Infor LN Sizing and Technical Infrastructure background

Sizing Documentation

Publication Information

Release: Infor LN

Publication date: November 25, 2014

Document code: B0065D

Contents

About this guide

This document provides an overview of the infrastructure planning requirements for an Infor LN implementation; it discusses the Infor benchmark methodology, sizing methodology, hardware, operating system, database, disk, and network requirements.

This document does not replace the sizing guide, but is intended to provide high-level sizing background information without too much technical details.

The scope of this document does not include additional IT environment issues such as availability, redundancy, disaster recovery, backup/recovery procedures, and work requirements for Infor LN.

In providing this sizing information, Infor assumes no liability whatsoever for results that differ from estimates. Infor recommends that you confirm the data and sizing based on your own implementation experience and in consultation with the selected platform vendor.

Performance warranties

Infor performs sizing thoroughly, but cannot guarantee optimal performance because there are many variables in the sizing procedure. Infor assumes no liability for any damages or extra costs that result from recommendations based on this sizing guide.

Intended audience

This document is intended for technical consultants of Infor, partners and customers who are responsible for sizing and implementing Infor software.

Related documents

Document number	Document title
B0071 US	Infor LN Minimum Hardware Requirements on Windows/SQL Server
B0036 US	Infor LN Configuration and Planning Questionnaire
B0052 US	Infor LN Sizing White Paper
B0080 US	Infor Ming.le LN UI Plugin 11.1

Contacting Infor

If you have questions about Infor products, go to the Infor Xtreme Support portal at www.infor.com/inforxtreme.

If we update this document after the product release, we will post the new version on this Web site. We recommend that you check this Web site periodically for updated documentation.

If you have comments about Infor documentation, contact documentation@infor.com.

Chapter 1 Infor LN Architecture

1

Introduction

Infor LN can be configured in many different ways; each way has its advantages and disadvantages. During benchmarks, different setups are used. To know how the Benchmark and Sizing methodology should be used, a good understanding of the Infor LN architecture is required. This chapter explains which components of the Infor LN architecture are important and where these components are used in different types of configurations.

In general, the Infor LN architecture consists of three layers:

- Presentation Layer
- Application Layer
- Database Layer

The connection between these layers is shown in the following figure.

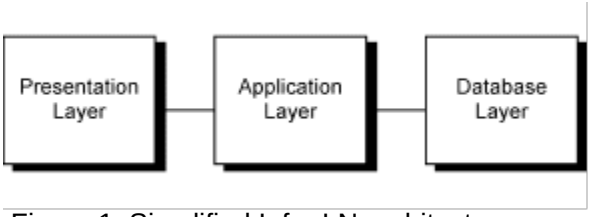

Figure 1: Simplified Infor LN architecture

Each of these layers contains one or more physical components.

Layer	Component
Presentation	Infor Windows (WebTop) or Infor Web UI. The ASCII and X-Windows interface are no longer supported with Infor Baan5.0x and above.
Application	Infor LN Virtual Machine. The most important component is the bshell, which works as an interpreter on the Infor LN objects. These Infor LN objects contain the compiled Infor LN 4GL code of the programs. The bshell interprets the 4GL code and can be seen as the engine behind the

	presentation layer. The bshell sends display information to the presentation layer and receives user response from it. On the other side, it has a connection to the database layer where it sends queries to and receives the results.
Database	For each supported database there is a driver. This component, known as the database driver, is the most important part of the database layer. The database driver transforms the Infor LN queries into database-dependent queries and takes care of the communication between database and application layer.

When the different layers are transformed into components, the following can be made.

