tell you the patterns of usage over days, weeks, or months, and help you locate and diagnose bottlenecks that may slow your system down.

In Performance Monitor, a counter is simply a chart of a single aspect of a specific system over time. This paper will focus on using various groups of counters—including Windows-, Laserfiche-, and SQL Server-specific counters—each of which can be further categorized based on performance object.

Windows-specific counters can be used to monitor performance objects related to various hardware and operating system features. For example, the **%Processor Time** counter is a Windows-specific counter that looks at the percentage of the processor in use and graphs it over time.

Laserfiche-specific counters are used to trace performance objects specific to your Laserfiche system, such as your Laserfiche Server, repositories, and more. The **Read-only sessions** counter is an example of a Laserfiche-specific counter: it tracks the number of read-only Laserfiche licenses connected to the Laserfiche Server and graphs that number over time.

5.2.2 Hardware Monitoring

One of the best ways to ensure that your system doesn't become overloaded is to monitor your hardware on a regular basis. Keeping regular performance logs will help you identify trends and establish a baseline for system use. Performance Monitor is an excellent tool for identifying hardware bottlenecks and determining where upgrades would be most useful.

You can run Performance Monitor as often as you choose; however, depending on your server load, you may need to run it more or less often to adequately monitor its performance. For a heavily used server, we recommend roughly once a week. For less trafficked servers, once a month or once a quarter is sufficient.

Performance Monitor logs are most useful when they include several days of data, which helps to ensure that the data accurately reflects normal use patterns and minimizes the impact of anomalous days on your data. These system load logs can serve as both troubleshooting tools and evidence to justify requests for hardware upgrades or additional servers.

Laserfiche Certified Professional Program. Not For Distribution or Resale.

When creating Performance Monitor logs, it is important to consider the interval at which you want to collect data. Generally, this decision is best informed by what purpose the log will serve. For example, if you are running Performance Monitor to gather baseline values for your system, sampling data at one minute intervals is sufficient. If you are logging information for a stress test, however, you should use a 3-5 second interval to collect more granular data.

For more information about how to create performance logs and set intervals for data sampling, see Creating Counter Logs below under the Setting Up Performance Monitor section.

Professional Copy. Not For Distribution or Resale.

5.2.2.1 Windows-Specific Counters

The tables below list some of the most useful and relevant Windows-specific counters, including descriptions of what each counter tracks and recommended threshold levels to measure your system's performance against.

Counter	**Avg. (or Current) Disk Queue Length**
Performance Object	**Physical Disk**
What it Tracks	The average (or current) number of both read and write requests waiting to be processed by the selected disk during the sample interval.
Recommended Threshold Levels	Ideally, **less than 2 times** the number of spindles making up the physical disk.

Counter	***%Disk Read Time / %Disk Write Time***
Performance Object	**Physical Disk**
What it Tracks	*The percentage of elapsed time that the selected disk was busy servicing read/write requests*
Recommended Threshold Levels	*For best performance, the Read and Write percentages* ***should not exceed 60% combined for more than 10 minutes at a time****.* ***Note:*** *Depending on what actions you are performing, the recommended 60% threshold may not apply. For example, if you perform a large scale batch import, you would expect (and want) the Disk Write levels to be as close to 100% as possible, while the Disk Read levels remained low. For large scale exports, the inverse would be true.* ***Note:*** *You can also measure the combined Disk Read/Write percentages using the* ***%Disk Time*** *counter.*

	Similarly, this counter alone should not exceed 60% for more than 10 minutes at a time.

Counter	**%Processor Time**
Performance Object	**Processor**
What it Tracks	The percentage of elapsed time that the processor spends executing a non-idle thread. It is calculated by measuring the time that the idle thread is active in the sample interval, and subtracting that time from interval duration. **Note:** Each processor has an idle thread that consumes cycles when no other threads are ready to run. This counter is the primary indicator of processor activity and displays the average percentage of busy time observed during the sample interval. It is calculated by monitoring the time that the service is inactive, and subtracting that value from 100%.
Recommended Threshold Levels	For best performance, this counter **should not exceed 80% for more than 10 minutes at a time.**

Counter	**Pages/sec**
Performance Object	**Memory**
What it Tracks	The rate at which pages are read from or written to disk. This counter is a primary indicator of the type of faults that can cause system-wide delays.
Recommended Threshold Levels	Take note of **sustained activity (10 pages/sec or higher) on the order of minutes**. A high count here requires further evaluation; however, it can sometimes indicate that you need more memory.

Counter	**Page Writes/sec**
Performance Object	**Memory**
What it Tracks	The rate at which pages are written to disk to free up space in the physical memory. **Note:** This counter shows write operations without regard to the number of pages written in each

Laserfiche Professional Program. For Distribution or Resale.

	operation.
Recommended Threshold Levels	When you evaluate **Page Writes/sec** values, it is important to consider them in the context of both **Page Reads/sec** and **Pages/sec** values. This will provide insight into what type of paging your system is performing, as well as indicate if the system is experiencing virtual memory problems, such as thrashing, that can cause frequent faults.

Counter	**Available Bytes**
Performance Object	**Memory**
What it Tracks	The amount of physical memory (in bytes) immediately available for system use or allocation to a process.
Recommended Threshold Levels	Avoid values **less than 4-10 MBs**.

Counter	**Processor Queue Length**
Performance Object	**System**
What it Tracks	The number of threads in the processor queue. Unlike the disk counters, this counter only shows ready threads, not threads that are running. **Note:** There is a single queue for processor time even for computers with multiple processors. If your computer has multiple processors, divide this value by the number of processors servicing the workload to measure against recommended threshold.
Recommended Threshold Levels	For best performance, this value **should not exceed 1-2 threads per processor for more than 10 minutes at a time**.

Laserfiche Professional Program. Not For Distribution or Resale.

5.2.3 Laserfiche Monitoring

With Laserfiche 7.1 and above, you can use Performance Monitor to look at specific aspects of your Laserfiche system. For instance, you might want to compare the number of read-write versus read-only licenses in use at different times to decide if you need to acquire additional licenses. You could also monitor your repositories to see how many documents or pages are being created per second to determine which repositories are the most heavily used.

5.2.3.1 Laserfiche-Specific Counters

Laserfiche uses four performance objects—Laserfiche Server, Laserfiche Repository, Laserfiche Full-Text Search Server, and Laserfiche Full-Text Search Catalog—each of which has its own set of counters to track its performance. The tables below outline the various counters associated with each performance object and what they track.

Laserfiche Server

Counter	What it Tracks
Retrieval Licenses	The number of retrieval licenses currently in use on the server
Full Licenses	The number of full-featured licenses currently in use on the server
Read-Only Sessions	The number of read-only sessions connected to the server
Read-Write Sessions	The number of full read-write sessions connected to the server
Mounted Repositories	The number of mounted (and available) repositories
Named User Licenses	The number of named user licenses currently in use
Public Portal Licenses	The number of public portal licenses currently in use
Laserfiche Repository	Repository-specific information related to Laserfiche

Laserfiche Repository

When you select a Laserfiche Repository counter, you can also select an instance. Each repository that you have created has its own instance; the **_Total** instance is a sum of all of these instances.

Counter	What it Tracks
Read-Only Sessions	The number of read-only sessions logged in to the repository
Read-Write Sessions	The number of full read-write sessions logged in to the repository
Total Sessions	The total number of sessions logged in to the repository
Pages Created/sec	The number of pages created per second for the repository
Documents Created/sec	The number of documents created per second for the

Laserfiche Certified Professional Program. Not For Distribution or Resale.

	repository
Entries Created/sec	The number of entries created per second for the repository
CLI Calls/sec	The number of CLI (Call Level Interface) calls per second for the repository **Note:** This counter scales with the number of calls made to the SQL database through the Laserfiche Server, allowing you to categorize what kind of work the Laserfiche Server is performing and identify when it is using the database extensively.
Database Connections	The number of database connections for the repository

Laserfiche Full-Text Search (LFFTS) Server

Counter	What it Tracks
Concurrent searches	The number of concurrent searches on the search engine
Documents indexed per second	The number of documents indexed per second on the search engine
Pages indexed per second	The number of pages indexed per second on the search engine
Words indexed per second	The number of words indexed per second on the search engine

Laserfiche Full-Text Search (LFFTS) Catalog

When you select an LFFTS Catalog counter, you can also select an instance. Each repository (and subsequent search catalog) that you have created has its own instance; the **_Total** instance is a sum of all of these instances.

