
Licensing Handbook



for OpenSpeech Recognizer 2.0
and Speechify 3.0

Document History

Date	Release Name
July 2003	Second Edition – for OpenSpeech Recognizer 2.0 and Speechify 3.0
June 2002	First Edition, for OpenSpeech Recognizer 1.1 GA

Notice

Copyright © 2001–2003 by SpeechWorks International, Inc. All rights reserved.

The information in this document is subject to change without notice.

Use of this document is subject to certain restrictions and limitations set forth in a license agreement entered into between SpeechWorks International, Inc. and purchaser of the OpenSpeech Recognizer software. Please refer to the SpeechWorks license agreement for license use rights and restrictions.

Portions of the OpenSpeech Recognizer Software are subject to copyrights of AT&T Corp., E-Speech Corporations, and Bell Communications Research, Inc.

SpeechWorks is a registered trademark, and OpenSpeech, Speechify, Speechify Solo, SpeechSecure, SpeechSite, SpeechWorks Here, DialogModules, SMARTRecognizer, the SpeechWorks logo, and the SpeechWorks Here logo are trademarks of SpeechWorks International, Inc. in the United States and other countries. All other trademarks are the property of their respective owners.

Published by:

SpeechWorks International, Inc.
695 Atlantic Avenue
Boston, MA 02111

Table of Contents

Preface	v
About the OpenSpeech Recognizer	1-vi
Roadmap to available documentation	1-vii
Support services	1-viii
 1. About SpeechWorks Licenses	
Overview of license options	1-2
Evaluation Software Development Kit (SDK) licenses	1-2
SDK licenses (perpetual licenses)	1-2
Runtime licenses	1-2
OSR speech licenses	1-3
Speechify licenses	1-3
Obtaining and managing licenses	1-4
Viewing licenses	1-4
Generating licenses and downloading license files	1-4
Making changes to generated licenses	1-5
Adding licenses during operation	1-6
Modifying (replacing) licenses	1-6
Dealing with license problems	1-7
 2. License Architectures	
License server configurations	2-9
Distributed license servers and floating licenses	2-9
Failure recovery	2-10
License servers and system architectures	2-12
OSR all-in-one (in-the-skins) architecture	2-12
OSR client-server architecture	2-13
Speechify client-server architecture	2-15
Configuration tasks for the system administrator	2-16
Balancing license server load	2-18
Redundant license servers	2-18
Server-configured redundancy	2-19
Client-configured redundancy	2-23
Combining server-configured and client-configured redundancy	2-23
Sharing a license server with OSR and Speechify	2-25

Platform integration design for licensing	2-26
Overview of licensing modes.....	2-26
How to set licensing modes.....	2-27
OSR licensing modes	2-27
Speechify licensing modes	2-27
Default OSR licensing	2-27
Explicit OSR licensing.....	2-28
Default Speechify licensing.....	2-30
Explicit Speechify licensing.....	2-30

3. Configuring Licensing on Windows

License server installation	3-33
Configuring and starting the license server	3-35
Manually starting and stopping the license server.....	3-36
Configuring OSR on Windows.....	3-37
Setting up SWILicenseServerList	3-37
Configuring Speechify on Windows	3-38
Configuring Speechify manually	3-38

4. Configuring Licensing on Linux

Starting the FLEXlm license server	4-41
Manually starting and stopping the license server.....	4-41
Automatically starting the license server	4-42
Creating log files.....	4-42
Running the license server on a separate host.....	4-43
Configuring OSR on Linux.....	4-44
Setting up SWILicenseServerList	4-44
Configuring Speechify on Linux.....	4-45

Appendix A: License Server Host IDs

Appendix B: License Files

Filename and status after initial installation	B-50
Sample license file for OSR.....	B-51
Sample license file for Speechify.....	B-52
Sample shared (merged) license file for OSR and Speechify	B-53



Preface

This guide describes the details of license enforcement and is written for:

- ❑ Integrators who develop a platform for use with the SpeechWorks® OpenSpeech™ Recognizer, Speechify™ text-to-speech system, or both
- ❑ System operators who set up and administer those systems

OSR and Speechify use third-party license management software to allocate and free licenses. For OSR, you need licenses for each telephony channel (or port) that initializes a recognizer (SWIrec) or speech detector (SWIep); you do not need separate licenses for each component. For Speechify, you need licenses for each port that might simultaneously play Speechify generated audio, and licenses for extra voices. SpeechWorks controls the initial generation and distribution of these license files (also known as keys).

New and changed information

Second edition

For OSR, license functionality was first introduced for the 1.1 release. There are no significant changes to the functionality for the OSR 2.0 release except that DTMF-only licensing has been removed from the product.

All Speechify information is new for this edition. For Speechify, licensing was introduced for the Speechify 3.0 release.

Updated information:

- ❑ Added explanations and illustrations throughout this guide.
- ❑ Merged in information about OSR Client-Server licensing.
- ❑ Merged in information about Speechify licensing.
- ❑ Clarifications about license server quorums. See [page 2-20](#) and [2-24](#).
- ❑ Added path to show where to find lmutil. See [page A-47](#).

Welcome to SpeechWorks

OpenSpeech Recognizer (OSR) is a speaker independent, continuous speech recognition software system. It can be used as a component of a speech recognition application in conjunction with other computer telephony functions such as text-to-speech. OSR allows a computer equipped with a voice processing board or other telephony interfaces to accurately recognize and process a caller's spoken response or DTMF (TouchTone) input. OSR consists of a speech recognition engine, a comprehensive API library, sample programs, and a complete set of documentation to enable you to develop a speech recognition application quickly and easily.

About the OpenSpeech Recognizer

OpenSpeech Recognizer is a SpeechWorks product that provides a pure recognition interface for VoiceXML browsers and for non-VoiceXML platforms that need access to a low-level recognition resource. Major attributes of this product include:

- ❑ “All-in-one” or client-server architecture.
- ❑ An interface for controlling recognition and managing grammars.
- ❑ An interface for delivering audio to the recognizer.
- ❑ An interface for speech detection that can be separated from the recognizer
- ❑ A wide variety of built-in grammars.
- ❑ Event and waveform logging.
- ❑ Offline grammar compilation and debugging tools.
- ❑ Support for recognition of multiple and simultaneous languages.
- ❑ Acoustic adaptation to automatically improve recognition accuracy
- ❑ Optimized features for VoiceXML browsers: URI resolution, SRGS grammar format,¹ XML result format, parallel grammars, and grammar caching.

The run-time component of the OpenSpeech Recognizer supports a C API which can be used by C or C++-based programming environments.

The OpenSpeech Recognizer API includes C functions for the following operations (see the *OSR Reference Manual*):

- ❑ Initialization and object creation of recognizer and speech detector resources.
- ❑ Configuration and management of grammars.
- ❑ Acoustic state input and output
- ❑ Configuration of the recognizer and logging of recognition activities.
- ❑ Delivery of audio to the speech detector and receipt of status results.
- ❑ Delivery of audio to the recognizer and receipt of status results.
- ❑ Invocation of recognition and delivery of recognition results.

Roadmap to available documentation

Recommended reading:

Self-starters can learn about SpeechWorks technologies by using the product documentation. The following documents are available for OpenSpeech Recognizer:

- ❑ The *OSR Client/Server Operator's Guide* provides a central location for OSR client-server concepts, installation, and configuration.
- ❑ The *OSR Developer's Guide* describes grammars: using the built-in grammars and creating your own.
- ❑ The *OSR Licensing Handbook* describes the process for getting licenses to run the OSR software, details about configuring your license server, and related topics.
- ❑ The *OSR 2.0 Migration Guide* is required reading for application developers and platform integrators who are currently working with OSR 1.*n* systems.
- ❑ The *OSR Platform Integration Manual* provides details on integrating the Recognizer with your voice processing platform or VoiceXML browser. It describes the required functionality of the Recognizer, and gives suggestions for managing grammars, caching, and logging.
- ❑ The *OSR Reference Manual* provides details about OSR speech detector and recognizer API functions and configuration parameters.

1. SRGS is the Speech Recognition Grammar Specification Version 1.0 (W3C Candidate Recommendation 26 June 2002)

- ❑ The *OSR Release Notes* cover specific issues for a release, including descriptions of new features, known functional constraints, and migration issues.

The following documents are available for developers using Speechify™ text-to-speech applications:

- ❑ The *Speechify User's Guide* provides installation, programming, and reference information about the Speechify product.
- ❑ There is a *Speechify Language Supplement* for each supported language. These supplements contain language-specific reference information for application developers.
- ❑ The *Speechify 3.0 Migration Guide* is required reading for application developers and platform integrators who are currently working with Speechify 2.*n* systems.

Review the release notes distributed with each product for the latest information, restrictions, and known problems.

Support services

To receive technical support from SpeechWorks International, Inc.:

- ❑ Visit the Knowledge Base or ask a question at:

http://www.speechworks.com/training/tech_support.cfm

- ❑ Ask for “technical support” at +1 617 428-4444

How this guide is organized

The chapters of this document cover these topics:

Chapter 1 covers the types of licenses available and the process for obtaining and managing licenses, and describes the license files.

Chapter 2 covers license architectures and enforcement, including allocating and freeing licenses implicitly and explicitly, and load balancing.

[Chapter 3](#) describes how to install and configure a Windows license server.

[Chapter 4](#) describes how to install and configure a Linux license server.

[Appendix A](#) describes how to find the hostid for your license server machine.

[Appendix B](#) describes a sample license file.

This guide also includes an index.



About SpeechWorks Licenses

This chapter discusses the various types of licenses and how to get them.

Legal ramifications

Your use of SpeechWorks products is governed by your license agreement with SpeechWorks. Neither this document nor the actual mechanisms of license key enforcement are a substitute for that license agreement.

General requirements

The general license requirements are as follows:

1. Obtain a license file from the SpeechWorks licensing fulfillment website. For information, see “Obtaining and managing licenses” on [page 1-4](#).
2. Install the FLEXlm license server software (included with the SpeechWorks installation media) on a machine capable of communicating with the OSR or Speechify machine via TCP/IP. All references to the “license server” in this handbook refer to the FLEXlm license server. See [Chapter 2](#) for a discussion of where to install the license server.
3. Run the license server and configure the system. See the discussions in [Chapter 3](#) (Windows) and [Chapter 4](#) (Linux).

Overview of license options

Evaluation Software Development Kit (SDK) licenses

An evaluation license permits a single developer to use the SpeechWorks SDK for a period of 45 days from the date of issue. For OSR, you are allotted 4 ports of recognition and speech detection (endpointing).

For Speechify, you are allotted a speak license and a voice license for every voice that is currently available. Speechify's standard evaluation license is limited to 4 ports for a period of 45 days from the date of issue.

There is also an extended evaluation option for both products, providing n ports (where n can be any number) for a period of 90 days.

These ports may only be used for *internal* development. To continue development with the OSR or Speechify SDK after the evaluation period has expired, you must license a full SDK. To deploy a speech system, you must purchase the desired number of runtime ports.

