Guidance for Laserfiche Deployments on Microsoft Azure

January 2023

Laserfiche®

Contents

Introduction

This document provides general guidance to customers and Solution Providers on deploying Laserfiche within a Microsoft Azure environment. The recommendations here are broadly applicable for systems with up to 1,000 active users with document-centric processes.

Infrastructure

Servers

Operating Systems

Laserfiche 11 supports Windows Server 2016, 2019, and 2022. We strongly recommend Windows Server 2019 or 2022 for any new deployments. Please see the Laserfiche Support site's Supported OS List for the latest information.

Databases

Laserfiche 11 supports Microsoft SQL Server 2016, 2017, 2019, and 2022 (Standard or Enterprise edition) on Azure VMs, as well as Azure SQL Managed Instance and Azure SQL Database. We recommend using Azure SQL Managed Instance to take advantage of its automatic backups, patching, and general reliability, with the lowest configuration complexity. Use "Memory optimized Premium-series" or "Premium-series" instances for Production if available, as they use current generation hardware with significantly higher performance than Standard-series instances. We recommend a minimum Managed Instance size of 4 vCore for up to 500 users, and a minimum of 8 vCore for 500 to 1,000 users.

If running SQL Server on Azure VMs, please consult Microsoft's VM size: Performance best practices for SQL Server on Azure VMs documentation for the latest guidance.

Instances

At the time of writing, the General Purpose Das_v5 (3rd Generation AMD EPYC 7763v CPUs) and Ds_v5 (3rd Generation Intel® Xeon® Platinum 8370C (Ice Lake) CPUs) instance series are ideal for most Laserfiche components and provide the best price/performance for Production systems. We recommend a minimum instance size of D4(a)s_v5 (4 vCPU, 16 GB RAM) with Accelerated Networking enabled for each Production server. Laserfiche Full-Text Search and SQL Server are memory-intensive components. Production servers hosting these components should use E-series E(a)s_v5 VMs, which have double the memory-to-CPU-core ratios of the D-series.

If Azure releases newer generations of D- and E-series VMs with subsequent generations of x86 AMD and Intel CPUs (Ds_v6, v7, etc.), we recommend using the latest available generation in those series.

B-series ("burstable") instances are generally suitable for cost-effective non-Production environments. Do not use B-series instances to host Laserfiche Quick Fields Agent as that component can have high sustained CPU usage (by design) if schedules run on a frequent basis even if there is no actual content to process.

Server Configuration

The table below describes a baseline Production environment for up to 500 total users with moderate activity levels.

Server Role	Azure Instance	Laserfiche Applications
Content Server	D4as_v5	Laserfiche Content Repository Server
Search & Audit Server	E4as_v5	Laserfiche Full-Text Search, Audit Trail
Workflow Server	D4as_v5	Laserfiche Workflow
Web Server	D4as_v5	Laserfiche Web Client, Forms, Directory Server
Capture Server (optional)	D4as_v5	Laserfiche Quick Fields Agent, Import Agent
SQL Server (Managed Instance)	E4s_v5 (4 vCore)	Laserfiche databases

Non-Production environments can either mirror Production while using smaller or B-series instances or scale down to two to three servers ("web" and "app"). Non-Prod environments can share a SQL instance to save on costs. If sharing a SQL instance between non-Prod environments, create separate SQL accounts for each environment and only grant them access to the databases for their environment.

Larger systems with over 500 users and/or heavy activity levels that need performance scaling should generally first increase SQL Server resources to 8 vCPU and 128 GB RAM or higher. The performance of most core Laserfiche components, especially Laserfiche Server repositories, Forms, and Workflow is primarily driven by SQL.

Server Storage Configuration

All Production servers should use 128 GB Premium SSD (P10) Azure Managed Disks for their OS disks. These provide 500 IOPS (3,500 burst) and 100 MB/s throughput (170 MB/s burst), which is broadly sufficient for Windows OS and Laserfiche applications. The Azure VM Availability SLA only covers VMs using Premium Managed Disks for their OS.

Non-production servers can use either Premium SSDs or Standard SSDs for their OS disks.

Repository Storage Configuration

Use Managed Disks attached to the Content Server as data disks to store Laserfiche repository volumes in almost all circumstances. They provide the simplest configuration, best performance, and greatest compatibility with Azure backup and recovery services.