Counter	What it Tracks
Concurrent searches	The number of searches in progress
Documents in the catalog	The number of documents in the LFFTS catalog
Documents indexed per second	The number of documents indexed per second
Pages indexed per second	The number of pages indexed per second
Terms in the catalog	The number of terms in the LFFTS catalog
Time since last indexed	The time (in seconds) that has elapsed since the catalog was last indexed

Laserfiche Certified Professional Program. Not for Distribution or Resale.

Time since last searched	The time (in seconds) that has elapsed since the catalog was last searched
Words indexed per second	The number of words indexed per second

5.2.3.2 Monitoring Laserfiche Workflow

Laserfiche Workflow is slightly different from the rest of the Laserfiche system in that it uses the Windows Workflow Foundation performance object, which comes with its own set of counters that can be used to track its performance, as well as the performance of Laserfiche Workflow.

Some of the most useful performance counters for Laserfiche Workflow, however, are obtained from Windows Messaging Queue, which must be installed manually before you install the Workflow Server and the Workflow Subscriber. For information on installing the Windows Messaging Queue component with Workflow, see the Laserfiche Workflow help files.

The tables below list some of the most useful counters provided by the Windows Messaging Queue performance objects (MSMQ Service and MSMQ Queue), including descriptions of what each counter tracks and recommended threshold levels to measure your system's performance against.

Windows Messaging Queue

Counter	**MSMQ Incoming Messages**
Performance Object	**MSMQ Service**
What it Tracks	The total number of incoming messages placed in queues on the selected computer by the Message Queuing service. This counter will give you an idea of the load placed on your Workflow server.
Recommended Threshold Levels	Watch for a **spiking number of incoming messages**, which may signal a runaway workflow. **Note:** The **Incoming Messages/sec**, which tracks the rate at which incoming messages are placed in queues, can be used in the same manner.

Counter	**Messages in Queue**
Performance Object	**MSMQ Queue**
What it Tracks	The total number of messages that currently reside in the selected queue.

Laserfiche Professional Program Slides. For Distribution on Request.

	This counter will give you an idea of how fast messages in a particular queue are being cleared out. **Note:** To track Workflow-related messages, select the appropriate Workflow queue from the instance column.
Recommended Threshold Levels	Watch for an **increasing number of messages** in the queue, which suggests Workflow is falling behind and may signal: • A runaway workflow, or • A need for more efficient starting rules

5.2.4 SQL Server Monitoring

You can also use Performance Monitor to look at specific aspects of your database management system, including SQL Server. With regard to the Laserfiche system, SQL Server (in this case, Microsoft SQL Server) is responsible for communicating with the Laserfiche Server and performing certain tasks that it generates.

Given its central role, SQL Server issues—such as improper maintenance or design—are often directly related to some of the most common Laserfiche performance issues, especially in the case of large scale installations.

One of the best ways to prevent the kinds of issues that can lead to performance degradation is to consistently track SQL Server using the various counters provided by Microsoft Performance Monitor. In this section, we will provide an overview of some of the most useful and relevant SQL Server-specific counters.

For more information about monitoring SQL Server activity, including additional performance objects and counters not discussed here, see the Microsoft Developer Network article Using SQL Server Objects.

5.2.4.1 SQL Server-Specific Counters

The tables below list some of the most vital SQL Server-specific counters, including descriptions of what each counter tracks and recommended threshold levels to measure your system's performance against.

Counter	**Batch Requests/sec**
Performance Object	**SQLServer: SQL Statistics**
What it Tracks	The rate at which SQL batch requests are received by the server.
Recommended Threshold Levels	Ideally, this figure should be **no less than 75% of your rated network speed**.

Certified Professional Forums | For Distribution or Resale.

Counter	**Buffer Cache Hit Ratio**
Performance Object	**SQLServer: Buffer Manager**
What it Tracks	The percentage of pages that the SQL server is able to find in the buffer pool without reading from disk.
Recommended Threshold Levels	For best performance, this figure should be **over 95%** most of the time.

Counter	**Full Scans/sec**
Performance Object	**SQLServer: Access Methods**
What it Tracks	The rate at which unrestricted full scans are being performed. **Note:** These scans can either be base table or full index scans.
Recommended Threshold Levels	Given that most queries generated by the Laserfiche Server are covered by indexes, this counter generally **should not exceed half of all queries on the Laserfiche database instance**.

Counter	**Index Searches/sec**
Performance Object	**SQLServer: Access Methods**
What it Tracks	The rate at which index searches are being performed. **Note:** Index searches are used to start range scans, single index record fetches, and reposition within an index.
Recommended Threshold Levels	Ideally, the value of this counter **should be in higher proportion of all searches than Full Scans/sec**. **Note:** This counter can be useful in determining what type of work the SQL server is performing when it is overloaded. For example, if the **Index Searches/sec** values are high while **Disk Queue Length** value for the database disk are also high, it indicates that your indexes may be fragmented and need to be rebuilt more often.

Professional Program - Not For Distribution or Resale.

Counter	Transactions/sec
Performance Object	SQLServer: Databases
What it Tracks	The rate at which transactions are started for the database.
Recommended Threshold Levels	In general, there is not a "recommended" level; however, this counter is important as a general metric for how much work SQL Server is performing.

Counter	Active Transactions
Performance Object	SQLServer: Databases
What it Tracks	The number of active update transactions for the database.
Recommended Threshold Levels	No recommended threshold; this counter is most useful for providing context for values from other SQL Server-specific counters.

Counter	Average Wait Time
Performance Object	SQLServer: Locks
What it Tracks	The average wait time (in milliseconds) for each lock request that resulted in a wait.
Recommended Threshold Levels	No recommended threshold; this counter is most useful for providing context for values from other SQL Server-specific counters.

Laserfiche Certified Professional Program. Not for Distribution or Resale.

5.2.5 Setting Up Performance Monitor

Before you can start using Performance Monitor to track your system's performance, you must configure it to include the various counters you want to use. Below you will find steps for adding counters and creating new performance logs.

To open Performance Monitor and start configuring, select **Run** from your desktop's **Start** menu, type *perfmon*, and click **OK**.

5.2.5.1 Adding Counters

By default, Performance Monitor will display counters measuring aspects of memory (**Pages/sec**), hard disk (**Avg. Disk Queue Length**), and processor load (**%Processor Time**). Depending on your specific performance monitoring goals, you may need to add additional counters.

To add a counter to Performance Monitor:

1. With the **System Monitor** node selected, click the plus sign icon (highlighted below) in the toolbar.

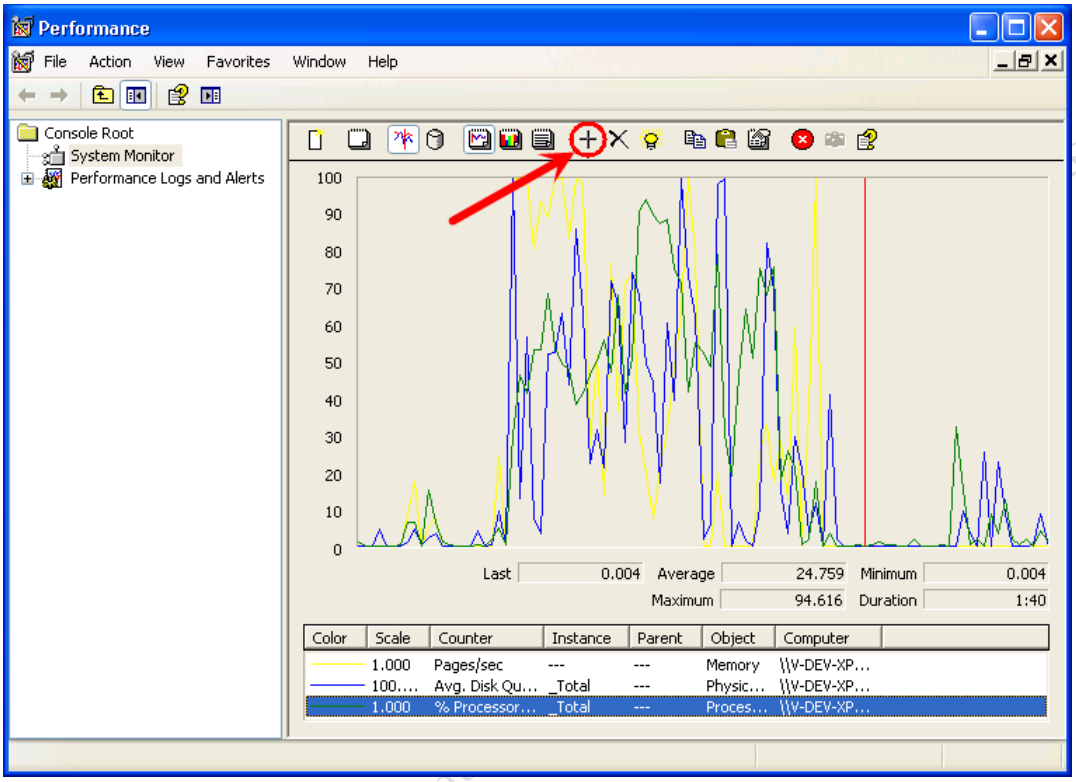

2. In the **Add Counters** dialog box, select the performance object you want to track from the **Performance Object** drop-down menu.