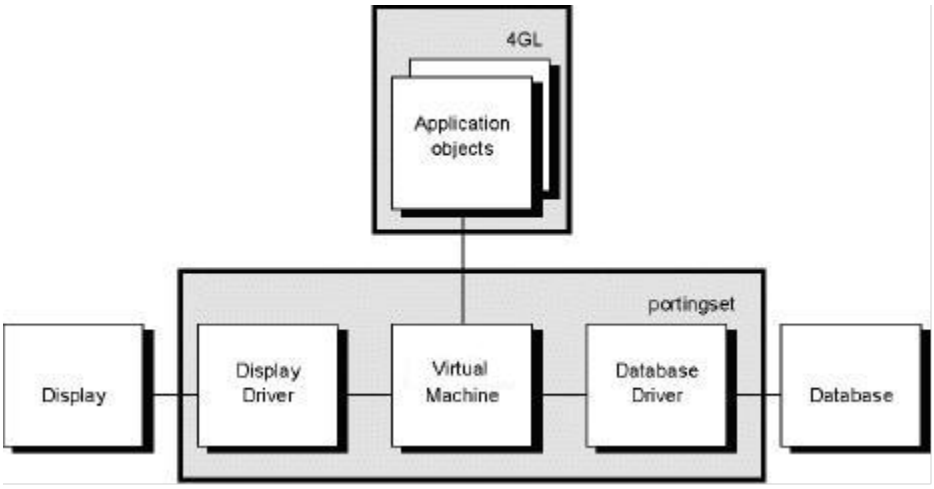

Figure 2: Infor LN architecture

The set of components is called a porting set, the content of which is hardware dependent.

This architecture is set up to run Infor LN. Infor LN contains the following software parts:

- Infor LN 4GL objects. Compiled 4GL application software that includes menus, forms, and reports.
- Data dictionary. Defines the data model in the application.

Infor LN contains multiple modules, each of which covers a logical part of the application. The Infor LN modules include the following:

- Infor LN Common Data
- Infor LN Manufacturing
- Infor LN Project
- Infor LN Service

- Infor LN Order Management

With this architecture, Infor LN is extremely open and can easily adapt new technologies. This is why, for example, Infor LN, which was originally designed as a UNIX application, can be ported to Microsoft Windows without changes to the application code.

Platform/database support

The main platforms/databases supported for Infor LN:

Supported OS	Oracle	DB2	SQL Server
HP-UX Itanium	☑		
IBM Power AIX	☑	☑	
Microsoft Windows	☑		☑
Linux (SuSE, Redhat)	☑	☑	
Oracle SPARC Solaris	☑		

For detailed information about the supported platforms and databases, contact your sales representative. For customers the platform support matrix is available via Infor Xtreme solution 1183466.

Configurations

The three most common deployment options of Infor LN are as follows:

- 3-tier deployment: the application and database components reside on separate physical servers.

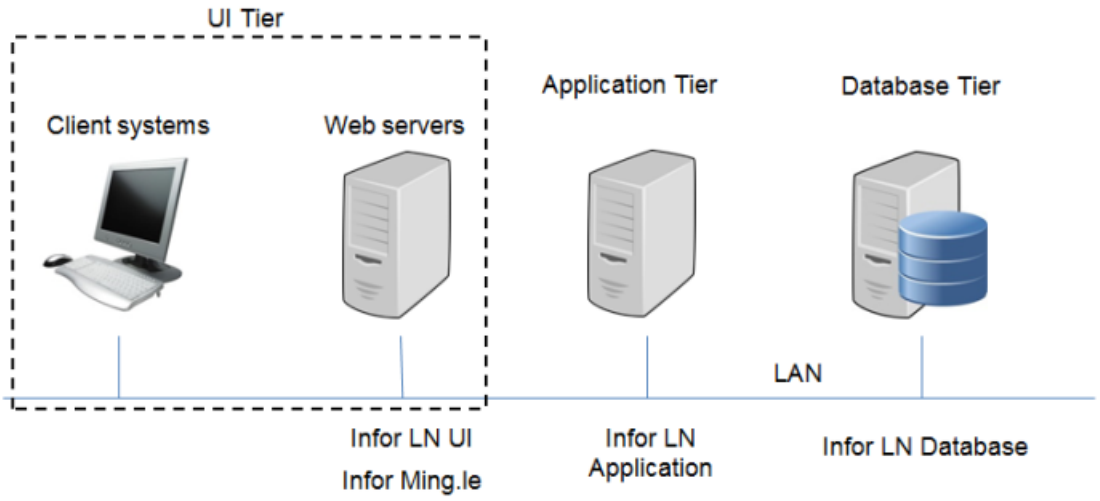

- 2-tier deployment: the application and database components share one single server.