SDK licenses (perpetual licenses)

A full SDK license permits a single developer to use the SpeechWorks SDK for an unlimited period of time. For OSR, you are allotted 4 ports of recognition and speech detection. For Speechify, you are allotted 4 ports of speaking. These ports may only be used for *internal* development. To deploy a production system, you must license the desired number of runtime ports.

Runtime licenses

Runtime licenses apply on a per-port basis:

- ❑ For OSR, they must be purchased for every channel of speech detection or recognition that will be deployed in a speech system. (Each license covers speech detection and recognition even though these are separate components.)

- ❑ For Speechify, they must be purchased for every voice and for every channel of output speech. Runtime licenses do not expire.

OSR speech licenses

OSR speech licenses allow you to invoke the recognizer and speech detector. The licenses grant full access to the recognizer and speech detector, and you can recognize from both voice and DTMF grammars.

Speechify licenses

Speechify licenses provide the ability to activate Speechify server ports. *Standard* Speechify licenses permit *n* simultaneous playbacks of Speechify generated audio for a single voice. *Uplift* licenses cover additional voices. Thus, Speechify has two different license types:

- ❑ Speak licenses limit the number of active server ports *across all voices*. The number of active Speechify ports across all voices cannot exceed this limit.
- ❑ Voice licenses limit the number of active server ports *for a specific voice*. The number of active Speechify ports for a specific voice cannot exceed this limit.

For example, if you purchase 96 ports of Speechify Tom U.S. English, you might pay an “uplift fee” on just 24 of those ports to obtain 24 ports of Speechify Karen Australian English. From a licensing perspective, the following scenarios are valid:

- ❑ 96 active ports playing Tom
- ❑ 24 active ports playing Karen
- ❑ 24 active ports playing Karen and 72 active ports playing Tom
- ❑ 1 active port playing Karen and 95 active ports playing Tom

These scenarios are invalid; they violate the licensing terms and return an error code from the appropriate API function:

- ❑ 97 active ports playing Tom (more than 96 simultaneous playbacks)
- ❑ 25 active ports playing Karen (more than 24 ports playing Karen playbacks)
- ❑ 24 active ports playing Karen and 73 active ports playing Tom (more than 96 simultaneous playbacks)

Obtaining and managing licenses

Each SpeechWorks customer agreement includes provisions for specific accounts and numbers of licenses. SpeechWorks sends License Authorization Codes (LACs) to each customer, and each LAC enables requests for license files for an account. The codes are sent via e-mail.

Viewing licenses

To view licenses, follow these steps:

1. Go to the License Fulfillment website: <http://licensing.speechworks.com> and enter the appropriate LAC (see step 1 on [page 1-5](#) for more details).

The “License Fulfillment Homepage” displays your license status in a text box:

- Pending – means that licenses exist but no license file has been generated.
 - Partial – means a license file has been generated for some ports. For example, if the original order is for 72 ports, you might generate licenses for 24 ports, and return later to generate licenses for additional machines.
 - Completed – means that license files have been created for all of the ports associated with this order. For example, if your OSR license agreement specifies 72 ports of speech detection and recognition, you can generate license files for any number of ports up to that maximum. (See “License server configurations” on [page 2-9](#) for information about using license servers.)
2. Choose the appropriate product category (OSR or Speechify), and click “View or Replace Licenses.” In response, the site displays currently generated licenses and hosts, and you can view and download the license files (as described above) by clicking “View the license” for each license.

Generating licenses and downloading license files

To use licenses, you must generate a license file. During generation, you associate licenses with a specific hostid.

To generate licenses, follow these steps:

1. Go to the License Fulfillment website: <http://licensing.speechworks.com> and enter the appropriate LAC. (Your company might have one LAC or it might have several LACs for different accounts. For example, the Speechify product uses a different LAC from OSR and a different LAC for each voice.)
2. The Licensing Fulfillment webpage shows you the products for which you can generate licenses. (There may be more than one product in your order.) Choose the appropriate product category (OSR or Speechify), and click Generate License.
3. Enter the hostid for your license server (hostid is also discussed in [Appendix A](#)), choose the desired quantity of licenses for your license server, and click “Generate License”. The next webpage displays the license file and instructions for saving the license to a local file.
4. Install the license file (with the expected name) in the directory specified in “Configuring and starting the license server” on [page 3-35](#).

Making changes to generated licenses

After you generate a license file, you can make changes in the future:

- If you generated some of the available licenses, you can return to the license fulfillment page and generate additional licenses.
 - To generate new licenses for a new license server (that does not have an existing license file), follow the steps in “Generating licenses and downloading license files”.
 - To generate new licenses for a license server that already has a license file, you must merge the new and existing license files. This is discussed below in “Adding licenses during operation”.
- If you have more licenses than needed for a license server, you can return those licenses (and free them for use on another server). This is discussed below in “Modifying (replacing) licenses”.

Adding licenses during operation

After you generate, install, and configure licenses for a license server, you can generate additional licenses for that server in the future. The steps are identical to the discussion in “Generating licenses and downloading license files” on [page 1-4](#) except that you finish by merging the new license file with the existing one. To summarize the steps:

1. Go to the License Fulfillment website and enter the appropriate LAC.
2. Select the desired licenses and click Generate License.
3. Enter the hostid for your license server, choose the desired quantity of licenses, and click “Generate License”. The next webpage displays the license file and instructions for saving the license to a local file.
4. Merge the new license file into the existing file in the target directory. To merge the files, open the files in a text editor and copy the INCREMENT lines from the new file to the existing. (Use the example on [page B-53](#) as a model for merging files.)

Modifying (replacing) licenses

You can return licenses that have been previously generated, and those licenses again become available for re-generation. For example, this is useful when replacing license server hardware because you can return licenses from the old server and re-generate them for the new.

When you replace licenses with this procedure, the license website creates a new License Authorization Code (LAC) for the replaced licenses, and the number of licenses associated with the original LAC is reduced by the same quantity:

1. From the homepage of the License Fulfillment website, click “View or Replace Licenses.” In response, the site displays currently generated licenses and hosts.
2. Click the appropriate “Replace the license” button. In response, the site requests an License Authorization Code, the license file, and an e-mail address for acknowledgement and further instructions. (The fields for the LAC and filename are automatically filled based on your previous selection.)
3. Fill in the text boxes, click “Submit,” and then confirm the information in the next screen that appears. In response to your submission, the SpeechWorks license database will be updated. You will receive the new LAC via e-mail, which can be used to generate the new replacement licenses.

Dealing with license problems

If you have problems obtaining or running with licenses, file an incident report and SpeechWorks will respond rapidly. To file a report, visit this website:

<http://www.speechworks.com/training/techsupport/private/login.cfm>



License Architectures

This chapter discusses system architecture options available for deployment:

- ❑ “License server configurations” (see below)
- ❑ “Platform integration design for licensing” on [page 2-26](#)

License server configurations

All references to the “license server” in this handbook refer to the FLEXlm license server.

Distributed license servers and floating licenses

SpeechWorks uses FLEXlm to implement a floating license model:

- ❑ License servers allocate licenses to license clients, which run on OSR or Speechify machines. There is no limit to the number of license servers that can be used on a network.
- ❑ A fixed number of licenses is available; this is known as the license “pool.” The pool is created when a license server starts and destroyed when the server stops. A single license server or many license servers can manage a given pool.

- ❑ License clients request licenses when a recognition or TTS event is needed. Licenses are not required to be dedicated to specific ports; they can be “floating” (any license can be allocated to any port). They can be released when the license is no longer needed.

One advantage of this model is that a single license server can manage licenses across a network of multiple license clients without knowing the hostid for each client. (Instead, the license clients point to the IP addresses of their license servers.) Thus, a pool of recognition servers and TTS servers can share the same license pool.

Failure recovery


If OSR suddenly loses communication with its license server, it continues processing telephone calls normally for a period of 61-89 seconds. If a license server returns during that period, operation continues normally. If a license server does not return, OSR releases all its licenses and can no longer perform speech detection or recognition events for telephone calls until the license server again becomes available.

Likewise, if Speechify suddenly loses communication with its license server, it continues processing speak requests normally for 61-89 seconds. If a license server returns during that period, operation continues normally. If a license server does not return, Speechify releases all its licenses and can no longer perform speak requests.


Details on the 61-89 timeout second period:

- ❑ License clients and servers periodically communicate via a “heartbeat” mechanism: every 30 seconds, each license client sends a heartbeat message to the license servers.
- ❑ If a corresponding message is not returned from the server for two consecutive heartbeats, the client attempts to reconnect.
- ❑ Upon a second unanswered heartbeat, the client releases all its licenses, and all calls to API functions that require licenses return `SWIrec_ERROR_LICENSE_COMPROMISE` (for OSR) or `SWItts_NO_LICENSE` (for Speechify).

The following diagrams show the shortest and longest possible timeout periods until a reconnect completed (or until the compromise state is reached). The range is 61 to 89 seconds depending whether the network license server failure occurs immediately before or after a heartbeat. Here is the shortest timeout:



Here is the longest possible timeout period:



The abbreviations in the preceding figures indicate the following:

- ❑ F = failure
- ❑ H0 = heartbeat occurs before a disconnect (no effect)
- ❑ H1 = heartbeat 1, attempted reconnect 1
- ❑ H2 = heartbeat 2, attempted reconnect 2
- ❑ H3/C = heartbeat 3, enter compromise state
- ❑ t = time, in seconds

Once compromised, OSR and Speechify will remain in this state until a license server *with a sufficient number of licenses* becomes available. Once the license server returns to service, OSR (or Speechify) will resume normal operation (by automatically re-assigning licenses) after a period of time equal to the timeout period.

License servers and system architectures

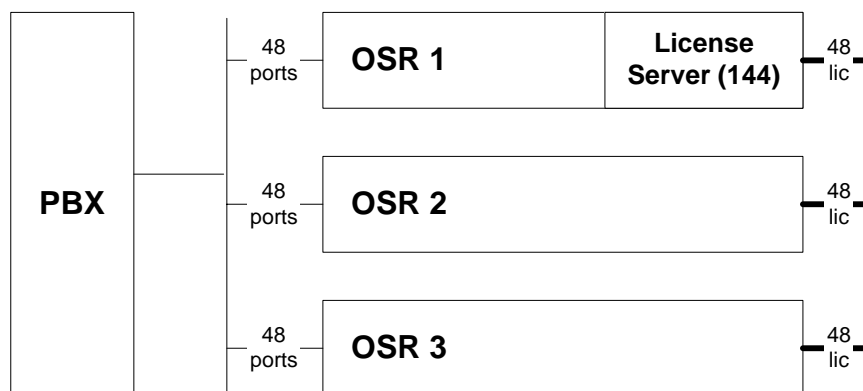
There are minor configuration differences for license servers depending on the product (OSR versus Speechify) and the architecture (all-in-one versus client-server).

The sections below describe these differences briefly, and are followed by a discussion of system administrator configuration tasks ([page 2-16](#)).

OSR all-in-one (in-the-skins) architecture

OSR “all-in-one” (or alternatively, “in-the-skins”) is an architecture where all of the components of the speech detector and recognizer are running on the same computer.