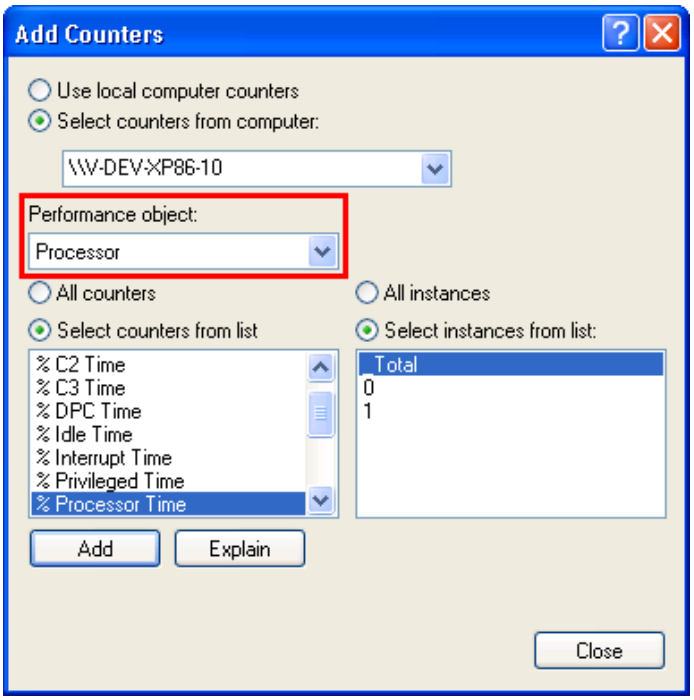

Note: When you select a particular performance object from the drop-down menu, the counters specific to that object will be displayed.

To help you find the counters outlined in this paper, each counter has been listed with the name of the performance objects under which you can find it.

Laserfiche Certified Professional

3. Select a counter from the list. To select more than one counter at a time, hold the CRTL key while selecting your counters. To select all available counters listed, select the **All counters** radio button.

   Tip: Click the "Explain" button to view a description of what each counter monitors.

4. **Optional:** Select an instance for the counter.

   Note: Instances are used for counters that could apply to multiple items. For example, if you select a counter from the Laserfiche Repositories performance object, the list of instances will include all the repositories available on the current computer.

   a. To monitor one instance, select that instance.

   b. To monitor more than one instance using a separate counter for each instance, press the CTRL key and select multiple instances.

   c. To monitor all instances at once using a single counter, select **_Total**.

5. Click **Add**.

6. Repeat steps 2-5 as necessary until you have added all the counters you want.

   Tip: Don't add too many counters; otherwise, they will be difficult to tell apart.

7. Click **Close**.

5.2.5.2 Creating Counter Logs

Counter logs allow you to collect and record performance data from your computer based on a range of performance objects, counters, and object instances. The information contained in your logs can help you establish usage patterns and diagnose bottlenecks that may be affecting your system.

To create a new counter log:

1. Expand Performance Logs and Alerts in the left-hand pane. Click on Counter Logs.

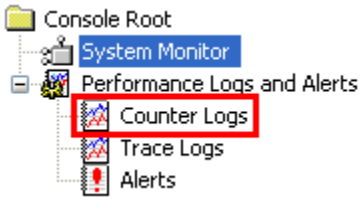

Laserfiche Certified Professional Program. Not For Distribution or Resale.

2. Right-click the **Counter Logs** node and select **New Log Settings**.

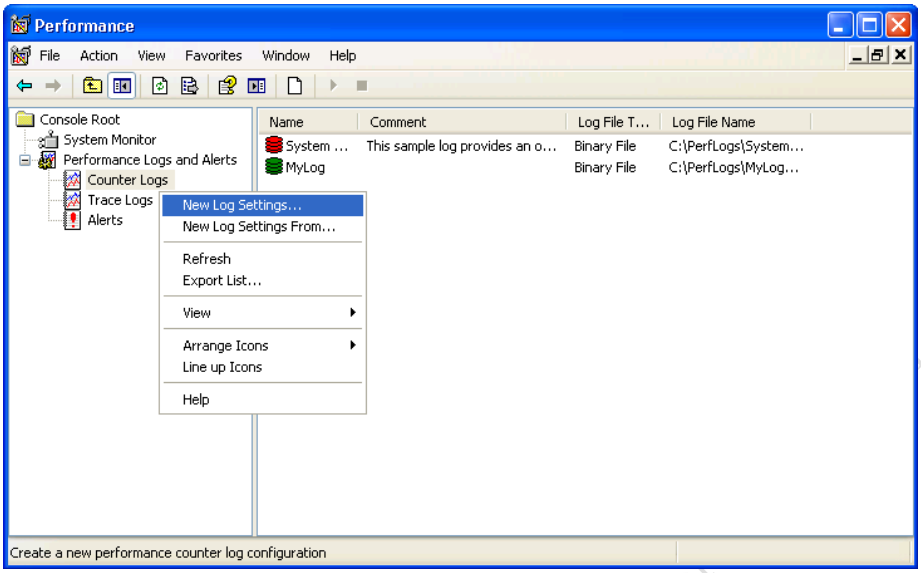

3. In the **New Log Settings** dialog box, specify a name for your log.
4. In the **Properties** dialog box, click **Add Counters** and the same steps as above for adding counters. To remove a particular counter, select it and click **Remove**.

> **Note:** To add all counters for a specific performance object, click **Add Objects** and select the ones you want to add.

5. Set the interval at which you want data to be sampled using the **Interval** and **Units** fields.
6. Click **OK**.

By default, Performance Monitor will begin adding data to the log as soon as you finish creating your log.

> **Note:** You can also configure additional log settings, such as log file type, comments, and schedules for gathering logs, using the **Log Files** and **Schedule** tabs in the Properties dialog box. For more information, see the Performance Monitor help files.

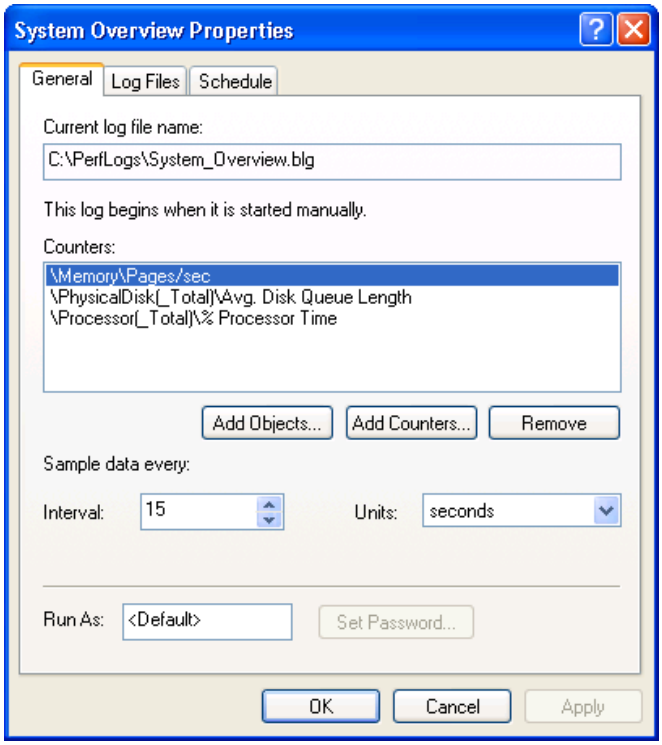

Laserfiche Certified Professional Program - Not For Distribution or Resale.

