

SharePlex for Oracle 9.1.2

Installation and Setup Guide for a SQL Server Source



© 2018 Quest Software Inc. ALL RIGHTS RESERVED.

This guide contains proprietary information protected by copyright. The software described in this guide is furnished under a software license or nondisclosure agreement. This software may be used or copied only in accordance with the terms of the applicable agreement. No part of this guide may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording for any purpose other than the purchaser's personal use without the written permission of Quest Software Inc.

The information in this document is provided in connection with Quest Software products. No license, express or implied, by estoppel or otherwise, to any intellectual property right is granted by this document or in connection with the sale of Quest Software products. EXCEPT AS SET FORTH IN THE TERMS AND CONDITIONS AS SPECIFIED IN THE LICENSE AGREEMENT FOR THIS PRODUCT, QUEST SOFTWARE ASSUMES NO LIABILITY WHATSOEVER AND DISCLAIMS ANY EXPRESS, IMPLIED OR STATUTORY WARRANTY RELATING TO ITS PRODUCTS INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NON-INFRINGEMENT. IN NO EVENT SHALL QUEST SOFTWARE BE LIABLE FOR ANY DIRECT, INDIRECT, CONSEQUENTIAL, PUNITIVE, SPECIAL OR INCIDENTAL DAMAGES (INCLUDING, WITHOUT LIMITATION, DAMAGES FOR LOSS OF PROFITS, BUSINESS INTERRUPTION OR LOSS OF INFORMATION) ARISING OUT OF THE USE OR INABILITY TO USE THIS DOCUMENT, EVEN IF QUEST SOFTWARE HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. Quest Software makes no representations or warranties with respect to the accuracy or completeness of the contents of this document and reserves the right to make changes to specifications and product descriptions at any time without notice. Quest Software does not make any commitment to update the information contained in this document.

If you have any questions regarding your potential use of this material, contact:

Quest Software Inc.

Attn: LEGAL Dept

4 Polaris Way

Aliso Viejo, CA 92656

Refer to our Web site (<https://www.quest.com>) for regional and international office information.

Patents

Quest Software is proud of our advanced technology. Patents and pending patents may apply to this product. For the most current information about applicable patents for this product, please visit our website at <https://www.quest.com/legal>.

Trademarks

Quest, the Quest logo, SharePlex, and Join the Innovation are trademarks and registered trademarks of Quest Software Inc. For a complete list of Quest marks, visit <https://www.quest.com/legal/trademark-information.aspx>. All other trademarks and registered trademarks are property of their respective owners.

SharePlex Installation and Setup Guide for a SQL Server Source

Updated - 4/27/2018

Version - 9.1.2

Contents

About this Guide	8
Conventions used in this guide	9
Basic system requirements	10
System	10
Internet protocol	10
System Requirements — SQL Server Capture	12
Supported versions and targets	12
On-premises SQL Server source to on-premises SQL Server target	12
On-premises SQL Server source to other on-premises targets	13
On-premises SQL Server source to all cloud target types	14
Cloud SQL Server source to all cloud target types	14
Conditions of support	15
Supported operations	15
Supported data types — SQL Server to SQL Server	15
Supported data types — SQL Server to Oracle and Open Target	16
SQL Server to MySQL supported data types	16
SQL Server to Oracle supported data types	18
SQL Server to PostgreSQL supported data types	19
SQL Server to XML supported data types	20
Supported SharePlex features	21
SharePlex preinstallation checklist	23
Overview	23
Network checklist	23
Installer checklist	25
Windows system checklist	26
SQL Server source checklist	27
Open Target checklist	30
Download the SharePlex installer	34
About the SharePlex installers	34
Linux and Unix	34
Windows	35
Where to get SharePlex installers	35
Install SharePlex on Linux and UNIX	36

Install SharePlex on Linux/Unix for Oracle Database	36
Run the installer in interactive mode	37
Next steps	38
Install SharePlex on Linux/Unix for Open Target Databases	39
Run the installer in interactive mode	39
Next steps	40
Install SharePlex on Windows	42
Install SharePlex on Windows	42
Change global settings in the MKS toolkit	44
Assign SharePlex users to security groups	46
Overview	46
About the SharePlex security groups	46
Description of the SharePlex security groups	46
Create and populate SharePlex groups on Unix and Linux	47
Create and populate SharePlex groups on Windows	48
Set up replication between SQL Server databases	49
About these instructions	49
Configure SharePlex on the source	49
Run Database Setup	50
Ensure that all tables have a primary key	50
Ensure varchar(max) length compatibility	50
Configure replication	50
Configure SharePlex on the target	51
Set up replication from SQL Server to a different target type	52
About these instructions	52
Set up replication from SQL Server to Oracle	53
Review column names	53
Ensure column length compatibility	53
Configure SharePlex on the source	54
Run Database Setup	54
Ensure that all tables have a primary key	54
Configure replication	54
Configure SharePlex on the target	55
Set up replication from SQL Server to MySQL	55
Review column names	56
Configure SharePlex on the source	56
Run Database Setup	56
Ensure that all tables have a primary key	56
Configure replication	56
Configure SharePlex on the target	57