The illustration below shows an example system running three copies of OSR (one copy per machine) with a maximum of 48 ports of speech recognition each. One way to deploy this system, is to run a single license server on one of the OSR machines:



When the license server starts, it creates a pool of licenses (in this example, 144 licenses). The server allocates individual licenses to ports controlled by license clients on a first-come, first-serve basis. To perform the allocation, the license server uses “default” or “explicit” mode as described on [page 2-26](#).

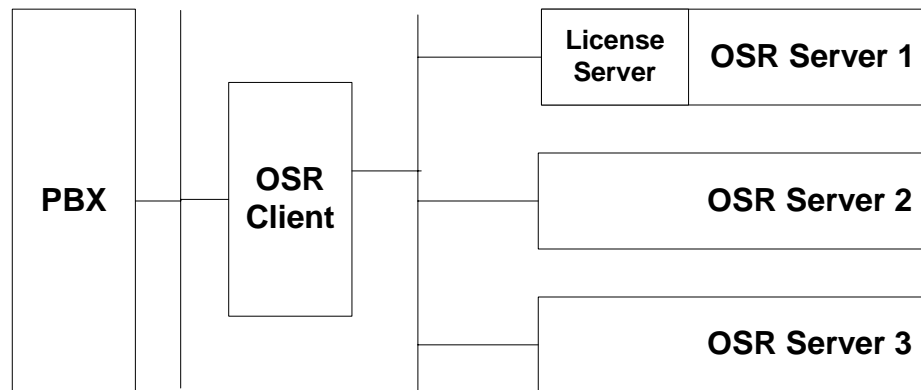
The license server maintains a count of allocated licenses across the network. It does not monitor port usage on a per-machine basis, but rather identifies that there cannot be more than 72 ports active at any one time. In this example, the equal distribution of 48 licenses per machine is arbitrary (having more licenses than available ports is unnecessary; having less results in under-utilization of resources).

The license server can run on any of these OSR machines or on another machine in the network. *The license server must be on the same logical subnet as the OSR machines.*

OSR client-server architecture

OSR “client-server” is an architecture where the speech detector and recognizer are distributed across a network of two or more computers running the TCP/IP protocol.

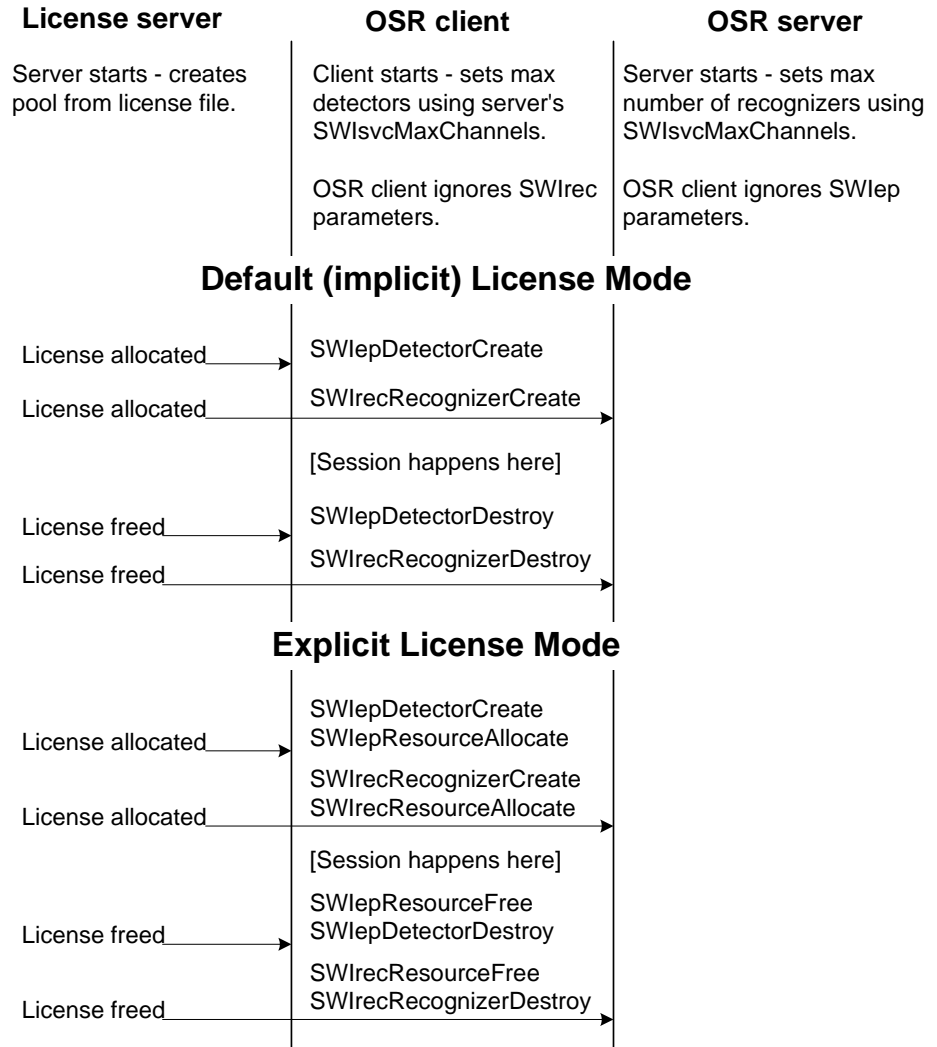
The licensing model for the OSR client-server environment is essentially the same as the preceding all-in-one approach. The example below shows an OSR client using three OSR servers. One of the OSR servers is also running a License server:



Although the speech detector and the recognizers are on different machines in the OSR client-server environment, this does not create licensing problems because license servers allocate licenses independently for speech detector and recognizer ports. (This does not reduce the number of licenses by half because each license file grants an equal number of licenses for both.)

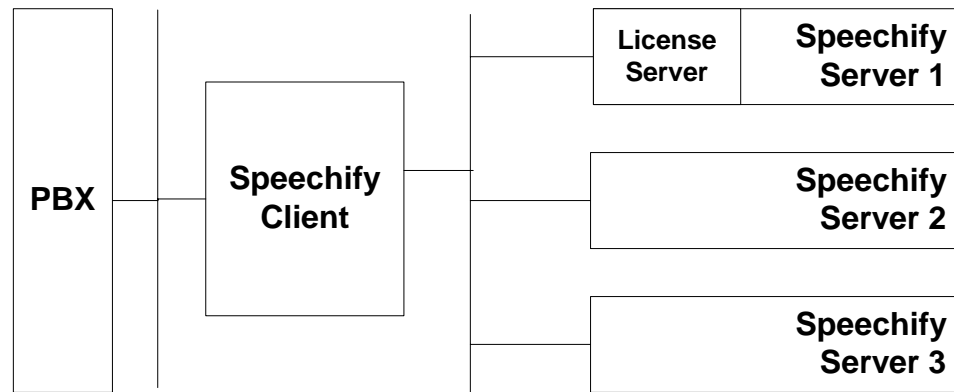
When the license server starts, it creates a pool of licenses. The server allocates individual licenses to ports controlled by license clients on a first-come, first-serve basis. To perform the allocation, the license server uses “default” or “explicit” mode as described on [page 2-26](#).

The license server maintains a count of allocated licenses across the network. It does not monitor port usage on a per-machine basis, but rather identifies that there cannot be more than 72 ports active at any one time.



Speechify client-server architecture

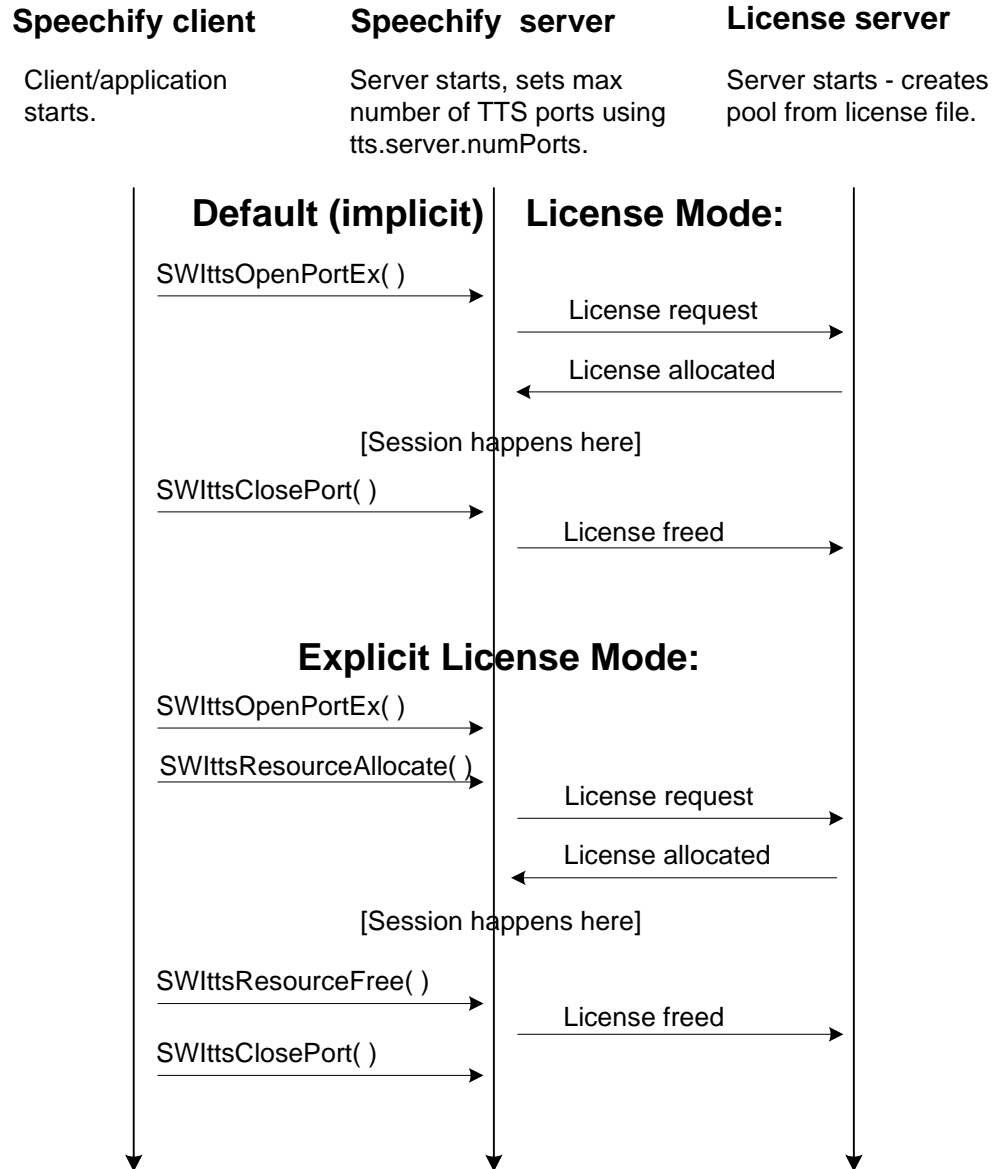
The licensing model for Speechify is essentially the same as for the preceding OSR architectures:



The Speechify client does not require licenses.