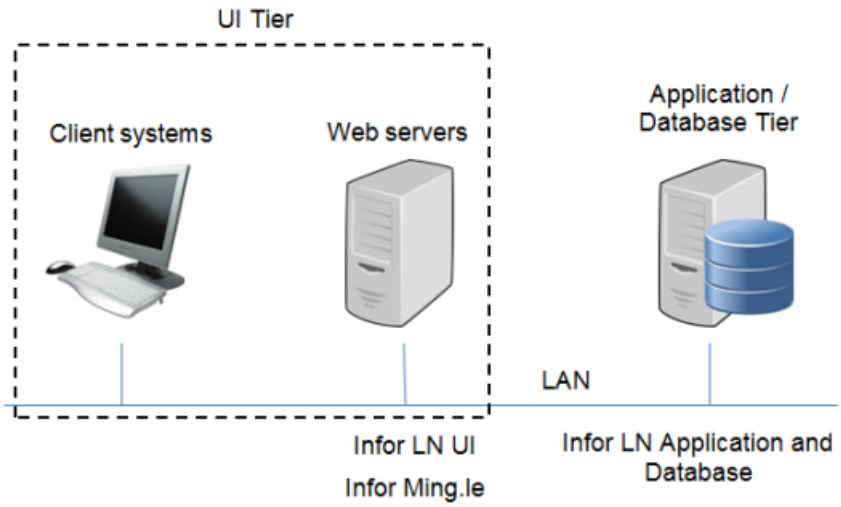

- Single-tier deployment: application, database, and Web components share one single server.

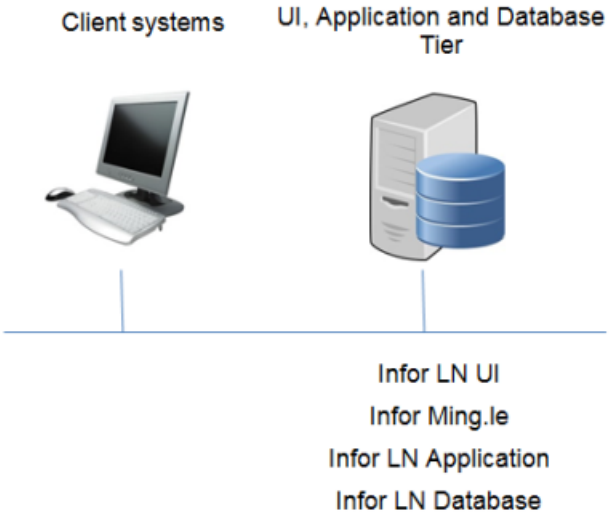

Chapter 2 Infor LN Benchmark Methodology

2

Infor developed its own benchmark methodology to collect Infor LN sizing information. The methodology is described here.

The Infor LN Benchmark Methodology allows you to measure the performance of various computer system configurations. The unit of measurement is Infor LN Benchmark Reference User (BRU), which is the output of the Infor LN Load Factor Sizing Model.

The goals of the SBM were that the methodology should be the following:

- Reliable
- Easy to reproduce
- Consistent during reproduction

The Infor benchmark methodology is designed to simulate and measure a peak load. An automated tool simulates users executing real Infor LN application processes and sessions. Each application process has multiple steps and different load profiles that are relative to other processes. The system under test is said to be saturated when user response times exceed a predetermined threshold. At this point in time, the number of sessions and the mix of sessions are reported, which provides a measurement of load. The benchmarks are performed jointly by Infor and the hardware vendors.

LoadRunner

Infor selected LoadRunner to perform the benchmarks. LoadRunner is a load testing tool developed by HP Mercury Interactive. Virtual users are used to put load on the benchmarked configuration. The virtual user is a "C" program that sends API calls to the Infor LN virtual machine. The virtual user acts as a real user and uses the Infor Windows user interface. LoadRunner uses a controller system that controls all virtual users and gathers the results. A host system runs the virtual users. The controller and the host system are Microsoft Windows systems. No actual Infor Windows user interface installation is required to simulate the users; the API calls handle the interfacing. The following figure shows how various parts of the benchmark setup are implemented.