5.3 Laserfiche Search

Recently, Laserfiche was asked to help improve the speed of searching for one of our customers. Using the knowledge and methods mentioned previously in this manual, the problem spots were diagnosed and corrected to improve the speed of searches.

5.3.1 Laserfiche Search Performance

In general, the time it takes to get search results isn't dependent on the search process itself. Instead, if there is a delay in retrieving search results, it is probably due to the time it takes to load and parse the results themselves. Thus, even though a full-text search returning ten results and a full-text search returning five thousand results both have a search process that searches the entire repository, the latter will take significantly longer to complete because the Client must display all five thousand results. Additionally, the Client must evaluate the security rights for each document to determine if the user can see it or not, prior to displaying the search results.

The best way to speed up a search is to return fewer results. You can do so by making your search more exact. For example, if you know that your repository contains many different kinds of reports, and you want a finance report specifically, search for "Finance Report" instead of simply "Report." The search will return results faster, and as an added bonus, you have fewer results to sort through yourself. You can also make a search more precise by adding a constraint: if you know roughly when the document was created, you can add a Creation Date search to reduce the number of documents that will be returned. Refining your search in this way puts slightly more load on the server, but if it significantly reduces the number of results returned, the tradeoff is frequently worth it.

Certain search types are slower than others. Some search types require the server to perform calculations on the fly, which can slow down searching. These search types are the Has Pages search and the following Records Management searches: Vital Record, Cutoff and Disposition. Additionally, the Path column in Search may be slow to load.

Fuzzy searches are useful for helping users account for typos, or when they are unsure how a search term should be spelled. However, when you set the fuzzy search criteria, use the lowest percentage of the word or number of incorrect letters possible. As you increase the "fuzziness" of the search, the time to run the search increases, the number of results increases, and the time to display the results increases. This results both in slow performance and in many irrelevant results being returned.

A full-text search uses search index files to locate documents. These search index files can fragment, just like your computer's hard drive. Re-indexing the repository occasionally will recreate these files and reduce search slowness due to fragmentation. The frequency with which you will need to do this depends partly on

Laserfiche Certified Professional Program | Not for Distribution or Resale.

the size of your repository and your scan or import volume. Re-indexing every month or two is sufficient for most installations, although very high-volume repositories may benefit from biweekly or even weekly re-indexing. You can re-index at off-peak hours to reduce the server load.

5.3.2 Architecture Overview

Searching in Laserfiche has several elements, which are handled separately by different server components. An understanding of these pieces and how they relate informs our approach to troubleshooting and best practices for optimization.

- **Laserfiche Full-Text Indexing and Search Service (LFFTS):** An independent service that handles full-text searching of the repository.
- **SQL Search:** Retrieves searches for things other than full-text, such as entry names, creation dates, field values, annotation, tags, etc., that are stored in the SQL database.
- **Laserfiche Server** The Laserfiche Server merges the search results, determines which columns to display, and checks that the user has the rights to the data being returned.

5.3.3 Best Practices

The following best practice recommendations will keep search times lower so that users can find information more quickly.

5.3.3.1 Optimize Laserfiche Client Settings for Search Performance

Laserfiche Client settings can be customized according to a variety of needs. Of course, given that a particular organization or user may have multiple needs, these may have to be balanced. If the goal is the fastest possible search performance, the following settings are recommended:

5.3.3.2 Make Searches Faster

- Consider adjusting the **Stop Words** list in **Indexing Properties** in the Laserfiche Administration Console to exclude frequently searched words that do not add meaning when a full-text search is performed.
- In Laserfiche 8.1.1 or higher, optimization was introduced to make searching in string fields significantly faster than before. Therefore, the best practice is to define string fields to be shorter than 100 characters whenever possible.

> **Example:** In a Laserfiche test environment, exact-match cross-field searches with short search phrases improved from taking 0.11 seconds to only 0.02 seconds (an 80% improvement) after the optimization.

Laserfiche Certified Professional Program. Not For Distribution or Resale.

- Disable the **Fuzzy Search** option by default, and enable it only when necessary. A quick way to find out if anyone has this option enabled is to execute the following SELECT query in the SQL database.

```
select * from trustee_attr where attr_name = '[SEARCH]Option'
and attr_val not like '0'
```

  If you do use fuzzy search, note that the greater the level of fuzziness, the more results will be returned, and the less likely they are to be relevant to the query. Setting the fuzziness to a lower percentage or smaller number of letters will return fewer, more relevant results that can be displayed faster.

- Disable the **Include Shortcuts in Search Results** option. A quick way to find out if anyone has this option enabled is to execute the following SELECT query in the SQL database.

```
select * from trustee_attr where attr_name = '[SEARCH]IncludeShortcuts'
and attr_val like 'Yes'
```

5.3.3.3 Have Results Returned Faster

- Folder filter expressions are a newer security concept that are not being used that much. Give the **Bypass Filter Expressions** privilege to the **Everyone** Group if no filter expression is used in the repository.
- Give **Bypass Browse** privilege to administrators who should have access to all entries in the repository or already have the **Manage Entry Access** privilege.
- Make entry access security configuration as simple as possible.
- Do not display resource-intensive columns. As of Laserfiche 8.1.1, these are **Path** and **Security**. These settings are stored in the system as user attributes and apply to both the thick and the Web clients. They can be changed by the individual users or, in version 8.0 and above, centrally managed by an administrator.

5.3.4 Design a System for Best Search Performance

It is a best practice to place the full-text search catalog on the fastest hard drive possible, and separate from the drive where Laserfiche volumes are stored.

The full-text search catalog must remain local to the Laserfiche Server, but it is recommended that the volumes be moved to a network storage device. Doing so may affect the speed of indexing, however, which will make it even more important to perform the indexing after hours, as described below.

5.3.5 Perform Full-Text Indexing at a Convenient Time

Indexing can be a resource-intensive process. A best practice is to avoid full-text indexing while searching, as much as possible. For example:

Laserfiche Certified Professional Program. Not For Distribution or Resale.

- **Plan heavy scanning and document creation during non-production hours.**
- **Scheduled indexing**: Scheduled indexing is a new feature offered in Laserfiche version 8.1.1 or higher. It allows administrators to schedule indexing to exclude certain peak business hours. Note, however, that if this is done, files imported during the excluded time will not be returned in full-text searches until they have been indexed.

5.3.6 Monitor the System for Heavy Load

As discussed previously, you can monitor key performance counters on the servers and make sure they do not deviate from the baseline.

- **Check the Indexing Load:** You can determine if indexing is the cause of slow full-text searches by checking the index queue size in the Laserfiche Administration Console as well as the server performance counters related to indexing.
- **Monitor Concurrent Search Load**: A best practice is to periodically monitor the number of concurrent searches using the **Concurrent Searches** counter. If the number is frequently high, it may be time to perform a search performance benchmark test with multiple concurrent searches and analyze performance bottlenecks in such conditions.
- **Monitor Load from Other Applications Hosted on the Same SQL Server:** If the production SQL server also hosts other databases than the Laserfiche production database, activities in all the hosted databases can be checked at random times in the Activity Monitor in SQL Server Management Studio. If the load becomes too heavy, a dedicated SQL server should be used to host the Laserfiche database.
- **Check for Busy Network Traffic:** Compare the network performance counters to the **Current Bandwidth** counter for the Laserfiche, SQL, and Web servers to see if there is a network bottleneck.

5.3.7 Perform System Maintenance

Monitoring and periodically defragmenting key components of the system can improve search performance.

- **Check SQL indexes:** Rebuild if fragmentation is greater than 40%. Reorganize if fragmentation is between 10% and 40%.
- **Monitor key SQL tables in the Laserfiche database:** Monitor tables such as *toc*, *doc*, and *propval*. Defragment if necessary.
- **Monitor full-text search performance:** After reindexing, the search performance should be regularly monitored.

Laserfiche Certified Professional Program. Not for Distribution or Resale.