On the Speechify server, licenses are allocated differently depending on the licensing mode (see [page 2-26](#)):

- ❑ Default (implicit) mode – The license server *allocates licenses as Speechify clients connect to the server* (one license per client `SWIttsOpenPortEx()` operation). The licenses are freed when the client disconnects or the Speechify server stops.
- ❑ Explicit mode – The platform calls `SWIttsResourceAllocate()` to explicitly *allocate individual licenses* at the start of each session. The platform frees the licenses at the end of each session by calling `SWIttsResourceFree()`.



Configuration tasks for the system administrator

Tasks for OSR all-in-one

For configuration, the administrator must do the following:

- ❑ Provide the hostid of each computer running the license server to SpeechWorks (sent via the website). See “Generating licenses and downloading license files” on [page 1-4](#).

- ❑ Provide the host name (not the hostid) of the computer running the license server to all the machines running OSR. The name is either entered into the SpeechWorks.cfg or into an environment variable. See “Configuring OSR on Windows” on [page 3-37](#) or “Configuring OSR on Linux” on [page 4-44](#) for details .
- ❑ Decide and configure the desired license mode for the speech detector and the recognizer. See “Overview of licensing modes” on [page 2-26](#).

Tasks for OSR client-server

For OSR client-server configuration, the only difference from OSR all-in-one is that the licensing mode is configured on both the OSR client and OSR server machines

- ❑ On the OSR client, configure the license mode for the speech detector.
- ❑ On the OSR server, configure the license mode for the recognizer.

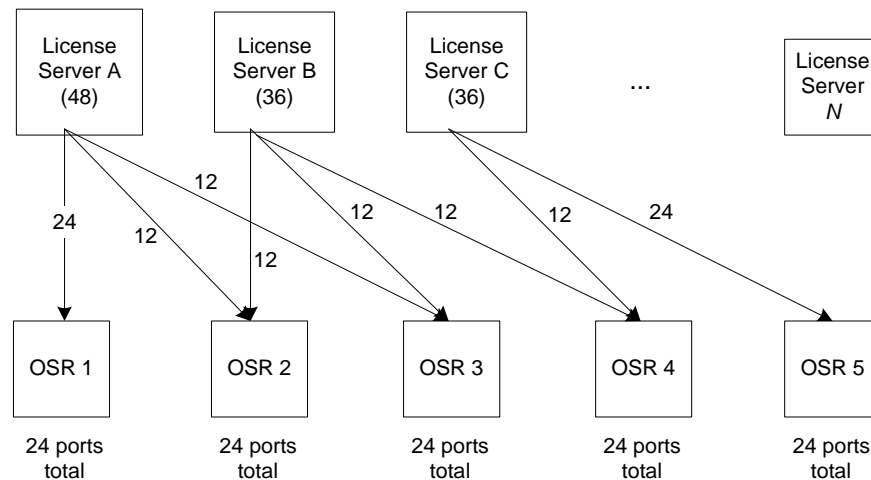
Tasks for Speechify client-server

For Speechify client-server configuration, the only difference from OSR all-in-one is that the licensing mode is configured on the Speechify server machine.

Balancing license server load

Large systems that require many license servers can implement a load balancing architecture to ensure that a primary license server does not become overloaded and thus cause delays.

Load balancing is accomplished by limiting the number of licenses each license server on the network can issue. The network administrator must input the host names of all license servers into the `SpeechWorks.cfg` file or environment of each OSR or Speechify system. (Note that the picture indicates OSR but applies equally to Speechify.)



In addition, the above configuration can be made redundant (see next section below).

Redundant license servers

If you choose a single license server architecture, you risk that a failure of the host machine (or of the license server daemon itself) will introduce a single point of failure in your license network and deny licenses to all machines.

To avoid this risk, FLEXlm offers the ability to support redundancy in two different ways: client-configured redundancy, and server-configured redundancy:

- ❑ For client-configured redundancy, the licenses are divided into multiple license pools, each administered by a single license server.

- ❑ For server-configured redundancy, three license servers administer a single license pool.

You can employ both methods together for maximum redundancy.

Server-configured redundancy

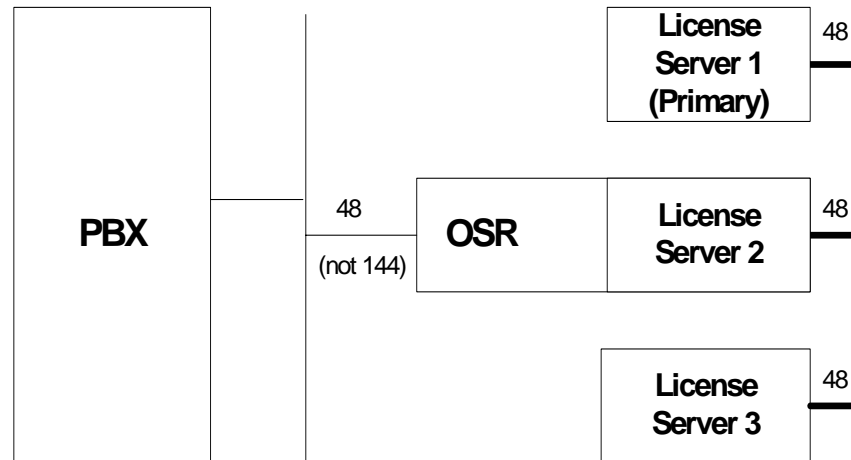
Server-configured redundancy requires three license servers operating in what is known as a quorum. While sharing a single license pool, each of these license servers must be on a separate box, and must also be on the same logical subnet. (This means that two additional computers are required when providing redundancy for a license server.) The entire license pool remains available as long as a quorum of two of the original three license servers is running normally.

Characteristics for server-configured redundancy:

- ❑ The primary server does not need to be a dedicated machine; however, any other processes running on the machine will reduce the memory and CPU cycles available to the server. Also, failure of other processes could affect the functioning of the primary server.
- ❑ To avoid any memory constraints with respect to any SpeechWorks product (OSR or Speechify), avoid defining the primary license server on a machine that is running that product.
- ❑ The secondary and tertiary license servers can be located on SpeechWorks product machines or on another computer on the subnet. The backup license servers are seldom used, and they do not need to be dedicated machines.
- ❑ If a license server fails, the FLEXlm logs indicate that licenses were provided from a backup license server. These logs are distinct from SpeechWorks logs. Because the redundancy is invisible to the SpeechWorks software, no SpeechWorks logs indicate that there was a failure.

Example of simple server-configured redundancy

The following illustration shows a simple configuration for server-configured redundancy. This is an OSR all-in-one system (but the picture applies equally to Speechify).



Any 3 machines on the same subnet can be formed into a quorum:

- ❑ When creating the license file, enter the hostid for all 3 machines. In the illustration, the quorum is formed by servers 1, 2, and 3 for a total of 48 ports.
- ❑ When configuring `SWILicenseServerList` for OSR or `tts.server.licenseServerList` for Speechify, enter the 3 machines in the same order as in the license file.

For example, if your quorum machine names are `nicosia`, `arctic`, and `nepal` (and this is the order they appear in the license file), then they might appear as follows in the `SWILicenseServerList` for OSR:

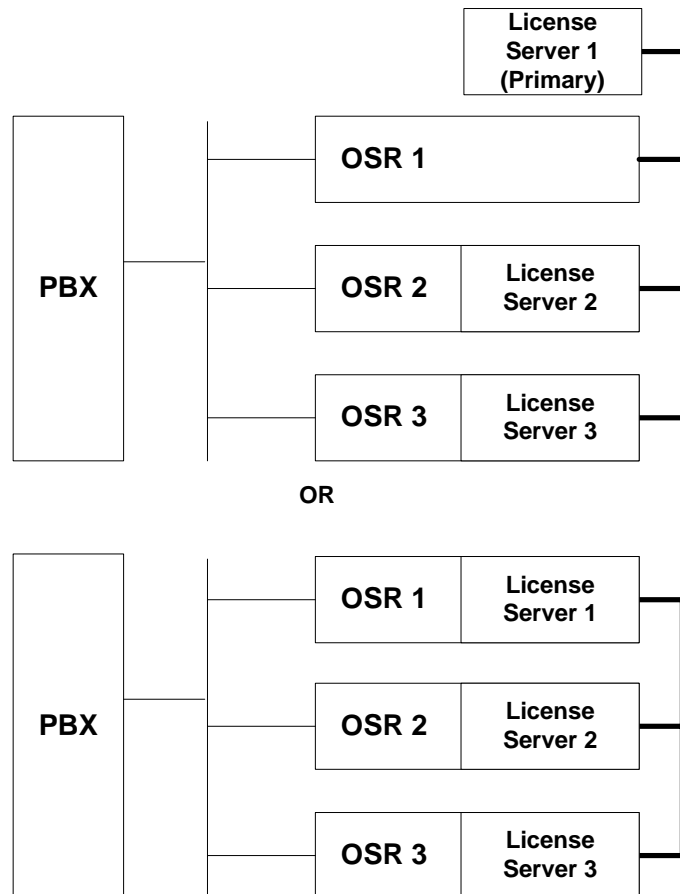
```
28000@nicosia;28000@arctic;28000@nepal
```

In the Speechify configuration file:

```
<param name="tts.server.licenseServerList">  
  <namedValue name="1"> 28000@nicosia </namedValue>  
  <namedValue name="2"> 28000@arctic </namedValue>  
  <namedValue name="3"> 28000@nepal </namedValue>  
</param>
```

Example configuration with multiple OSR recognizers

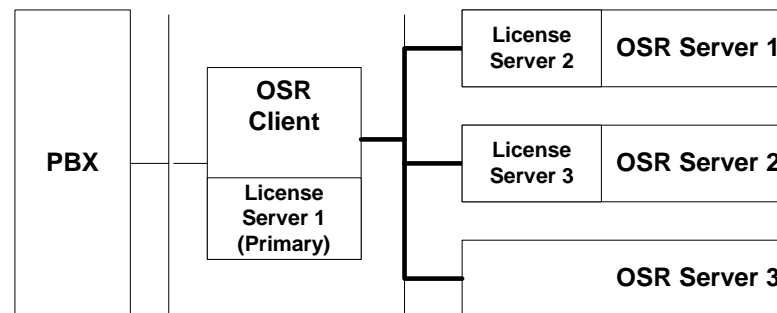
This following illustration shows two configurations for server-configured redundancy. These are all-in-one systems. The first example is preferred (with the primary server on a separate machine):



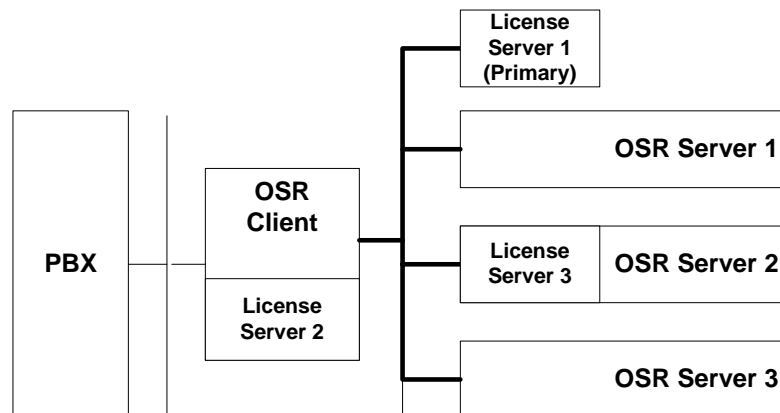
Example configuration with Speechify or OSR Client-Server

Licensing for Speechify or OSR client-server is similar, but with more choices because you can run license servers on the client, server, or on a separate machine. In every case, the license server provides licenses for both the endpointer (SWIepAPI) and the recognizer (SWIrecAPI).