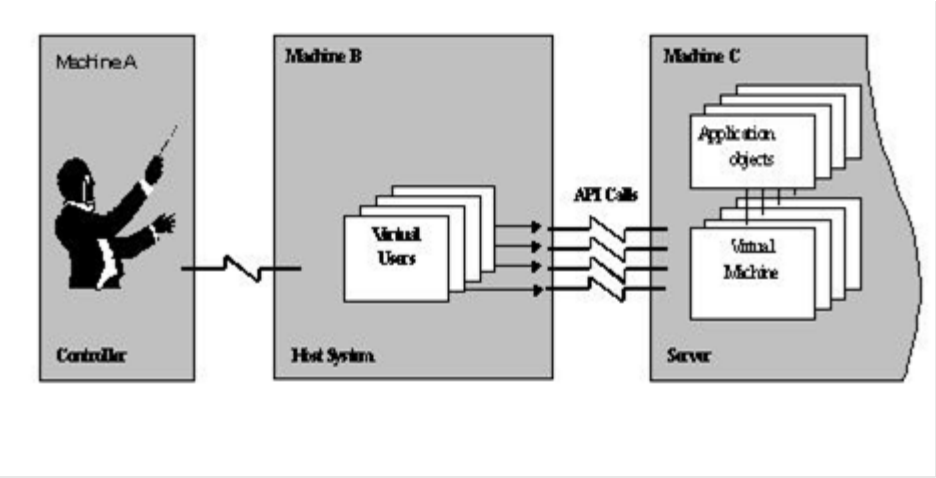

Figure 3: Overview LoadRunner network connection

User profiles in the Infor LN benchmarks

To generate a realistic load on the System-Under-Test (SUT), a realistic mixture of user profiles has been made. These profiles use a broad range of data in the underlying database and a broad range of sessions in the Infor LN applications.

Various scripts have been developed, each of which handles a unique set of tasks. The following tables list the benchmark scripts that are used for the various Infor LN versions and the main activities in these scripts.

Script name	Main activities
Service	- Create call - Transfer call - Estimate materials and labor - Assign engineer - Release service order - Generate/Release inbound advice - Enter call diagnostic/actual problem - Freeze/Confirm shipments loads - Set material lines to ‘costed’ - Insert actual hours - Set labor lines to ‘costed’ - Report service order complete - Service order costs - Accept call - Service invoicing

Finance	- Create financial batch - Create transaction - Create two documents - Create five document lines per document
PurchaseOrder	- Insert one purchase order with five order lines - Approve purchase order - Release to warehousing - Warehouse receipts - Process purchase order
SalesOrder	- Insert one sales order with five order lines - Confirm Order - Print Order - Release order to warehousing - Generate outbound advice - Release outbound advice - Confirm shipments - Release to Invoicing - Create Billing Request - Process Invoice - Process sales order
Browsing	- Startup successively eight different lookup sessions and search and browse through the data

Each of the benchmark user profiles has a unique load factor. Together with its contribution to the total mix of user profiles, an overall load factor of the benchmark is determined.

The following figures show the mix of user profiles for the various Infor LN benchmarks.

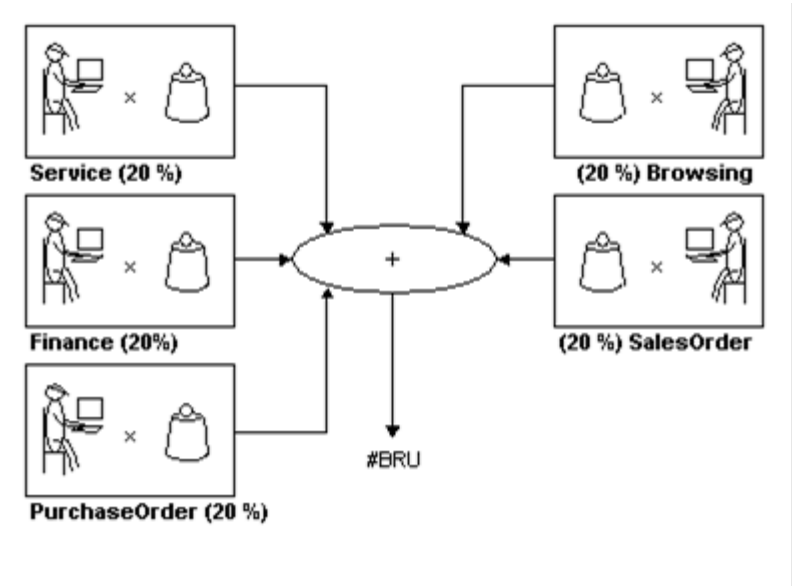

Figure 4: Mix of user profiles in the Infor LN benchmark

Performance criteria of Infor LN benchmark

In an Infor LN benchmark, the SUT is fed by several users, based on the profiles according to the mixture as described previously. The number of users in the benchmark increases to a particular amount of load that the system becomes fully saturated. Therefore, all the CPU power is used to serve the Infor LN software. To achieve the best results, the following rules are used during the benchmarks:

- The SUT has sufficient memory (never memory bounded)
- The SUT has no I/O problems (never I/O bounded)
- In a database server benchmark, the application servers are always oversized (never more than 70 percent CPU utilization), so the servers become transparent for the database server.
- In an application server benchmark, the database server is always oversized.

If the system is CPU-bounded (no more CPU power left), the response times of the individual transactions of the Infor LN software increase. LoadRunner can measure response times of end-to-end transactions (from user interface to database and back). In each script, the response time criteria of these transactions, mainly important transactions such as writing a sales order line, searching a record, and printing an invoice, are defined. A few of these transactions are defined to be critical and have a defined upper limit. When one of these transactions nears the upper limit, the maximum number of users is reached in a benchmark.

Chapter 3 Infor LN Sizing Methodology

3

Infor LN Load Factor Model

The Performance and Benchmarking Center of Infor has developed the Infor LN Sizing Methodology, which encompasses the following elements:

- Infor LN Load Factor Model
- Batches/Integrations
- Infor LN implementations
- Benchmark data

The following figure illustrates the Infor LN sizing methodology.

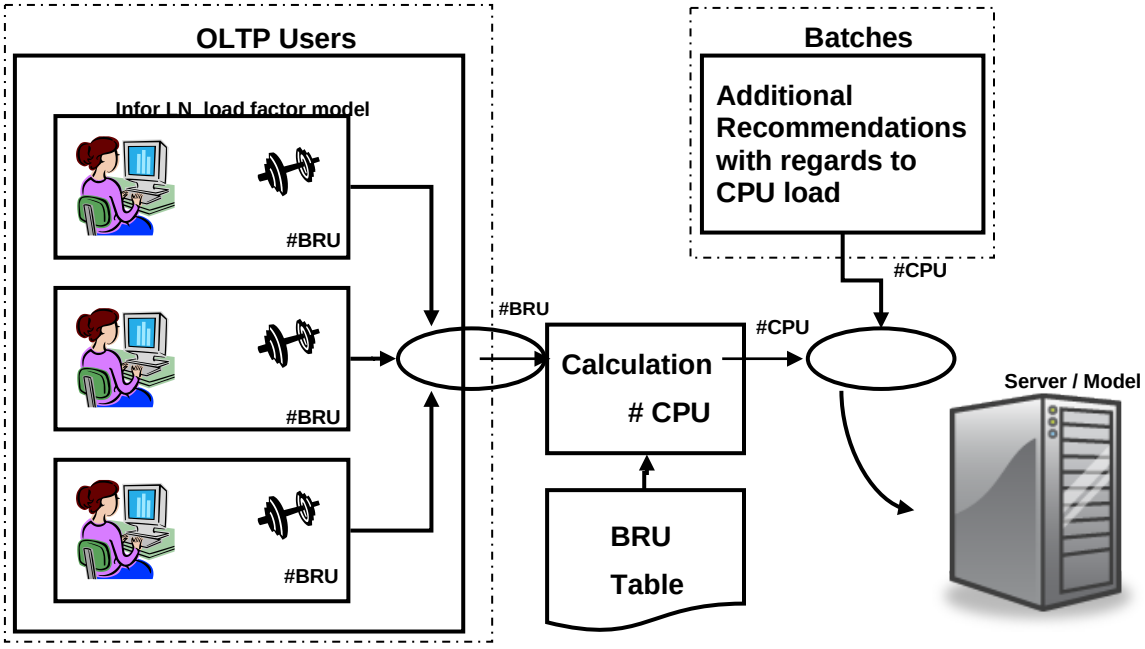

Infor LN users have a different workload. The Infor LN software can be characterized by several different user profiles. Different users characterize different user profiles and generate a particular

load. Currently, these user profiles correspond to the Infor LN packages, including Manufacturing, Distribution, and so on. The number of concurrent users is multiplied with the corresponding load factor.