- **Remove full-text search index fragmentation:** Search index fragmentation can be a factor in search performance. One tool available to perform this defragmentation is the Laserfiche 8 Quick Reindex utility (QRcmd.exe), a free component that can be installed with Laserfiche Server 8. It quickly regenerates the search index files for a Laserfiche repository.

5.3.8 Troubleshoot Hardware Performance if Necessary

In most cases of reduced search performance, after checking the items discussed above, we should have a good idea of the cause of the slowdown. If necessary, we can troubleshoot on the hardware resources on the SQL Server and Laserfiche Server. This step is not recommended at the beginning of the investigation. Higher level issues, as listed above, should be exhausted first. For example, if there is a missing index in the database, it may cause the SQL Server to run full-table scans and appear to have an Input/Output (I/O) bottleneck.

If it is suspected to be a SQL Server performance issue, there is a good troubleshooting guide, *Troubleshooting Performance Problems in SQL Server 2008,* at http://download.microsoft.com/download/D/B/D/DBDE7972-1EB9-470A-BA18-58849DB3EB3B/TShootPerfProbs2008.docx. Most of the same performance counters and analysis used in troubleshooting SQL Server performance can also be applied to the identifying Laserfiche Server performance bottlenecks.

5.3.9 Keep Up to Date on Laserfiche Releases and Fixes

Laserfiche periodically releases Server and Client patch updates and version updates that include performance optimizations. It is recommended to periodically check the Laserfiche Support site at https://support.laserfiche.com for announcements of product hotfixes and releases. Current release and patch information is always available at https://support.laserfiche.com/KB/1012113.

5.3.10 Optimize Search Types

Laserfiche Advanced Audit Trail can be used to monitor searches and their duration. Searches using overly broad criteria may return more search results than needed and slow down the results of others' searches.

Searches can be directly specified by individual users or aided by integrations and customizations that use Laserfiche advanced search syntax. As much as possible, users should be trained to conduct effective searches, which will help them better find the information they are looking for as well as reducing the load on the system.

- When conducting a full-text search, specifying more detail will return more relevant results in a shorter time.
- An "AND" Search will return fewer and more specific results than an "OR" search.

Laserfiche Certified Professional Program. Not for Distribution or Resale.

- For field searches, searching within a single field is much faster than searching across all fields.

5.3.10.1 Adjust Customizations to Use the Most Efficient Search Syntax

Integrations and other customizations that incorporate Laserfiche search syntax should be adjusted to make sure their code reflects the suggestions provided above. Any other operations that also use advanced search syntax rather than the search interface should also follow these recommendations.

5.3.11 Adjust Customizations to Display Results Efficiently

Optimizing customizations to use the most efficient methods of retrieving and displaying search results can improve the speed of many searches, thereby potentially improving performance throughout the repository.

5.3.12 Defragment Search Index

Production tests of Laserfiche search catalogs have revealed fragmentation levels as high as 18.42%—which was reduced to zero after reindexing using the Quick Reindex Utility. In the same system, the average full-text search duration (for search phrases selected for test purposes) dropped from 13 seconds to 5 seconds, showing performance 2.4 times as fast. The search catalog will automatically undergo an optimization procedure if it reaches a certain level of fragmentation, but you can also reindex manually if you are finding that searching is slower than expected.

5.3.13 Use Faster CPU for SQL Server

Measurements were taken of MS SQL Server's **%Processor Time per CPU** during a field search using Microsoft Performance Monitor. The only difference between the two SQL Servers was the speed and type of CPU.

	Production SQL Server	**Test SQL Server**
Processor Name	Intel Xeon E5405	Intel Xeon x5365
Speed	2GHz	3GHz
# of Processors	2	2
# of Cores/processor	4	4

On the production server, **%Processor Time per CPU** surged up to 100% for all eight CPU cores, indicating a CPU bottleneck on the SQL Server during field searches. On the Laserfiche test SQL Server, however, **%Processor Time per CPU** went up to 100% for only one CPU core, but stayed below 65% for the other seven CPU cores. Compared to the production SQL Server, the Laserfiche test SQL Server was much less constrained by the CPU bottleneck, thanks to the faster processors.

Laserfiche Certified Professional Program. Not For Distribution or Resale.

In fact, a comparison of search performance test results conducted in the production system and the Laserfiche test system showed that field searches were much faster in the test system; the average field search duration in the Laserfiche test system was half of that in the production system. It was recommended that a higher speed CPU be used for the production SQL server.

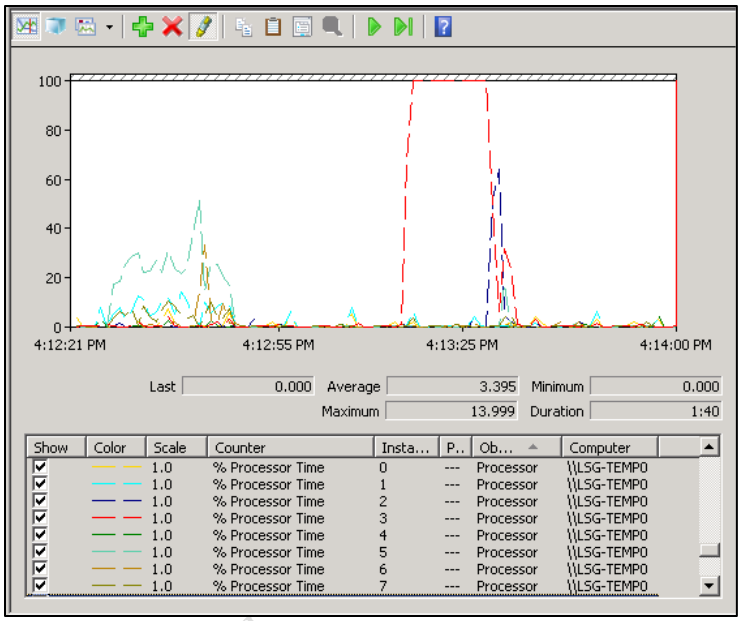

Laserfiche Certified Professional Program. Not For

6.0 Troubleshooting

Most problems with Laserfiche can be diagnosed and corrected without the need to contact Laserfiche Technical Support. Even if you do decide to contact Technical Support, knowing common troubleshooting procedures will often speed the problem resolution process along.

Regardless of the problem you are experiencing, follow the steps outlined below to try to resolve it:

1. Try performing the action again, as the problem may have been transient. If it resolves itself after trying again but occasionally reappears, you may want to continue troubleshooting.
2. Restart the Laserfiche program and try again. Sometimes restarting the Laserfiche Server service will help.
3. In some cases, rebooting the computer will fix the problem. Sometimes logging out of Windows and logging back in will be sufficient. If rebooting fixes the problem, the fix may often turn out to be temporary and the problem may manifest itself again later.
4. If there is an error message or code, try looking up the error in the Laserfiche Knowledge Base. If there is no error message or code, look in the Windows Event Viewer on the Laserfiche Server machine, as there may be an error logged there. Otherwise, search the Laserfiche Knowledge Base using relevant keywords—don't give up after a single search, but try several different keywords.
5. Search the Laserfiche Support Forums for posts identifying a similar problem. You can also ask questions there after performing your own troubleshooting.

   Note: The forums are not appropriate for complicated cases or cases requiring quick turnaround. Different people monitor the forums on an irregular but frequent basis without any guarantee of assistance.

6. Try repairing your Laserfiche installations.
7. Check if this manual has a specific series of troubleshooting steps to follow for your problem. If so, follow it; otherwise use a general troubleshooting strategy.

If you are a user or administrator of the Laserfiche software, your reseller might have given you instructions on what to do if you encounter a problem when using Laserfiche. Your reseller's instructions should take precedence over the advice given in this paper, as they are your support contact and will know best what to do in your particular situation.

Laserfiche Certified Professional Program. Not For Distribution or Resale.

6.1 The General Software Troubleshooting Strategy

If the basic troubleshooting steps haven't solved the problem, it's time for the general software troubleshooting strategy. First of all, determine if there has been a change of configuration recently in any of your computer or network systems. Sometimes the problem is due to software not being installed or a configuration change not being applied correctly. However, if something was previously working and then suddenly stops working, the first question to ask is "what changed?" If you can determine what changed recently in your computer or network system, first verify that the changed system is in the correct state. Once you've verified the change was correctly applied, try to undo the change if you can to see if that fixes the problem.