Two OSR examples are shown below (but the pictures apply equally to Speechify). The second example is the preferred configuration:



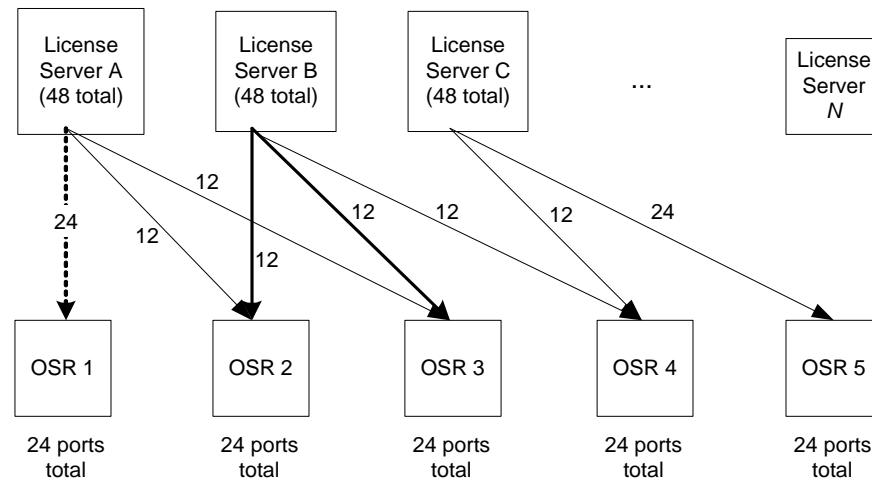
OR



Client-configured redundancy

You can add redundancy to load balanced systems by allowing all license servers on the network to allocate licenses to all license clients. (See “Balancing license server load” on [page 2-18](#) for more information about load balancing.) For client-configured redundancy, licenses are divided into multiple license pools, each administered by a single license server.

There is no practical limit to the number of license servers that you can designate for each license client machine.



For example, if License Server A fails, OSR no longer runs, but 2 and 3 continue with limited ports (12 and 12), since 2 and 3 are partially served by License Server B.

You can make this configuration more redundant by specifying the hostnames of *all* license servers for each OSR machine; if any license server fails, the others will have licenses available.

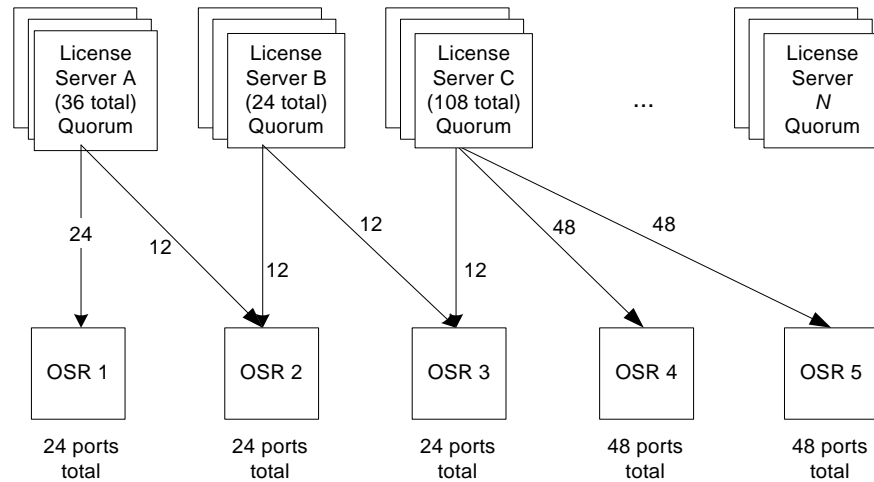
(Note that this example shows OSR but applies the same to Speechify.)

Combining server-configured and client-configured redundancy

You can combine both redundancy methods:

- ❑ Designate multiple primary license servers for each OSR or Speechify machine as you would do in “Client-configured redundancy” above. (See also, “Configuring OSR on Windows” on [page 3-37](#) and “Configuring Speechify on Windows” on [page 3-38](#).)
- ❑ Create a quorum of three servers for each primary server. This helps ensure that if a primary server crashes, you do not lose its available license pool.

For example:



When setting the license servers with `SWLicenseServerList`, you must keep all the quorum servers together in the list. For example, for OSR the list might appear as follows:

```
27000@tatooine;28000@nicosia;28000@arctic;28000@nepal; 27000@  
hoth;28000@dagobah;28000@endor;28000@naboo
```

For Speechify:

```
<param name="tts.server.licenseServerList">  
  <namedValue name="1"> 27000@tatooine </namedValue>  
  <namedValue name="2"> 28001@nicosia </namedValue>  
  <namedValue name="3"> 28001@arctic </namedValue>  
  <namedValue name="4"> 28001@nepal </namedValue>  
  <namedValue name="5"> 27000@hoth </namedValue>  
  <namedValue name="6"> 28001@dagobah </namedValue>  
  <namedValue name="7"> 28001@endor </namedValue>  
  <namedValue name="8"> 28001@naboo </namedValue>  
</param>
```

In both examples, bold indicates members of a quorum. Given this list, the SpeechWorks product would request licenses in the following order: tatooine, nicosia, hoth, and dagobah. If nicosia becomes unavailable, the SpeechWorks product would search arctic for licenses after tatooine. (Finally, if both nicosia and arctic are unavailable, then nepal **would not be used** because the quorum has been compromised.)

Sharing a license server with OSR and Speechify

The same license server instance may serve OSR and Speechify licenses. However, by default, installing OSR and Speechify installs separate license servers. This separate license server installation includes different software installation areas and different service names for the Globetrotter FLEXlm tools on Windows. To use a single license server for both OSR and Speechify, follow these steps:

1. For a network license server that serves only licenses (and doesn't run either product), install the license server from either one of the products (it doesn't matter which product).
2. Manually combine the OSR and Speechify licenses into a single file by opening both license files, then copy and paste the text of the INCREMENT lines from one file into the bottom of the other. For an example, see [page B-53](#).
3. Use the Globetrotter FLEXlm tools from one of the products (it doesn't matter which product) to configure the licensing server as normal – pointing it at the license file that now contains both the OSR and Speechify licenses.

To configure licensing, first configure OSR to set the licensing mode and the license server list (see “Configuring OSR on Windows” on [page 3-37](#) or “Configuring OSR on Linux” on [page 4-44](#)). Next, configure Speechify to set the licensing mode and license server list. (See “Configuring Speechify on Windows” on [page 3-38](#) or “Configuring Speechify on Linux” on [page 4-45](#).)

Platform integration design for licensing

While the previous information in this chapter focused on the physical architecture of license servers and SpeechWorks product machines, the following sections describe the software mechanisms that trigger the individual license allocations.

Overview of licensing modes

Both OSR and Speechify run in one of the following license allocation modes; in essence, the modes determine how long a client retains an issued license:

- *Default* licensing – For OSR, licenses are allocated and freed automatically as speech detector and recognizer resources are created. For Speechify, speak and voice licenses are allocated and freed automatically as each Speechify engine resource is created and destroyed by calling `SWIttsOpenPortEx()` and `SWIttsClosePort()` respectively. (The “default” mode is the default.)
- *Explicit* licensing – Licenses are allocated and freed by the platform developer using four API functions. To use explicit licensing for OSR, the platform must call the following functions:
 - `SWIepResourceAllocate()`
 - `SWIepResourceFree()`
 - `SWIrecResourceAllocate()`
 - `SWIrecResourceFree()`

Speechify uses these functions:

- `SWIttsResourceAllocate()`
- `SWIttsResourceFree()`

The platform integrator must decide which (or both) of the modes to support. If both modes are supported, the integration can allow application developers to decide which mode to use based on the needs of individual applications.

The mode is controlled by a configuration parameter (see below).

How to set licensing modes

OSR licensing modes

For OSR, the default licensing mode is specified in the `Baseline.xml` file. You can override the mode in a `user.xml` configuration file. *You cannot set these parameters dynamically via `SWIrecSetParameter()`.*

Specify the licensing mode separately for speech detection (SWIep) and recognition (SWIrec) with the `swiep_licensing_mode` and `swirec_licensing_mode` parameters.

Since the modes are specified separately for SWIep and SWIrec licenses, you can, for example, configure OSR to have SWIep licenses allocated with the default mode (when speech detectors are created) while SWIrec licenses are allocated with platform control (using explicit licensing).

Speechify licensing modes

For Speechify, the default licensing mode is specified in the `SWIttsConfig.xml` file, with the `tts.server.licensingMode` parameter. See “Configuring Speechify on Windows” on [page 3-38](#) or “Configuring Speechify on Linux” on [page 4-45](#).

Default OSR licensing

In this mode, OSR automatically allocates a speech license when either `SWIepDetectorCreate()` or `SWIrecRecognizerCreate()` is called. OSR frees licenses when `SWIepDetectorDestroy()` or `SWIrecRecognizerDestroy()` is called.

If OSR cannot check out a valid license from the license server under this mode, the create functions return a failure (`SWIep_ERROR_NO_LICENSE` or `SWIrec_ERROR_NO_LICENSE`, respectively).

See the *OSR Reference Manual* for details on the API functions and return codes.

Considerations for using default licensing mode

You may want to use the default licensing scheme in these cases:

- ❑ You plan to use all created speech detectors and recognizers for actual detection and recognition. If you create extra detector or recognizer objects that you are not using for detection or recognition, you should select a different licensing mode. Otherwise, you can purchase more licenses than your application actually needs.
- ❑ If you are detecting barge-in or using SpeechWorks' speech detector for begin-of-speech detection, the speech detector is always "active," and therefore, is probably active on every channel on your system. In this case, you need a license for every channel. The default licensing scheme for the speech detector is appropriate in this case.
- ❑ You have already completed an integration with the OSR 1.*n* software and do not want to perform additional integration work. This method is fully backwards compatible with OSR 1.*n* and allows your integration code to execute as before.
- ❑ You are deploying your license server on a separate machine in a wide-area network (WAN). Across such a network, latency is a factor and frequent round-trip communication to and from the license server is not advisable. Assuming that you do not create and destroy detectors and recognizers frequently, access to the license server is minimized.