The total user load that is added to the load of the batches, and the way in which Infor LN is implemented, provides an idea of the system capacity requirements. To select a hardware configuration that can handle this workload, you can use the sizing tables in the sizing guide.

Batches

Batches are programs without user interaction. Batches often run for several hours. Examples of batches include the creation of planning MPS/MRP runs and end of year sessions. Heavy print and query sessions are batches. Integrations that constantly exchange data are also counted as batches.

When the batch runs, every Infor LN batch takes a complete CPU or one thread/core when a physical CPU consists of multiple cores. All components of the previously described porting set are single-threaded applications and will run in order; therefore they require a maximum of one CPU. An exception is batches where parallel bshells are involved.

Batches are difficult to size because the effect on the system is not always clear. The amount of data that must be processed has a considerable effect on the duration. If you run batches during normal system operations concurrently with interactive users, you can use the following guidelines:

UNIX/Linux and Microsoft Windows Platform:

- A Symmetric Multi Processor (SMP) machine is required as a server, both in two-tier and three-tier C/S. An SMP machine is required to prevent slow response times for other Infor LN users.
- For each concurrent batch, an additional CPU is required. In a two-tier environment, one additional processor (CPU) is needed. In a three-tier configuration, for each set of two batches, one CPU must be added to the application server and one CPU must be added to the database server to handle these batch runs.

Exception guidelines for hyper-threading and multicore CPUs

- CPUs use hyper-threading (Intel specific) technology or are equipped with multiple cores. The OS counts these CPUs as the number of threads or cores that exist. Each thread or core can now run a concurrent batch.

Example 1

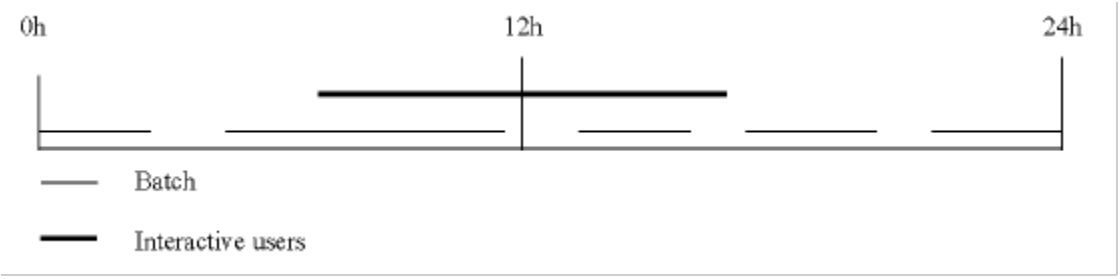

Company A has four batch tasks. One batch runs in the morning, two batches run sequentially in the afternoon, and a large batch runs at night. The company has chosen a two-tier solution. One extra processor (CPU) is enough to handle these batches, because there is no overlap between the batches.

Rule of thumb for average number of concurrent batches

In many presales situations, customers want a high-level estimate of the hardware costs so that they can estimate the overall implementation budget. Often, limited sizing information is available during this stage of the project. In such a situation, you must estimate the number of concurrent batches. The following rule of thumb is applicable for up to 200 concurrent users:

- Per 50 concurrent users, reserve capacity for one batch.

CPU speed and cores

Using high clock speed processors will result in the runtimes of these jobs being reduced.

CPUs use hyper-threading technology (Intel) and are equipped with multiple cores. The OS counts these CPUs as the number of threads or cores that exist; each thread or core can run a concurrent batch job.

Other CPU Performance factors

Other factors can affect the performance ratings of the systems.

- The Infor LN audit server can track and record changes to the database. In most implementations, auditing is used primarily on parameters and critical business tables to log user changes and timestamps to the audit database. If all the database tables are audited, there is an additional load of approximately seven percent.
- Field level authorization rule processing and filtering can degrade performance by 20 percent.

- There are processes, such as enterprise planning, that can be configured to use the multiple virtual machine feature to multithread typically single-threaded processes. Multiple virtual machines allow the work load to span multiple virtual machines that then run on multiple CPU processor cores to complete the task.