The second method is to use one or more programs to trace the actions of the software and to analyze the trace output until you find where the problem occurs. From there, go backwards in the trace log and determine why the problem occurs. Once you've determined a probable cause, record the current configuration settings and modify them to see if the problem is fixed. Repeat this process until the problem is fixed. If a particular tracing tool does not reveal any problem, use another, as the problem may be with a different component or with another aspect of the component's operation. You should be familiar with how Windows, networks, and Laserfiche works to identify the component that is the most probable source of the problem.

6.2 List of Useful Tools

Programs for troubleshooting usually trace and display activity so you can see what the misbehaving program is actually doing and where it's going wrong. An error message or a single log message often only tells you the end result, but the following tools can give you an indication of how the program got to that state. The following tools are all free, well documented, and supplied by third parties.

Tool	Description	Best for
Windows Event Viewer	The Windows Event Log is often the first place you should look for information about what has gone wrong.	• Laserfiche Server and Search Engine • Workflow • Web Access • WebLink • Services in general • Program refuses to start • Installation problems
Sysinternals Process	Process Monitor displays file system, registry, process and	• Wide range of problems with Laserfiche clients

Laserfiche Certified Professional Program. Not For Distribution or Resale.

Monitor	thread, and some network activity in real-time. It supports customizable filters and can save logs for later analysis. This well-known program is a successor to the venerable Filemon and Regmon utilities.	• Access denied errors in Windows
Sysinternals Process Explorer	Process Explorer is like a much more powerful version of Windows Task Manager. It is especially useful to see which process is hogging up the CPU or has a particular file locked.	• A generally handy utility • CPU hog • Identifying which process has a file locked • Identifying spyware or other programs running in the background
Sysinternals TCPView	TCPView is similar to the netstat utility that ships with Windows and shows all the open TCP and UDP endpoints and which process opened them.	
Wireshark	Wireshark is an open-source network packet capture and analysis tool. It is indispensible for troubleshooting network problems.	• Client/Server communication problems such as failure to log in • Kerberos • All kinds of network related issues
Microsoft Network Monitor	Microsoft Network Monitor is a network packet capture and analysis tool similar to Wireshark. The latest version is a free download and very capable. This is a good alternative to Wireshark for users leery of open-source software.	• Client/Server communication problems such as failure to log in • Kerberos • All kinds of network related issues
WinHTTP Tracing	Laserfiche clients use an HTTP client library built into Windows known as WinHTTP to communicate with the Laserfiche server. On Windows XP and Windows Server 2003 WinHttpTraceCfg.exe is a utility that ships with the Windows Server 2003 Resource Kit and can enable WinHTTP tracing. In Windows	• Problems with Laserfiche logging in • Errors with a single Laserfiche client after a certain server request

Laserfiche Certified Professional Program. Not For Distribution or Resale.

	Vista and Windows 7 the netsh winhttp command is used to configure WinHTTP tracing.	
Microsoft SQL Server Profiler	SQL Profiler is a tool that ships with SQL Server and displays in real-time the SQL statements that an SQL Server instance is executing, along with execution statistics and other details. SQL Profiler supports filtering by events and by conditions on columns and saving logs for later analysis.	• Performance problems, except for full-text search performance

Laserfiche Certified Professional Program. Not for Distribution or Resale.

6.3 Error Logging

All errors, warnings, and informational messages produced by the Laserfiche 8 Server are logged in the Windows Application Event Viewer on the Server computer. This allows you to view all relevant error messages, as well as information about the current state of the Laserfiche Server on boot-up, in a central location, and sort, export, and save them using the Event Viewer's tools.

Occasionally Laserfiche Support may ask you for these logs if they are helping to diagnose a problem. If they do, please be sure to send the Event Viewer contents as an EV file, which will give them a complete picture for the issue.

Some processes (specifically briefcase and volume actions) produce additional diagnostic logs in the repository directory. When these logs are created, their creation is also logged in the Application Event Viewer, so that you can find and refer to them.

6.4 Compatibility

In general, newer versions of the Laserfiche Server are compatible with older versions of Laserfiche client applications, because we endeavor to ensure that the Server's communication methods will be backwards-compatible. For example, the Laserfiche 8 Server can be accessed with Laserfiche 6 and 7 Clients. However, since servers can only include the communication interfaces that are in existence when they are released, a server cannot necessarily interact with a client that has a newer version. A version 8 Client can't be used with a Laserfiche 7 Server. In addition, when using an older Client application with a newer Server, you will not be able to take advantage of the new features of the Server.

6.5 Repository Creation and Registration

Creating and registering repositories includes a number of tasks, listed below, that are run automatically. The Laserfiche Server service account needs the appropriate rights to accomplish these tasks.

6.5.1 Repository Creation

Creating a new repository pre-creates the three repository components—the first repository volume, the database, and the search catalog—and prepares them to be populated with data when the first documents are imported and scanned. Repository creation consists of the following tasks, which are handled automatically by the Repository Creation Wizard:

- Preparing the database schema for use with Laserfiche. (If you are using Microsoft SQL Server, repository creation can also create the database for you; if you are using Oracle, you must create an empty user schema for Laserfiche to use.)
- Creating a default volume structure.
- Creating an empty search catalog.
- Creating the repository directory.
- Registering the repository with the Laserfiche Server.

6.5.2 Repository Registration

Repository registration makes the Laserfiche Server aware of the repository, so that it can load it and make it available to client applications. Repository registration consists of the following tasks, which are handled automatically by the Repository Registration Wizard:

- Adding the repository registration information, including the repository's network location, to the Server computer registry.
- Instructing the Server to load the repository.
- Attaching the index catalog. (If the index catalog does not exist or cannot be found, registration will still complete successfully and present the user with a message to that effect; the user will then need to create or attach the index catalog in the Administration Console before full-text searching and indexing will function.)

6.6 Repository Migration and Updating

When upgrading from one version of Laserfiche to another, in most cases you will need to allow Laserfiche to update your database schema. Depending on the specific upgrade you are performing, this will mean either a migration or a database update.

6.6.1 Repository Migration

Upgrades from one major version to another require a repository migration. For example, upgrading from Laserfiche 7 to Laserfiche 8, or Laserfiche 6 to Laserfiche 8, require migrations. Migrations are performed using a separate migration utility,

Laserfiche Certified Professional Program - Not For Distribution or Resale.

which can be obtained on the Laserfiche Support Site or from the Server installation autorun package.

A migration primarily affects the database used with the repository to be migrated. Migration does not touch the original Laserfiche 6 or 7 database; instead, it creates a new database, and then copies data from the original database to the new database, "translating" it into the new schema as it does so. The end result is an unchanged Laserfiche 6 or 7 database, and a new database that contains the same information as the old database, but updated to the new Laserfiche 8 format. If the migration fails for any reason, it can be re-run on the old database, which was not modified.

Volume files are not affected by the repository migration, although you will need to verify that the volume path locations remain correct relative to the new repository location. If you cannot locate your documents after migrating, the volume paths should be updated in the Administration Console.

The search index files are discarded when migrating from Laserfiche 6 or 7 to Laserfiche 8, because they use a completely new format in Laserfiche 8. Instead, a new index catalog will be created after the migration, and the migration process will prompt you to allow it to reindex all previously-indexed documents to populate the new catalog.

The migration also keeps a migration log, which will track both errors and informational messages. Note that the migration log is verbose: unless a message is specifically categorized as an error, it is informational and can be safely ignored.

After migrating, it is a good idea to test your repository and make sure everything works properly. It is also vital to update your back up plan—specifically, be sure that you do not continue to back up your now-unused Laserfiche 6 or 7 database, instead of your Laserfiche 8 database!

See the Laserfiche 8 Migration Quick Reference for more information.

6.6.2 Database Updates

When upgrading from one minor version to another (for instance, from Laserfiche 8.0.0 to Laserfiche 8.1.2), you do not need to perform a repository migration. Instead, the Laserfiche Server will handle a database update for you. The database is updated in place; no additional database needs to be created. Once again, the volume files remain untouched. The Server will update the database as soon as it loads the repository after an upgrade. Since the database is being modified, it is good practice to ensure a good backup prior to installing the update.

6.7 Troubleshooting Common Problems

The following is a list of common Laserfiche problems and how to troubleshoot them.

Laserfiche Certified Professional Program. Not for distribution or resale.