Explicit OSR licensing

With this mode, the platform integrator decides when to allocate and free licenses for speech detectors and recognizers. Since every call to `SWIepDetectorCreate()` or `SWIrecRecognizerCreate()` creates a separate instance of a detector or recognizer, each instance requires a license when executing.

Four API functions are available; two each for the recognizer and speech detector (for function details and examples, see the *OSR Reference Manual*):

- ❑ `SWIepResourceAllocate()` and `SWIepResourceFree()`
- ❑ `SWIrecResourceAllocate()` and `SWIrecResourceFree()`

Use `SWIepResourceAllocate()` and `SWIrecResourceAllocate()` to allocate licenses for detector and recognizer objects.

Allocation can fail for these reasons:

- ❑ Your entire license pool for the chosen feature has been exhausted.
- ❑ The license server for your license client(s) is unavailable. As described earlier, license clients can point to multiple license servers, each of which may issue their own pre-allocated pool of licenses, so all license servers must in fact be unavailable or out of licenses for allocation to fail.

Considerations for using explicit licensing mode

You may want to employ the explicit method when:

- ❑ You want to have complete control over when a license is bound to your resources. This method can be used to guarantee that a license is always available and to avoid the situation described above in the default mode. For example, you may want to guarantee that for the duration of any call, a license is always available and cannot be freed and re-used by another call. To do this, at the beginning of your call (such as when the platform detects an incoming call), call `SWIrecResourceAllocate()` (or `SWIepResourceAllocate()`), and after hang-up has been detected, call the corresponding `Free()` functions. This is only one example of how these functions can be used.
- ❑ You have a pool of OSR licenses to use for multiple OSR servers or clients, but want to explicitly decide which licenses go to which machines (and only to those machines).
- ❑ You are running load tests that simulate speech densities that are atypical for normally deployed speech systems. In this case, processor and memory usage is extremely high, even if not indicative of a real-world deployment. Frequent round trips to the license server, although negligible in overall performance impact, may affect your performance numbers. You can use the explicit mode to ensure that the license server is accessed only once (regardless of your use of any other API function).
- ❑ You require dedicated ports for different text-to-speech licenses. For example, if you purchase more than one Speechify voice, you could use explicit mode to determine which ports are assigned to which voices.

Default Speechify licensing

In this mode, Speechify automatically allocates one speak license and one voice license when `SWIttsOpenPortEx()` is called. Speechify frees both licenses when `SWIttsClosePort()` is called.

If Speechify cannot check out a valid license from the license server under this mode, `SWIttsOpenPortEx()` returns a failure (`SWItts_NO_LICENSE`).

Considerations for using default licensing mode

You may want to use the default licensing scheme in these cases:

- ❑ You plan to use all created engine instances for actual speak operations. If you create extra engine objects that you are not using, you should select a different licensing mode. Otherwise, you can purchase more licenses than your application actually needs.
- ❑ You have already completed an integration with Speechify 2.x and do not want to perform additional integration work. This method is fully backwards compatible with Speechify 2.x and allows your integration code to execute as before.
- ❑ You are deploying your license server on a separate machine in a wide-area network (WAN). Across such a network, latency is a factor and frequent roundtrip communication to and from the license server is not advisable. Assuming that you do not create and destroy engines frequently, access to the license server is minimized.

Explicit Speechify licensing

With this mode, the platform integrator decides when to allocate and free licenses. Two API functions are available. For function details and examples, see the *Speechify User's Guide*.

- ❑ `SWIttsResourceAllocate()`
- ❑ `SWIttsResourceFree()`

Before calling `SWIttsSpeak()`, call `SWIttsResourceAllocate()` on a Speechify port resource to allocate one speak license and one voice license. Call `SWIttsResourceFree()` on that port to free both licenses. If you are using the explicit

licensing mode and do not call `SWIttsResourceAllocate()` before `SWIttsSpeak()`, the speak request fails (with `SWItts_NO_LICENSE`). If you call `SWIttsResourceFree()` while there is an active speak request (i.e., before `SWItts_cbStopped` or `SWItts_cbEnd` event delivery), the release fails (with `SWItts_MUST_BE_IDLE`).

Normally, your explicit allocation should span multiple `SWIttsSpeak()` operations. For example, you might allocate a license when a call comes in, use that port for the entire call, then free the license when the call ends.

Allocation can fail for these reasons:

- ❑ Your entire license pool has been exhausted.
- ❑ The license server for your license client is unavailable. As shown previously, license clients can point to multiple license servers, each of which may issue their own pre-allocated pool of licenses, so all license servers must in fact be unavailable or out of licenses for allocation to fail.

You do not need to allocate licenses for any `SWItts` API function except `SWIttsSpeak()`, not even for loading and activating dictionaries.

Considerations for using explicit licensing mode

You may want to employ the explicit method when:

- ❑ You want to have complete control over when a license is bound to your resources. This method can be used to guarantee that a license is always available and to avoid the situation described above in method I. For example, you may want to guarantee that for the duration of any call, a license is always available and cannot be freed and re-used by another call. To do this, at the beginning of your call (such as when the platform detects an incoming call), call `SWIttsResourceAllocate()`, and after hang-up has been detected, call `SWIttsResourceFree()`. This is only one example of how these functions can be used.
- ❑ You have a pool of Speechify licenses to use for multiple servers or clients, but want to explicitly decide which licenses go to which machines (and only to those machines).
- ❑ You are running load tests that simulate speak densities that are atypical for normally deployed speech systems. In this case, processor and memory usage is extremely high, even if not indicative of a real-world deployment. Frequent round trips to the license server, although negligible in overall performance impact, may affect your performance numbers.



Configuring Licensing on Windows

License server installation

Install from the SpeechWorks installation CD

The OSR and Speechify installation procedures include the needed files for licensing.

During installation, one step of the procedure lists software components and asks you to choose which ones to install on the current machine. For licensing, the needed option is “SpeechWorks Licensing Service.” Thus, during a product installation, you have these options for licensing:

- ❑ Install OSR or Speechify with a license server (default).
- ❑ Install OSR or Speechify without a license server (by de-selecting the “3rd Party Licensing Components” option).
- ❑ Install only a license server (de-select all options except the “3rd Party Licensing Components” option).



NOTE

Do not install the licensing component for both Speechify and OSR on the same system, as they will conflict. Instead, only install one (such as the OSR licensing component), then combine your OSR and Speechify licenses into a single license file and start the licensing service as normal. (See “Sharing a license server with OSR and

Speechify” on [page 2-25](#) for details.) Otherwise on reboot you will get a Windows error saying one or more services could not be started with the OSR or Speechify Licensing Service not starting.

If by accident you do this and get these errors, you can resolve this by disabling the Speechify licensing service, then configuring the OSR licensing service to use a combined license file. To disable the Speechify licensing service, open Control Panel >> Administrative Tools >> Services. Double-click Speechify Licensing Service, change the Startup Type from Automatic to Disabled, and click OK.

Configure licensing after the installation

After completing the software installation, follow these steps before running a speech application:

1. Obtain a valid license file from SpeechWorks and place it on your system. For details, see “Obtaining and managing licenses” on [page 1-4](#).
2. Configure the license server (this is optional if you store the license file in the recommended default location). See the next section below.
3. Start the license server (this can done with a reboot of the system if you instruct the operating system to do so). See the next section below.
4. For OSR, set up the SWILicenseServerList variable on each license client. See “Configuring OSR on Windows” on [page 3-37](#). For Speechify, set up the tts.server.licenseServerList parameter for each Speechify server. See “Configuring Speechify on Windows” on [page 3-38](#).

Once the steps above are complete, the product software is operational.

Configuring and starting the license server

Configuring, starting, and stopping is simple; there are very few steps, and most are optional.

Default configuration

When you install the license components for the first time, a Windows service called “OSR Licensing Service” or “Speechify Licensing Service” is created. By default, each service is configured to start automatically with every system reboot.

If you do not need to change the default, then the remaining tasks are:

1. Copy the license file you received from SpeechWorks to the default location (for OSR or Speechify, appropriately):

```
$SWISRSDK\flexlm\license folder\osr.lic  
$SWITTSSDK\flexlm\license folder\speechify.lic
```

2. Reboot the machine (or start the service manually via the Windows Services control panel applet).

Once the license server is running, it writes logs of its operations in these default locations:

```
$SWISRSDK\flexlm\license folder\osr-lic.log  
$SWITTSSDK\flexlm\license folder\speechify-lic.log
```

Changing the configuration

To change the default location of the license file, the license log file, or the auto-start configuration for the service, you can use the **lmtools** graphical tool or the **installs.exe** command line interface. For OSR, both tools are located in \$SWISRSDK\flexlm\components. For information on using the tools, see the *FLEXlm End Users Guide* in \$SWISRSDK\flexlm\components\htmlman\TOC.htm. (For Speechify, use \$SWITTSSDK instead of \$SWISRSDK.)

The following instructions perform these tasks using lmtools:

1. Copy the license file you received from SpeechWorks to the desired location.
2. Run the **lmtools** application.

3. Click the “Server/License File” tab, and select “Configuration Using Services” if it is not selected. (It should be selected by default.) If you do not do this, you cannot proceed using the steps below; the tabs at the top of the application will be absent or different.
4. Click the “Config Services” tab to edit the service called “OSR Licensing Service” or “Speechify Licensing Service” as appropriate. Verify the file paths of `osr.lic` or `speechify.lic` (the license file installed in step 1), `lmgrd.exe` (the license server executable), and the debug log file. Change the paths if desired.
5. Also in the “Config Services” tab, select the “Use Services” check box. SpeechWorks recommends that you also select the “Start Server at Power Up” check box so that the service starts automatically with every reboot of the machine.
6. Click “Save Service,” then click “Yes” to confirm.
7. To start the service, click the “Start/Stop/Reread” tab, and click “Start Server.”

By repeating the configuration steps above, you can create a new log file or change the default location of the license file.

Manually starting and stopping the license server

Although SpeechWorks recommends that you configure the license server to automatically start upon reboot of the machine, you can start it manually with the **lmgrd** command and the name of your license file:¹

```
lmgrd -c osr.lic  
lmgrd -c speechify.lic
```

Similarly, SpeechWorks recommends that the license server always remain running, but you can stop it manually with the appropriate **lmutil** command:

```
lmutil lmdown -c osr.lic  
lmutil lmdown -c speechify.lic
```

Above, when you enter the `lmgrd` and `lmutil` commands, you may need to add path information to the commands and/or the license filenames depending on the location of your working directory.

1. Run `lmgrd` and `lmutil` from the `$SWISRSdk\flexlm\components` directory (for OSR) or `$SWITTSSdk\flexlm\components` (for Speechify).

For information on `lmgrd` and `lmutil` options, see the *FLEXlm End Users Guide*.

Configuring OSR on Windows

After a license server is running with a valid OSR license file on a host in the subnet, you must ensure that the `SWILicenseServerList` variable is set on every license client machine. (As explained below, you might not need to change the default setting that is established during OSR installation.) Then, once you start the license manager service, your OSR application is able to retrieve valid licenses and run properly.