General

Infor carries out benchmarks to collect sizing information. For more information, refer to the 0. The tested system during the benchmark is said to be saturated when user response times exceed a pre-determined threshold. At this point in time, the number of sessions and the mix of session are reported, which provides a measurement of load.

Therefore, Infor LN sizings are sized to the highest expected load scenario for customers' systems. In addition to the online transaction processing, concurrent batch, integration or reporting jobs must also be considered. The peak load on the server is estimated at 70 percent of capacity on the server based on industry best practices regarding sizing.

Scalability for the application servers can be accommodated by adding additional servers to the hardware architecture. For the database server, scalability is typically addressed by adding processors and memory to the existing database. Distributed databases are supported by Infor LN, but most clients find the management of distributed databases to be administratively intensive.

Chapter 4 Sizing information

4

Memory Requirements

Memory is an important factor during the sizing. Insufficient system memory can lead to poor performance. Memory figures in this chapter are based on **concurrent** users. In previous sizing guides, the formulas were based on **connected** users. This has changed because too many mistakes were made, which led to under sizing.

The memory requirements for Infor LN are estimated to be 120 MB per concurrent user for a 2-tier environment. In a 3-tier environment, the application server would require 95 MB per concurrent user and the database server would require 48 MB per concurrent user. The estimated requirement for the operating system is 4 GB.

The described formulas are based on three tasks per user, such as entering sales orders or production orders. For each additional task, an extra amount of internal memory per user must be added.

To reduce the risk of poor system performance, you are advised to have a server with more than 4 GB of memory. In practice there can be tools, daemons, agents, and other installed software which are not taken care of in the sizing.

Network Requirements

Network requirements are divided into two components: user to application server and application server to database server.

Client – Application Server Connection

The user to application server connection is defined in bandwidth requirements based on a concurrent user and is defined in kilobits per second (Kbps).

The main factor that limits network performance is latency. For acceptable performance, Infor recommends that a maximum network latency of 200 ms round trip is required. Network latency between 200–400 ms is noticeable on response times, but workable for most situations. Network latency above 400 ms is not recommended.

The minimum bandwidth that is required for the Infor LN application including the UI tier, as tested during network performance benchmarks, is 4 KB (40 Kbps) on a minimum 256 KB line.

Some margin must be kept for peak loads, including print output.

Application Server – Database Server Connection

The bandwidth that is required between the application and database servers is highly dependent on the applications that are running and the number of connections. There are no metrics here to compare the required bandwidth.

Most customers have installed a separate LAN between the application server and the database server. The LAN is typically a Gigabit LAN.

In a 2-tier environment using backplane speeds, the performance of this connection is generally measured to be 20–40 percent better than connecting through dedicated network cards, which is recommended for a 3-tier environment.

Disk and IO Requirements

50 GB are recommended for the application server/database server. An estimate of the database growth is 0.6 GB per user per year.

If more users use the system, more disks are required to reach performance requirements. For an idea of the minimum number of disks for data, the following rule of thumb can be used:

$$I_{tot} = n/20 + 4$$

Where:

I_{tot} = Number of disks required for data storage

n = Number of concurrent users

More and more SAN and NAS solutions are used instead of solutions based on separate disks. The SAN/NAS is often shared with other applications. You cannot influence the number of disks;

therefore, the number of IOs per second is a more useful parameter for disk sizing. As a rule thumb, a modern disk can handle 200 sequential or 100 random IOs. The IO characteristic of Infor LN is generally random. The number of required IOs can be calculated as follows:

#disks*100

or

$I_{IOs} = (n/20 + 4) * 100$

Where:

I_{IOs} = Total number of IOs

n = Number of concurrent users

Infor LN UI Requirements

For sizing information, deployment guidelines and tuning of Infor LN UI, your sales representative can provide document *Infor Ming.le - LN UI plugin Sizing Guide (B0080 US)*.

Sample Infor LN configuration

For some sample hardware configurations see *Infor LN Minimum Hardware Requirements on Window/SQL Server (B0071 US)*. There are examples for 1–25 concurrent users, 25–50 concurrent users and 50–100 concurrent users.

For larger configurations or configurations based on other platforms contact your Infor representative.