Set up replication from SQL Server to PostgreSQL	57
Review column names	57
Configure SharePlex on the source	57
Run Database Setup	58
Ensure that all tables have a primary key	58
Configure replication	58
Configure SharePlex on the target	59
Set up replication from SQL Server to Kafka	59
Overview	59
Guidelines for posting to Kafka	59
Configure SharePlex on the source	59
Run Database Setup	59
Ensure that all tables have a primary key	59
Configure replication	60
Configure SharePlex on the target	60
Set the Kafka record format	61
View and change Kafka settings	61
Set recovery options	63
Set up replication to a cloud-hosted Open Target database	64
Installation in an IaaS (accessible) environment	64
Installation in a PaaS (non-accessible) environment	64
Basic SharePlex demonstration - all platforms	67
Overview	67
What you will learn	68
Prework for the demonstrations	68
Tables used in the demonstrations	68
INSERT scripts	68
Create and activate a configuration	68
Create a configuration file	69
Configuration templates	69
Activate the configuration	70
Troubleshooting Tips	70
Demonstration of replication	70
Verify replication startup	70
Verify replication of large data volumes	71
Verify queuing and continuity of replication	71
Verify SharePlex capture recovery	72
Demonstration of data compare and repair	72
Demonstration of named post queues	73
Clean up the replication environment	73
Create a configuration file	74
Configuration templates	74

Activate the configuration	75
Generate data	75
View the post queues	75
Solve installation problems	76
Overview	76
Solve license utility problems	76
Solve installer problems	77
Solve Database Setup problems	78
Oracle setup issues	78
SQL Server setup issues	80
Solve database connection problems	80
Solve SharePlex startup problems	81
Solve sp_ctrl problems	81
Solve host connection problems	82
How to find the ORACLE_SID and ORACLE_HOME	82
Remove SharePlex from a system	84
Overview	84
Remove SharePlex from Unix/Linux	84
Remove SharePlex from Windows	85
Files removed by the uninstall program	85
Files not removed by the uninstall program	85
Remove SharePlex from the system	85
Database Setup Utilities	88
Database Setup for MySQL	88
Overview	88
Supported databases	88
Guidelines for use	89
Required privileges	89
Run Database Setup for MySQL	89
Database Setup for Oracle	90
Overview	90
Supported databases	91
When to run Oracle Setup	91
Supported Oracle Connections	91
Required run privileges	92
Privileges granted to SharePlex	92
Requirements	93
SharePlex schema storage requirements	93
Run Database Setup for Oracle	94

Database Setup for PostgreSQL	98
Overview	98
Supported databases	98
Guidelines for use	99
Required privileges	99
Run Database Setup for PostgreSQL	100
Database Setup for SQL Server	101
Overview	101
Supported databases	101
Guidelines for use	102
Required privileges	102
Run Database Setup for SQL Server	102
Database Setup for Tiberio	104
Overview	104
Supported databases	104
Guidelines for use	104
Required privileges	104
Run Database Setup for Tiberio	104
General SharePlex Utilities	106
SharePlex license utilities	106
Description	106
License utilities on Unix and Linux	106
License utility on Windows	108
Install the SharePlex service	109
Advanced SharePlex installer options	110
Install SharePlex as Root	112
Run the installer in unattended mode	113
Enter responses in the file	113
Next Steps	114
SharePlex installed items	115
SharePlex Directories	115
SharePlex Installed Objects	116
MKS Toolkit	118
SharePlex Registry Entries	118
About us	119
Contacting Quest	119
Technical support resources	119

About this Guide

This manual provides instructions for installing and setting up SharePlex on an Oracle a SQL Server source database and all supported target types.

This manual also contains some demonstration tutorials that show how certain features of SharePlex work.

Other SharePlex documentation

For the complete SharePlex documentation set, go to <https://support.quest.com/shareplex>.

Conventions used in this guide

Conventions used in this manual

The following typographic conventions are used in this guide.

- **Bold** represents required components of a command or option that must be typed as shown.
- *Italics* represent variables defined, named or entered by the user.
- {Braces} enclose required arguments.
- [Brackets] represent optional command components and may also be used in example command strings to emphasize required user defined variables in long strings.

Example:

reconcile queue {*queue*name} for {*datasource-datadest*} [**on** *host*]

- A vertical bar, or “pipe” character, (|) within brackets or braces indicates that you can use only one of the enclosed components.

Example:

abort service {*service* | **all**}

Names of commands, programs, directories and files are expressed in **Bold**.

Other names are expressed in capital letters using the default font.

Examples:

The **sp_ctrl** program is located in the **bin** directory.

Open the **oramsglst** file.