There is no one-size-fits-all repository storage configuration. The ideal configuration for your scenario will depend on priorities between cost, performance, and business objective. Note that for most customers, repository storage is only a small portion of overall infrastructure costs (often around 10%), so there are often only small absolute gains from cost-optimizing storage before the 2 TB level.

Premium SSD storage is approximately $0.12 USD per GB per month.
Standard SSD storage is approximately $0.08 USD per GB per month.
Standard HDD storage is approximately $0.04 USD per GB per month.

Laserfiche repository storage workloads tend to be more IOPS intensive than throughput intensive, especially for document-centric usage. This is because Laserfiche repositories store content on the backend volumes in the form of many small files. For example, a 100-page scanned document stored as TIFs will have 100 TIF files (one for each page), 100 text files (one for the OCR text from each page), and associated thumbnails for those images.

Using Premium SSDs for repository storage will provide the best performance at the highest cost. Premium SSDs have the best IOPS performance by far. If your repository is expected to be 2 TB or less and is not entirely archival content, use Premium SSDs. The smallest size disk we recommend in

Production is a P20, which has 512 GiB, 2,300 IOPS (3,500 burst), and 150 MB/s throughput (170 MB/s burst).

If you have between 2 and 8 TB, consider a mix of an "Active" volume on a 2 TB or smaller Premium SSD and an "Archive" volume on a Standard SSD holding the remainder. You can use Laserfiche Workflow's Migrate Volume activity to schedule moving content from the Active volume to the Archive volume using criteria of your choice, such as "Last modified more than two years ago".

If you have 8 TB or more of content, consider using E60 (8 TiB) or larger Standard SSDs. At the E60 size, Standard SSDs get a massive jump in performance from the lower sizes which are capped at 500 IOPS (1,000 burst) and 60 MB/s (250 burst). An E60 has 2,000 IOPS and 400 MB/s throughput: similar IOPS to a P20 Premium SSD and over double the throughput.

Standard HDDs are suitable only for very infrequently accessed content. For example, they are a good fit for storing many TBs of archival content that must be retained for decades or indefinitely. If their performance characteristics are suitable for the workload, Standard HDDs are by far the most cost-effective repository volume storage option.

Note that Azure allows changing disk types in-place. If you are migrating a large volume of content into a Laserfiche system on Azure and do not plan to use all Premium SSDs, we highly recommend temporarily changing the Standard SSDs or Standard HDDs to Premium SSDs while the migration is running to avoid storage bottlenecks, then changing the disk type back once the migration is complete. Consider temporarily increasing the Laserfiche Server's instance size as well to avoid compute bottlenecks.

SQL Databases

Use of Azure SQL Managed Instance greatly simplifies many database configurations. Azure SQL Managed Instance databases always automatically use the Full recovery model to enable the service's automatic point-in-time recovery backups. They also automatically manage storage configuration, though advanced manual configuration is possible for performance tuning with large databases. Most Laserfiche databases never reach the size where this storage tuning would provide any benefit.

Laserfiche databases should use the default SQL Server collation of SQL_Latin1_General_CP1_CI_AS.

Customers should configure SQL database Maintenance Plans to regularly rebuild fragmented indexes and update outdated statistics. We recommend Ola Hallengren's well-regarded SQL Server Index and Statistics Maintenance scripts with default settings as a starting point. Running the index and statistics maintenance plan on at least a weekly basis is appropriate.

Backup and Business Continuity

impacts with many small files) in all scenarios.

For backing up Laserfiche systems in Azure, customers should at minimum use the following that apply to their environment:

- Azure Backup (or equivalent snapshot-based 3rd party backup solution) with a Geo-Redundant Storage (GRS) Backup Vault for Production system VMs and file shares with minimum 30-day retention:

    - Managed Disks hosting repository and/or Forms file volumes (1-4x daily)
    - SQL Server running in Azure VMs with Laserfiche application databases
- Azure SQL Database & SQL Managed Instance:
    - Automated backups: "Both SQL Database and SQL Managed Instance use SQL Server technology to create full backups every week, differential backups every 12-24 hours, and transaction log backups every 5 to 10 minutes. The frequency of transaction log backups is based on the compute size and the amount of database activity. When you restore a database, the service determines which full, differential, and transaction log backups need to be restored."
        - We highly recommend setting the retention for these automated backups to the 35-day maximum.
    - Long-term Retention for databases that must be retained beyond the 7-35 days provided by the automated backups described above.