6.7.1 Server Will Not Start

The Laserfiche Server (LFS) not starting is usually due to one of three reasons: a DLL or executable file is missing or corrupted, there is a TCP port conflict, or there is a licensing problem. The first place to check is the Windows Event Log. Broken installations can usually be fixed by running a repair operation from Add/Remove Programs. TCP listening port conflicts are corrected by configuring Laserfiche to listen on a different port than the currently configured setting, or by changing the configuration of the conflicting application.

Use the Sysinternals TCPView utility to see which programs are listening on which TCP ports. By default, the Laserfiche Server will listen on port 80 for HTTP requests, port 443 for HTTP over SSL (HTTPS), port 1888 for LFAPI, and port 5051 for notifications requests. On Windows XP 32-bit, the server will listen on port 5050 instead of 80 for HTTP requests to avoid a conflict with IIS 5.1.

One complexity to keep in mind is that lfs.exe does not directly listen on a TCP port for HTTP requests but instead directs HTTP.SYS to receive HTTP requests on the specific port and to forward the requests to lfs.exe. This is why you will not see lfs.exe listening on the configured HTTP port using TCPView but rather a process named "System," which is the Windows kernel.

Since IIS 6 and above use HTTP.SYS as well, it's possible for IIS and LFS to share port 80 without any conflicts. LFS reserves URLs beginning with "/lf" as belonging to it. Any Web sites you configure with IIS on a port that is shared with LFS need to avoid that URL prefix.

HTTP.SYS is also used for HTTP over SSL, but unlike plain HTTP, LFS does not share ports with IIS correctly for HTTPS connections. If you need to configure IIS to listen to port 443, configure LFS to listen on another port for HTTPS requests or disable HTTPS in LFS by setting the listening port for HTTPS to 0. To configure which ports LFS uses for HTTP and HTTPS, edit the DWORD registry values "Port" and "SSLPort" under the HKLM\SOFTWARE\Laserfiche\Engine\8.0\HttpServ key.

Occasionally, the server will start but terminate within a few minutes after starting. This can be caused by a license serial number conflict between two Laserfiche Servers on the same network. LFS will log serial number conflicts it detects to the Windows Event Log.

6.7.2 Cannot Attach Repository

Check the Windows Event Log to see what errors Laserfiche has logged when you cannot attach a repository. Usually the attach error will be accompanied by one or more database statement execution errors. Read through these carefully. If there is a database error that occurred during attach, determine if LFS was able to successfully log in to the database or if an error occurred when LFS tried to log in. Check to see that the database server is up and running, and, if it's on a remote machine, that it

Laserfiche Certified Professional Program. Not for Distribution or Resale.

isn't blocked by a firewall and is accepting remote connections on the expected TCP port. Ensure that the user account that the Laserfiche Server service is running as has sufficient rights to log in. If LFS has attempted to upgrade the schema, but the upgrade fails, contact Laserfiche Technical Support for assistance.

If the error is not database-related, it is most likely due to a problem with permissions. Ensure that the user account that the LFS service is running as can write to the HKLM\SOFTWARE\Laserfiche\Engine\8.0\Repositories registry key and the repository directory.

6.7.3 General Performance Problems

When users complain of poor browsing, document navigation, import, or just general usage performance, the problem is usually due to an overloaded database server. Use the Windows Performance Monitor to view the CPU utilization of the lfs.exe and lffts.exe processes and your database server application. You can also configure Performance Monitor to display the level of paging activity and disk I/O for each process. In most cases, the database server will be using most of the resources.

If the database server is running on the same machine as Laserfiche or other software, separating it on its own machine is often the step that can achieve the largest performance gains. Laserfiche strongly recommends using 64-bit versions of SQL Server and Oracle running on machines dedicated to running the database server with no other application or middleware software. For SQL Server, creating a maintenance plan that gathers statistics on all tables and indexes, rebuilds fragmented indexes, and checks for errors, and that is configured to run on a regular schedule, is strongly recommended. If any database customizations have been made, verify that all Laserfiche-created constraints and indexes are enabled.

In addition to hardware resources being overtaxed by the database server, virus scanners can also negatively impact performance.

6.7.4 Full-Text Indexing Is Impacting Performance

If full-text indexing occurs when users are working in the repository, users may see performance issues as system resources are diverted to the indexer. The simplest solution to this problem is to configure "indexing quiet periods" so that indexing takes place outside of the primary business hours. In some cases this is not possible, but configuring quiet periods to avoid indexing during peak hours may be a feasible alternative.

Another strategy is to import documents from Laserfiche using Quick Fields at scheduled times. This way automatic indexing can be enabled, but both import and indexing can occur during off-peak hours.

Laserfiche Certified Professional Program. Not For Distribution or Resale.

6.7.5 Running Out of User Licenses

A common complaint among users in many installations is that Laserfiche reports there are no user licenses available when attempting to log in. The two solutions to this problem are to buy more user licenses until supply meets demand or to convert to named user licensing. Named user licensing is predictable, guaranteed, and recommended by Laserfiche to avoid the uncertainty inherent in the concurrent licensing model.

The Laserfiche Server allows administrators to keep track of how many licenses are in use through performance counters. By keeping logs of license use, administrators can track how licenses are used throughout the day, the week, and over the course of a longer period of time. Audit Trail allows administrators to run reports to track how frequently individual users are logging in. These tools should allow for a more accurate prediction of a site's user license requirements than estimates based on hypothetical scenarios or guesses.

6.7.6 Kerberos Authentication with Delegation Does Not Work

The Troubleshooting Kerberos Problems page at http://technet.microsoft.com/en-us/library/cc728430%28WS.10%29.aspx is a good place to start. The "How the Kerberos 5 Authentication Protocol Works" at http://technet.microsoft.com/en-us/library/cc772815(WS.10).aspx is an excellent resource for understanding Kerberos V, and this knowledge can be essential for troubleshooting.

Verify that the client, Server, and domain controller have synchronized times and that the correct SPNs are registered. Use the SetSPN tool to verify that an HTTP/hostname.domainname.com SPN exists for the user account that the LFS service is running as, on the machine that LFS is running on. Kerberos requires a properly functioning DNS and Active Directory system. Ensure the basic infrastructure is working before continuing with Laserfiche specific investigation.

Once you have done the preceding you will probably need to use a network tracing tool like Wireshark or Microsoft Network Monitor to trace HTTP and Kerberos requests. Run the tracing on the client and attempt to log in to Laserfiche using Windows authentication. In the WWW-Authentication header in the requests and responses you will see an SPNEGO token that can contain either an NTLM packet or Kerberos ticket request. Both Wireshark and MS Netmon can decode this. In order to understand the ticket you must understand how Kerberos V works, for which we suggest reading the paper above.

In short, one should see a KRB_TGS_REQ request for the HTTP/hostname.domain.com SPN sent to the domain controller and a reply from the DC to the client in the form of a KRB_TGS_REP packet using the Kerberos V protocol. At this point the HTTP exchange will begin between the client and the Laserfiche Server. The client should embed a Base64 encoded KRB_AP_REQ request packet in the WWW-Authenticate header immediately after the initial "SPNEGO"

Laserfiche Certified Professional Program - Not For Distribution or Resale.

string in its HTTP request. The Server should then immediately authenticate the user and reply with a successful logon response.

If the KRB_TGS_REQ/KRB_TGS_REP exchange does not work correctly, the problem is probably with the service registration, and you should check the SPN registration. If the KRB_AP_REQ/KRB_AP_REP exchange embedded in the WWW-Authenticate header in the HTTP exchange between the Laserfiche Client and Server is not correct, check that the Kerberos ticket that the Client has cached is valid using KerbTray, klist or a similar tool.

Note: Be sure to watch the video on Configuring Kerberos.

6.7.7 User(s) Cannot Log In

If a user or users cannot log in to Laserfiche, usually the Laserfiche error code and message indicate the problem. If the error code is not useful, check to see if the problem is limited to a single client workstation, user account, or Windows profile. This can often provide clues to the problem. For example, if log in requests do not work on a particular workstation no matter the Laserfiche account or Windows profile, but work from other client workstations, the problem is probably a misbehaving anti-virus engine, or a router or firewall filtering out requests from that particular workstation. If the problem only occurs with a particular Laserfiche account, that account is probably denied log in access. Check the Laserfiche configuration again for account lockout and Windows account log on denies.