Setting up `SWILicenseServerList`

On Windows, `SWILicenseServerList` is defined in the `SpeechWorks.cfg` configuration file.

Set `SWILicenseServerList` to the port number and hostname of the license server machine using the following format:

```
port_number@hostname
```

This example is for a license server running on the current machine:

```
27000@osr-hostname
```

By default, the OSR installation sets `SWILicenseServerList` to specify the machine on which OSR is installed and the default FLEXlm port of 27000. In other words, `SWILicenseServerList` is set to `27000@localhost`. You must change this setting if there is no license server running on the installation machine, or if your primary license server is on another machine (recommended).

If you install OSR on host `groucho`, but want to run your license server on host `harpo`, you need to change the value of `SWILicenseServerList` on `groucho` (the OSR machine) to `27000@harpo`.

If you install OSR Client-Server software, you must change the default `SWILicenseServerList` setting on every OSR client and OSR server machine unless that machine is running the designated license server.

You can also configure each OSR machine to look for licenses from several different FLEXlm license servers that you have started. Set `SWILicenseServerList` to specify a list (separated by semi-colons) of port@server combinations. For example, if you have started license servers on both harpo and groucho, set the `SWILicenseServerList` parameter on your OSR host to `27000@harpo;27000@groucho`.

Configuring Speechify on Windows

After a license server is running with a valid Speechify license file on a host in the subnet, you must:

- ❑ Set the licensing mode.
- ❑ Set the `tts.server.licenseServerList` parameter on every Speechify server machine.

Then when you start the license manager service, your Speechify server is able to retrieve valid licenses and run properly.

Configuring Speechify manually

To set the licensing mode manually, set the value of the `tts.server.licensingMode` parameter in `$SWITTSSDK\config\SWIttsConfig.xml`:

```
<param name="tts.server.licensingMode">  
  <value> explicit </value>  
</param>
```

Also in `SWIttsConfig.xml`, set `tts.server.licenseServerList` to the port number and hostname of the license server machine using the following format:

```
<param name="tts.server.licenseServerList">  
  <namedValue name=""> port_number@hostname </namedValue>  
</param>
```

The `<namedValue>` "name" attribute values are ignored. By default, the Speechify installation sets `tts.server.licenseServerList` to `27000@localhost`. Thus, if you are running the license server on the same machine where Speechify is running, you do not need to change the default.

If you install Speechify on host groucho, but want to run your license server on host harpo, you need to change the value of `tts.server.licenseServerList` on groucho (the Speechify machine) to `27000@harpo`.

You can also configure Speechify to look for licenses from several different FLEXlm license servers that you have started. Set `tts.server.licenseServerList` to specify a list of `port@server` combinations. For example, if you have started license servers on both harpo and groucho, set the `tts.server.licenseServerList` parameter to:

```
<param name="tts.server.licenseServerList">  
  <namedValue name="harpo"> 27000@harpo </namedValue>  
  <namedValue name="groucho"> 27000@groucho </namedValue>  
</param>
```

Speechify looks for valid Speechify licenses on each of the servers in the order they are listed, going to the next server in the list only when valid licenses cannot be obtained from the previous server (the server is down, is out of licenses, or produces some other error).



Configuring Licensing on Linux

Starting the FLEXlm license server

Once you have installed OpenSpeech Recognizer or Speechify and obtained a valid license file from SpeechWorks International, you need to start the FLEXlm license server before you can run a speech application. The license server can be run either on the same host that is running OSR or Speechify, or on another host on the same subnet (recommended). You can start the license server manually with the `lmgrd` command and you can stop it with `lmutil` (see below for details), or configure the system to automatically start and restart the license server.

Manually starting and stopping the license server

To start the license server manually, use the **lmgrd** command with the name of your license file (for OSR or Speechify, appropriately):¹

```
lmgrd -c osr.lic  
lmgrd -c speechify.lic
```

1. Run `lmgrd` and `lmutil` from the `$SWISRSdk/flexlm/components` directory for OSR, and from the `$SWITTSSdk/flexlm/components` directory for Speechify.

To stop the license server manually via the command line, use the **lmutil** command with the name of your license file:

```
lmutil lmdown -c osr.lic  
lmutil lmdown -c speechify.lic
```

Above, when you enter the **lmgrd** and **lmutil** commands, you may need to add path information to the commands and/or the license filenames depending on the location of your working directory.

For information on **lmgrd** and **lmutil** options, see the *FLEXlm End Users Guide*.

Automatically starting the license server

To start the service automatically after rebooting your machine, add the appropriate command to the end of the file `/etc/rc.d/rc.local`:

```
lmgrd -c osr.lic  
lmgrd -c speechify.lic
```

Creating log files

By default, the license server does not write operational log files. To create logs, you must add the **-l** option to the **lmgrd** command line. For example:

```
lmgrd -c /your/path/osr.lic -l /your/log/path/osr-lic.log
```

Use the plus symbol (+) to append to an existing log. For example, use the symbol in the `/etc/rc.d/rc.local` file when automatically starting the licensing after reboots:

```
lmgrd -c /your/path/osr.lic -l +/your/log/path/osr-lic.log
```

Running the license server on a separate host

You can also run the license server software on a different host than the one running OSR or Speechify. To use a separate host as a license server, first install the license server software on the OSR or Speechify host via the SpeechWorks installation program. This configuration installs the server in the flexlm directory under the product root directory.

- ❑ For OSR, copy \$SWISRSDK/flexlm and its contents from the OSR host to the desired separate host machine. For Speechify, copy \$SWITTSSDK/flexlm and its contents from the Speechify host.
- ❑ For OSR, change the SWILicenseServerList parameter in the SpeechWorks.cfg file to the name of the new machine or machines. (See “Configuring OSR on Linux” below.) For Speechify, set up the tts.server.licenseServerList parameter for each server. (See “Configuring Speechify on Linux” below.)

Please remember that a license you receive from SpeechWorks is tied to a specific license server. You cannot run a license server with a license file that was created for another server.

Configuring OSR on Linux

After a license server is running with a valid OSR license file on a host in the subnet, you must ensure that the `SWILicenseServerList` variable is set on every license client machine. (As explained below, you might not need to change the default setting that is established during OSR installation.) Then, once you start the license manager service, your OSR application is able to retrieve valid licenses and run properly.

Setting up `SWILicenseServerList`

`SWILicenseServerList` is a parameter in the `SSWISRSdk/config/SpeechWorks.cfg` configuration file. By default, the value is loaded from the file unless you define the parameter as an environment variable before starting the application.

Set `SWILicenseServerList` to the port number and hostname of the license server machine using the following format:

```
port_number@hostname
```

This example is for a license server running on the current machine:

```
27000@localhost
```

By default, the OSR installation sets `SWILicenseServerList` to specify the machine on which OSR is installed and the default FLEXlm port of 27000. In other words, `SWILicenseServerList` is set to `27000@localhost`. You must change this setting if there is no license server running on the installation machine, or if your primary license server is on another machine (recommended).

Once you start the license manager service, your OSR application is able to retrieve valid licenses and run properly.

If you install OSR on host `groucho`, but want to run your license server on host `harpo`, you need to change the value of `SWILicenseServerList` on `groucho` (the OSR machine) to `27000@harpo`.

You can also have OSR look for licenses from several different FLEXlm license servers that you have started. Set `SWILicenseServerList` to specify a list (separated by colons) of `port@server` combinations. For example, if you have started license servers on both `harpo` and `groucho`, set the `SWILicenseServerList` parameter on your OSR host to `27000@harpo:27000@groucho`.

Configuring Speechify on Linux

After a license server is running with a valid Speechify license file on a host in the subnet, you must:

- ❑ set the licensing mode
- ❑ set the `tts.server.licenseServerList` parameter on every license client machine

Then when you start the license manager service, your Speechify server is able to retrieve valid licenses and run properly.

To set the licensing mode, set the value of the `tts.server.licensingMode` parameter in `$SWITTSSDK/config/SWIttsConfig.xml`:

```
<param name="tts.server.licensingMode">  
  <value> explicit </value>  
</param>
```

Also in `SWIttsConfig.xml`, set `tts.server.licenseServerList` to the port number and hostname of the license server machine using the following format:

```
<param name="tts.server.licenseServerList">  
  <namedValue name=""> port_number@hostname </namedValue>  
</param>
```

The `<namedValue>` “name” attribute values are ignored. By default, the Speechify installation sets `tts.server.licenseServerList` to `localhost@27000`. Thus, if you are running the license server on the same machine where Speechify is running, you do not need to change the default.

If you install Speechify on host `groucho`, but want to run your license server on host `harpo`, you need to change the value of `tts.server.licenseServerList` on `groucho` (the Speechify machine) to `27000@harpo`.

You can also configure Speechify to look for licenses from several different FLEXlm license servers that you have started. Set `tts.server.licenseServerList` to specify a list of `port@server` combinations. For example, if you have started license servers on both `harpo` and `groucho`, set the `tts.server.licenseServerList` parameter to:

```
<param name="tts.server.licenseServerList">  
  <namedValue name="harpo"> 27000@harpo </namedValue>  
  <namedValue name="groucho"> 27000@groucho </namedValue>  
</param>
```

Speechify looks for valid Speechify licenses on each of the servers in the order they are listed, going to the next server in the list only when valid licenses cannot be obtained from the previous server (the server is down, is out of licenses, or produces some other error).



License Server Host IDs

FLEXlm uses different machine identifications for different machine architectures. The program `lmhostid` prints the exact hostid that FLEXlm expects to use on that machine. See the *FLEXlm End Users Guide* for more details about hostids.

This table lists alternate methods to obtain the required hostid for each machine architecture. If you have not installed FLEXlm software, you can find the hostid with:

```
ipconfig /all
```

Typically, FLEXlm and its SpeechWorks licensing components run on the same machine. The available platforms are listed in the first portion of the following table. However, FLEXlm also runs on hardware platforms where the SpeechWorks components are not supported (shown in the second portion of the table), and it is possible to configure FLEXlm to run on one machine while the SpeechWorks components run on another.