Non-Production environments may have shorter backup retentions and use LRS/ZRS storage. If a Non-Production environment would be used for backup and disaster recovery testing, it should mirror at least the storage availability tier of Production.

In addition to the core backups above, customers with aggressive Recovery Point and Recovery Time Objectives (RPO/RTO) should consider using Azure Site Recovery (ASR). It is a Disaster Recovery as a Service (DRaaS) solution which provides continuous data replication for VMs (including Managed Disks) between Azure regions. ASR is cost effective because it does not require actively running VMs in the DR region until they're needed.

Use Azure SQL Database Auto-Failover Groups in conjunction with Azure Site Recovery. For more information see Overview of business continuity with Azure SQL Database.

If using Azure Files, configure shares to use GRS or GZRS storage. Ensure you have a through understanding of the data loss risks inherent to asynchronous replication. See Azure Files Disaster recovery and storage account failover for more details. We recommend customers who want to use ASR keep repository volumes on Managed Disks rather than Azure Files.

Please be aware that there are significant networking and interoperability considerations with an Azure Site Recovery deployment. Review, at minimum, the following Azure documentation before proceeding:

- About Azure VM disaster recovery (all sections, especially:)
    - Azure-to-Azure support matrix: Replicated machines - storage
    - Support for using Site Recovery with Azure Backup
- About networking for disaster recovery (all sections)
- About recovery plans
- How-to Guides: Azure to Azure Disaster Recovery (all sections)

Security

Encryption

Storage / Data at-rest

Azure Managed Disks are transparently encrypted using 256-bit AES encryption. More information is available here: Server-side encryption of Azure Disk Storage

Data in-transit

Configure Laserfiche applications to use HTTPS/TLS 1.2 as detailed in Configuring SSL/TLS Encryption in Laserfiche and online product documentation. At minimum, configure HTTPS for the Laserfiche web servers that end users will connect to. Use X.509 (SSL) certificates from either an internal Certificate Authority (CA) or a third-party trusted CA like DigiCert, Lets Encrypt, or GlobalSign as appropriate.

Network

Use Azure Network Security Groups (a stateful firewall) to restrict traffic as necessary. We recommend allowing all outbound traffic from any Laserfiche server to any other Laserfiche server and using only Inbound/Ingress rules. See Default Network Ports for Laserfiche Products for the inbound ports that Laserfiche applications require. Any application that lists a default port of 80 will use port 443 for HTTPS. Laserfiche applications will generally require outbound connectivity to Active Directory for LDAP (389/636) and your chosen SMTP endpoint for email notifications (25/465/587).

Endpoint Protection ("Anti-virus")

Laserfiche is broadly compatible with commercial endpoint protection solutions. We do not recommend any blanket exclusions of Laserfiche applications as doing so can unnecessarily weaken a customer's security posture. Customers should only apply AV exclusions in response to a specific observed performance issue or when instructed by Laserfiche support.

It is also appropriate to *temporarily* disable file scanning during a content migration, as the security scanners will often severely bottleneck the process when they try to scan millions of new files appearing. Watching for high CPU usage on the security agent can tip you off to an issue here.

Network Requirements

Laserfiche recommends end user workstations have at least an 100/25 Mbps down/up connection for the best experience, with 25/5 Mbps as an absolute minimum.

If end users are connecting to Laserfiche from VDI clients hosted in the same Azure data center in the same (or peered) VNET, there should be at least a 10 Gbps connection between the Laserfiche servers and VDI servers and thus no network performance concerns from the Laserfiche end. In this scenario, consult your VDI solution's networking recommendations as those are what will apply.

End User Client Requirements

Most Laserfiche functionality is accessible entirely through a supported modern web browser (Chrome, Edge, or Firefox).

Laserfiche has a few optional desktop client components that require Windows 10 or 11. These include the Laserfiche Office Integration, a COM plugin for Outlook, Word, Excel, etc. and Laserfiche Scanning,

which is required for the solution to interact with local scanners. All of these components can be deployed through unattended installation.

Laserfiche®

Guidance for Laserfiche Deployments on Microsoft Azure
February 2022

Author: Samuel Carson

Description:
This document provides general guidance to customers and Solution Providers on deploying Laserfiche within a Microsoft Azure environment. The recommendations within are broadly applicable for systems with up to 1,000 active users with document-centric processes.

Laserfiche
3443 Long Beach Blvd.
Long Beach, CA 90807
U.S.A

Phone: +1.562.988.1688
www.laserfiche.com