Check that the Server's host name contains the correct port number and that the Server does not have a firewall blocking lfs.exe from listening on that TCP port. Check that lfs.exe is running and has that repository registered. Check that the license lfs.exe is using is not expired and has sufficient user licenses.

If the above checks don't reveal anything misconfigured, the next step is to trace the HTTP exchange between the Laserfiche Client and Server during log in. The Laserfiche Client uses WinHTTP to communicate with the Server. WinHTTP is an HTTP client library that is a component of the Windows operating system and supports tracing all requests. If the client is running on Windows XP or Server 2003 use the WinHTTPTraceCfg.exe utility, which ships as part of the Windows Server 2003 Resource Kit to enable tracing.

Under Windows Vista, Server 2008, and later versions of Windows, the "netsh" program, which ships with Windows, should be used. Run "netsh winhttp set tracing trace-file-prefix="C:\temp\winhttp" level=verbose format=ansi state=enabled" to enable tracing, writing log files to the "C:\temp\winhttp" directory—though you can change the path to write files to whatever directory you like. When you are done attempting to log in with the client, run "netsh winhttp set tracing state=disabled" to stop tracing. You can then examine the log files to see what data the client sent to the Server and the response, if any, from the Server.

Laserfiche Certified Professional Program Training Material. Distribution or Resale.

6.7.8 Search Results are Missing Data or Contain Unexpected Results

If search results appear to be missing data, the problem is generally a result of one or more of the following conditions:

- Incorrect query
- Overly restrictive security access controls
- The search engine is not running or cannot be reached by the Laserfiche Server
- The search engine has not yet indexed the document
- The search catalog is corrupt and requires a rebuild

The last three causes are only relevant for full-text search queries of document text, as metadata searches do not use the search engine. To verify that the query is correct, either enable tracing of search requests in the Server or view the search query in the client by expanding the advanced search syntax node in the Search Pane.

6.7.9 MS SQL Server Transaction Log File Grows To a Very Large Size

If a SQL Server database's recovery model is set to Full or Bulk-logged and database backups are not regularly performed using SQL Server's integrated backup facility, the database's transaction log will grow without bound. If this occurs the simplest solution is to change the database's recovery model to Simple and then to perform a database shrink operation. This will shrink the transaction log and prevent the problem from reoccurring.

6.7.10 Error 9008 Is Returned When Performing an Operation

If a Laserfiche client application returns an error with 9008 as the error code, a database related error has occurred in the Server. This is a generic error indicating there is a connectivity problem between the Laserfiche Server and the database server or that there is a defect in the Laserfiche Server. The Windows event log on the Server machine should have more details about the error, generally logged in an event with message ID 10050. There might be more than one relevant event logged. It should be apparent from the error message if the problem is related to database connectivity. Otherwise, please contact Laserfiche technical support.

6.7.11 Error 9003 or 9004 Is Returned When Attempting to View Documents

If a Laserfiche client application returns an error with 9003 or 9004 as the error code when attempting to view document pages or an electronic document, mostly likely the volume paths are not configured correctly in the Laserfiche Server. Using the Laserfiche administration console, verify that the volume paths are configured correctly and that the volumes are accessible to the Laserfiche Server process. If the

Laserfiche Certified Professional Program – Not For Distribution or Resale.

volumes are stored on a network share, check that the user that the Laserfiche Server service runs as has been granted access to the network share. If problems persist, check the Windows event log for more details.

6.7.12 Named Device License Isn't Recognized When Logging In

If a workstation has been registered as a named device but the Laserfiche Server does not seem to recognize it and use the license when logging in from the workstation, an incorrect host ID was probably registered. Laserfiche uses host IDs to identify client machines for named user licensing: client applications pass it to the Server when logging in, and the Server checks it against the Server's named device database. The match has to be exact for the session to use the named device license.

To fix an incorrect host ID, first determine which host ID the client is passing to the Server and then register that host ID as the named device for the client workstation.

To determine the host ID that the client is passing to the Server during the log in process, enable WinHTTP tracing and attempt to log in. The X-LF-HostID header will contain as its value the host ID of the workstation.

Multi-homed workstations can have multiple host IDs, but the client will only send over a single host ID. Sometimes the host ID that one tool picks as the primary host ID will be different than the one that the Laserfiche client has chosen, and you simply need to register the correct one with Laserfiche Server (or License Manager, for Rio systems).

6.8 How to Control and Use Tracing

Laserfiche comes with an extensive set of built-in tracing facilities that supplement the facilities that are built into the operating system and provided by third party tools. Often, data from multiple sources will have to be gathered together and analyzed as a group with the need to correlate events recorded in one log with data in another. Knowing which tracing tools to employ in which situations, how to control what is traced, and how to filter and interpret trace output are all important skills in making the most out of tracing.

6.8.1 Laserfiche Server ETW Tracing

The Laserfiche Server and search engine, starting from version 8.0, support tracing many of its internal operations using a Windows subsystem called Event Tracing for Windows (ETW). The ETW subsystem is tightly integrated into the Windows kernel and many Microsoft server applications implement ETW providers. This allows trace output from the Laserfiche Server and search engine to be seamlessly integrated into trace output from the Windows kernel and Microsoft SQL Server.

The architecture of ETW divides responsibility for tracing across three types of programs: producers, consumers, and controllers. Producers are programs that generate trace output, consumers are programs that process and format trace output,

Laserfiche Certified Professional Program - For Distribution or Resale.

and controllers are programs that enable, disable, and control the tracing by producers.

In this architecture, the Laserfiche Server and search engine services are producers. Laserfiche ships with the Laserfiche Event Trace Viewer, an ETW consumer that can open, parse, and display ETW trace files. Consumers are required to view ETW trace files because the trace files are in a binary format and require processing to be human readable. Starting with Laserfiche 8.1.1, the Laserfiche Event Trace Viewer can also act as a controller and enable tracing of events in the Laserfiche Server. The Windows Performance Monitor can also act as a controller for any program that can act as an ETW producer.

The architecture of ETW is built around actions that are traced, called events, which are organized into classes or groupings of related event types—which is similar to how Laserfiche Auditing events and event classes are organized. The Laserfiche Event Trace Viewer, starting with version 8.1.1, allows tracing to be enabled at the event class level for the Laserfiche Server. Windows Performance Monitor only allows event tracing to be controlled at the producer level for applications, so either all or none of the events generated by an application producer will be enabled for logging.

While Laserfiche Server and search engine ETW tracing can be indispensible for troubleshooting, the trace output tends to be difficult to interpret because it is at a low-level and thus requires some knowledge about the internal logic of the Server and search engine. For end-users attempting to interpret Laserfiche ETW trace logs on their own, the server's CLI Calls event class is usually the most useful for troubleshooting purposes. This can be used to determine which call a 9013 error, "access denied" is originating from, for example.

ETW trace logs which are generated on Windows Vista, Windows Server 2008, or later operating systems can only be decoded and viewed on Windows Vista, Server 2008, or later versions of Windows, due to an ETW format change from the version used in Windows XP and Windows Server 2003. The newer versions of Windows can read ETW trace logs generated on computers running older versions of Windows.

When a controller makes a change to ETW tracing, the change usually takes effect immediately in the producing application with no system or application restart generally required. This is advantageous, as often for troubleshooting purposes one only wishes to enable tracing for a very short period of time, enabling it just before attempting to reproduce a problem and disabling it immediately afterwards. ETW tracing can have a significant negative performance impact.

Laserfiche Certified Professional Manual. Not For Distribution or Resale.

Laserfiche®

Laserfiche Administrator II Training Manual
June 2010

Authors: Pieter Nauta, Jennifer Dickson, Connie Anderson, Snow Tempest, Jeff Huang, Dave Haas
Technical Editor: Justin Pava
Editor: Tammy Kaehler

Laserfiche
3545 Long Beach Blvd.
Long Beach, CA 90807
U.S.A

Phone: +1.562.988.1688
www.laserfiche.com

Laserfiche is a trademark of Compulink Management Center, Inc. Various product and service names references herein may be trademarks of Compulink Management Center, Inc. All other products and service names mentioned may be trademarks of their respective owners.

Copyright © 2010 Laserfiche
All rights reserved

Laserfiche Certified Professional Program - Not For Distribution or Resale.