Hardware platform	Hostid	Command to run on license server	Example
SpeechWorks-supported and FLEXlm-supported:			
Linux (OSR and Speechify)	ethernet address	<code>/sbin/ifconfig eth0</code> Remove colons from HWaddr, e.g., 00:40:05:16:E5:25.	00400516E525
SUN (OSR client only)	32-bit hostid	<code>hostid</code>	170a3472
Windows (OSR and Speechify)	ethernet address	<code>lmutil lmhostid</code> Run <code>lmutil</code> from the SSWISRSdk\flexlm\components directory. Remove hyphens from the "Physical Address."	00B0A9DF9A32

Hardware platform	Hostid	Command to run on license server	Example
FLEXlm-supported only:			
AIX (RS/6000, PPC)	32-bit hostid	<code>uname -m</code> This command returns a 12-digit string, e.g., 000276513100. Remove the first two and the last two digits.	02765131
DEC Alpha	ethernet address	<code>netstat -i</code>	080020005532
HP	32-bit hostid	<code>uname -i</code> Convert to hex, or prepend with #.	778DA450 or #2005771344
	ethernet address	<code>lanscan</code> Station address without leading 0x.	0000F0050185
SCO	hostid string	<code>uname -x</code> Serial is SCO00354; prefix with "ID_STRING=".	ID_STRING=SCO00354
SGI	32-bit hostid	<code>/etc/sysinfo -s</code> Convert to hex, or prepend with #.	69064C3C or #1762020412



License Files

This appendix describes the details of the license files. As an end-user, you are allowed to make the following changes to the files (any other changes invalidate the license):

- ❑ Change the machine name (*not* the hostid) on the SERVER line of the license file. For example, this is useful if you rename the machine or if you run lmgrd on a different machine from the SpeechWorks components.
- ❑ Add a path to the vendor daemon on the VENDOR line of the license file. This is useful when swilmgrd is not stored in the default location and lmgrd cannot find that different directory).
- ❑ Change the port being used in the license file. This is useful when you are already using the default port for something else (for example, for another lmgrd process for a different product).
- ❑ Add an INCREMENT line if directed by SpeechWorks technical support or when merging OSR and Speechify license files ([page 2-25](#)). This line enables new features (for example, if you purchase a new Speechify voice). In these instances, you edit the license file and paste the provided text underneath the last FEATURE or INCREMENT line in your current file.
- ❑ Add new licenses to an existing license server (for example, to add new ports for a new application). For details, see “Making changes to generated licenses” on [page 1-5](#).

The format of the license file depends upon the type of license that is generated, based upon the chosen licensing policy; if you edit this file (e.g., to use the same license server for both OSR and Speechify), *be very careful*.

Filenames and status after initial installation

The default names of these files are:

Product	License file name
OSR	osr.lic
Speechify	speechify.lic

When you install OSR, the installation process specifies the hostname and port of the license server where the license is installed.

Parameter	Description	Default
SWILicenseServerList	Port number and hostname of the license server where the license file is installed.	27000@localhost

Speechify configures license servers in the configuration file, specifying the hostname and port of the license server where the license is installed. The default configuration value does not need to be changed unless you're using a remote license server.

Configuration parameter	Default
tts.server.licenseServerList	27000@ <i>speechify-hostname</i>

Assuming that the license server is up and running with a valid license, and the license server name and port are specified, nothing else needs to be done. See “Configuring OSR on Windows” on [page 3-37](#) or “Configuring Speechify on Windows” on [page 3-38](#) for more information about using a different license server or using a pool of license servers.

For more information about the license server, see “Configuring and starting the license server” on [page 3-35](#).

Sample license file for OSR

Below is a sample license file for OSR:

```
SERVER bellagio 0010a4c4d149 27000
VENDOR swilmgrd
USE_SERVER
INCREMENT osr_swirec swilmgrd 1.1 permanent 4 SIGN="0995
D853 6B3F CFB2 CA4C 4B7A 3506 41C2 DEDB A9A5 1A2E 9AFD
46D9 E60F 73DD 0293 A78E 8201 EDC6 5057 88F8 A0EA 4722
EF9D 446A 8552 8C23 8B02 B00F 8362"
INCREMENT osr_swiep swilmgrd 1.1 permanent 4 SIGN="0F89 429B
7967 0196 84A7 9726 BD54 E02E FE21 09E3 4D26 A947 648B
3DDB A9FA 1B62 E6E1 D6A5 8894 7668 F3B7 4573 01C1 29C6
B2B9 7AC1 6E44 26C8 88C6 E6CD"
```

The components of the license file include these lines:

- ❑ The SERVER line identifies the name of the server and its hostid. The final number on the line (27000) is the license server port (the port you specify when configuring OSR).
- ❑ The VENDOR line names the SpeechWorks vendor daemon. It can point to a different directory if swilmgrd and lmgrd are not in the same directory.
- ❑ USE_SERVER is explained in the *FLEXlm End Users Guide*.
- ❑ The INCREMENT line includes the name of the feature (e.g., osr_swirec or osr_swiep), the vendor daemon to use (swilmgrd), the version of the license (1.1), the expiration of the license (permanent), the number of licenses (4), and the signature that verifies the license is valid. You can add Speechify INCREMENT lines to this file below the OSR lines if you want to merge the two products under one license server (see “Sharing a license server with OSR and Speechify” on [page 2-25](#)).

Sample license file for Speechify

Below is a sample license file for Speechify:

```
SERVER venecia 0010a4c4d149 27000
VENDOR swilmgrd
USE_SERVER
INCREMENT speechify_switts swilmgrd 1.1 permanent 6 \
  SIGN2="0A34 731E \
    BBFC 05D6 2606 A624 4545 EDB8 46C7 A2B1 755A A338 \
    D9A7 5343 7CED 1FDF 7B9C D8A9 857A 7EAF 588D 2D33 9C29 \
    2F73 57BD 3D95 88A8 18B4 1DBA 642E"
INCREMENT speechify_switts_en_US_Tom swilmgrd 1.1 \
  permanent 6 SIGN2="0B41 76E5 3A03 3B7C 5979 DB49 577F \
    611B 1F09 C6D3 E775 B754 0F79 1305 3602 0F57 7CD9 DB76 \
    3295 EB02 C9AB B809 46D3 4596 0E65 4306 639A 33C0 \
    0764 260A"
```

The components of the license file include these lines:

- ❑ The SERVER line identifies the name of the server and its hostid. The final number on the line (27000) is the license server port (the port you specify when configuring Speechify).
- ❑ The VENDOR line names the SpeechWorks vendor daemon.
- ❑ USE_SERVER is explained in the *FLEXlm End Users Guide*.
- ❑ The INCREMENT line includes the name of the feature (e.g., speechify_switts or speechify_switts_en_US_tom), the vendor daemon to use (swilmgrd), the version of the license (1.1), the expiration of the license (permanent), the number of licenses (6), and the signature that verifies the license is valid. You can add Speechify INCREMENT lines to this file below the OSR lines if you want to merge the two products under one license server (see “Sharing a license server with OSR and Speechify” on [page 2-25](#)).

Sample shared (merged) license file for OSR and Speechify

Below is a sample license file that merges the sample OSR and Speechify files described on the previous pages. This allows the same server instance to provide licenses for both products.

In this example, the INCREMENT lines from the Speechify file were appended to the OSR file. Alternatively, the lines from OSR could have been added to the Speechify sample.

```
SERVER bellagio 0010a4c4d149 27000
VENDOR swilmgrd
USE_SERVER
INCREMENT osr_swirec swilmgrd 1.1 permanent 4 SIGN="0995
D853 6B3F CFB2 CA4C 4B7A 3506 41C2 DEDB A9A5 1A2E 9AFD
46D9 E60F 73DD 0293 A78E 8201 EDC6 5057 88F8 A0EA 4722
EF9D 446A 8552 8C23 8B02 B00F 8362"
INCREMENT osr_swiep swilmgrd 1.1 permanent 4 SIGN="0F89 429B
7967 0196 84A7 9726 BD54 E02E FE21 09E3 4D26 A947 648B
3DDB A9FA 1B62 E6E1 D6A5 8894 7668 F3B7 4573 01C1 29C6
B2B9 7AC1 6E44 26C8 88C6 E6CD"
INCREMENT speechify_switts swilmgrd 1.1 permanent 6 \
SIGN2="0A34 731E \
BBFC 05D6 2606 A624 4545 EDB8 46C7 A2B1 755A A338 \
D9A7 5343 7CED 1FDF 7B9C D8A9 857A 7EAF 588D 2D33 9C29 \
2F73 57BD 3D95 88A8 18B4 1DBA 642E"
INCREMENT speechify_switts_en_US_Tom swilmgrd 1.1 \
permanent 6 SIGN2="0B41 76E5 3A03 3B7C 5979 DB49 577F \
611B 1F09 C6D3 E775 B754 0F79 1305 3602 0F57 7CD9 DB76 \
3295 EB02 C9AB B809 46D3 4596 0E65 4306 639A 33C0 \
0764 260A"
```


Index

A

all-in-one configuration [2-12](#)

C

client-configured redundancy [2-23](#)

client-server configuration [2-13](#)

configuring OSR [3-37](#), [4-44](#)

D

default licensing [2-26](#), [2-27](#), [2-30](#)

design strategies for integration [2-26](#)

downloading license files [1-4](#)

E

evaluation license [1-2](#)

explicit allocation [2-28](#)

explicit licensing [2-26](#)

F

failure recovery

 recovery from server failures [2-10](#)

FLEXlm license server [3-38](#), [3-39](#), [4-41](#), [4-44](#), [4-45](#)

G

generating licenses [1-4](#)

H

hostid [A-47](#)

I

installing licensing [3-33](#)

installs.exe [3-35](#)

integration design strategies [2-26](#)

in-the-skins configuration [2-12](#)

L

license file [B-51](#), [B-52](#), [B-53](#)

license server distribution [2-9](#)

licensing

 default [2-26](#)

 explicit [2-26](#)

lmgrd [3-36](#), [4-41](#)

lmtools [3-35](#)

lmutil [3-36](#), [4-42](#), [A-47](#)

load balancing [2-18](#)

log file, for the license server [3-36](#)

log file, location of the license server log [3-35](#)

log file, location of the license server log (Linux)
[4-42](#)

M

managing licenses [1-4](#)

O

osr.lic [3-35](#), [3-36](#), [4-41](#), [4-42](#), [B-50](#)

osr-lic.log [3-35](#), [4-42](#)

P

platform integration strategies [2-26](#)

Q

quorum [2-19](#)

R

rc.local [4-42](#)

redundant license servers [2-18](#)

replacing licenses [1-6](#)

returning licenses [1-6](#)

runtime license [1-2](#)

S

SDK license [1-2](#)

server-configured redundancy [2-19](#), [2-23](#)

speak licenses [1-3](#)

speechify.lic [3-35](#), [3-36](#), [4-41](#), [4-42](#), [B-50](#)

speechify-lic.log [3-35](#), [4-42](#)

support services [viii](#)

swiep_licensing_mode [2-27](#)

SWIepDetectorCreate() [2-28](#)

SWIepResourceAllocate() [2-28](#)

SWIepResourceFree() [2-28](#)

SWILicenseServerList [3-37](#), [4-43](#), [4-44](#), [B-50](#)

swirec_licensing_mode [2-27](#)

SWIrecRecognizerCreate() [2-28](#)

SWIrecResourceAllocate() [2-28](#)

SWIrecResourceFree() [2-28](#)

SWIttsConfig.xml [2-27](#), [3-38](#), [4-45](#)

SWIttsOpenPortEx() [2-30](#)

SWIttsResourceAllocate() [2-30](#)

SWIttsResourceFree() [2-30](#)

SWIttsSpeak() [2-30](#)

T

tts.server.licenseServerList [3-38](#), [4-45](#), [B-50](#)

tts.server.licensingMode [2-27](#), [3-38](#), [4-45](#)

U

upgrading licenses [1-6](#)

V

viewing licenses [1-4](#)

voice licenses [1-3](#)