Find the value for ORACLE_HOME.

Click **Apply**.

System displays, such as prompts and command output, are expressed in a `monofaced` (fixed-space) font.

Examples:

```
sp_ctrl(sysA)>
```

```
User is a viewer (level=3)
```

Windows menu items, dialog boxes, and options within dialog boxes are expressed in **Bold**.

Example:

From the **File** menu, select **Print**.

System names are expressed generically or fictitiously. When necessary, the source system (or primary system) is referred to as *SysA*. Target systems (or secondary systems) are referred to as *SysB*, *SysC*, *SysD*, and so forth.

Basic system requirements

This chapter contains the basic system requirements for SharePlex. See also:

[System Requirements — SQL Server Capture](#) on page 12

System

Before installing SharePlex, ensure that your system meets the minimum hardware and software requirements.

- SharePlex processes are all 64-bit and can exceed 4 GB.
- Per process memory of greater than or equal to 256 MB is required. Depending on how you configure SharePlex, there can be one or more of the following processes on a system:
 - Capture
 - Read
 - Export
 - Import
 - Post
- See the Preinstallation Checklist in the SharePlex Installation Guide for additional system and/or database requirements.

Internet protocol

SharePlex supports IPv4 and IPv6 internet protocols. The following table shows the platforms for which SharePlex was tested with IPv6.

NOTE: If the `SP_SYS_HOST_NAME` environment variable is set to an IPV6 address on the source system, SharePlex on the target system must be version 9.0 or later.

Platform	Source	Target	SharePlex Source and Target versions	On-premises/Cloud
Windows	Link-Local IPV6	Link-Local IPV6	9.0 - 9.0	On-premises
Windows	Link-Local IPV6	Link-Local IPV6	9.0 - 8.6.4	On-premises
Windows	Public IPV6	Public IPV6	9.0 - 9.0	Cloud

System Requirements — SQL Server Capture

This chapter contains the requirements for capture from a SQL Server source database and replication to supported targets.

SharePlex support for SQL Server as a target is documented in [Set up replication from Oracle to SQL Server](#) on page 1.

Supported versions and targets

This section is divided into the following sections according to the way that a database is implemented (on-premises or cloud):

[On-premises SQL Server source to on-premises SQL Server target](#)

[On-premises SQL Server source to other on-premises targets](#)

[On-premises SQL Server source to all cloud target types](#)

[Cloud SQL Server source to all cloud target types](#)

NOTE: SharePlex supports both SQL Server Standard Edition and SQL Server Enterprise Edition.

On-premises SQL Server source to on-premises SQL Server target

This section shows the SQL Server versions that SharePlex supports when the **source and target databases are both SQL Server** and both are installed **on-premises**. SharePlex can replicate data between any of these database/platform combinations.

SQL Server Version (source or target)	Supported Platforms (source or target)
SQL Server 2008	Windows Server 2012

SQL Server Version (source or target)	Supported Platforms (source or target)
	Windows Server 2012 R2
SQL Server 2012	Windows Server 2012 Windows Server 2012 R2
SQL Server 2014	Windows Server 2012 Windows Server 2012 R2
SQL Server 2016	Windows Server 2012 Windows Server 2012 R2

On-premises SQL Server source to other on-premises targets

This section shows the supported non-SQL Server targets when the source SQL Server database is installed **on-premises** and the target is installed **on-premises**.

Source	Source Platform	Supported Targets
SQL Server 2014	Windows Server 2012	<p>Oracle 11g R2 (11.2.0.3) on:</p> <ul style="list-style-type: none"> • Windows Server 2012 R2 • Linux RHEL 6.2 • Linux RHEL 7.3 • AIX 6.1 • HP-UX 11.31 <p>Oracle 12c (12.1.0.1) on:</p> <ul style="list-style-type: none"> • Windows Server 2012 • Linux RHEL 6.6 • Linux RHEL 6.6 • Sparc Solaris 10 • Intel Solaris 11.2 • AIX 6.1 <p>Kafka 0.10 on Linux RHEL 7.3</p>
SQL Server 2016	Windows Server 2012	<p>PostgreSQL 9.6 on Linux 7.2</p> <p>PostgreSQL 10.0* on Windows 2012R2</p> <p>MySQL 5.7 on Linux 7.2</p>

***NOTE:** SharePlex supports the Community, Fujitsu, and EDB PostgreSQL implementations, with the exception that EDB Postgres Advanced Server version 10 is not supported by SharePlex.

On-premises SQL Server source to all cloud target types

This section shows what SharePlex supports when the source SQL Server database is **on-premises** and the target database is **cloud-based**.

Source source Oracle version	Supported cloud targets
SQL Server 2012	SQL Server 2012 on Azure IaaS SQL Server 2014 on Azure IaaS SQL Server 2012 on Azure PaaS SQL Server 2014 on Azure PaaS SQL Server 2016 on Azure PaaS
SQL Server 2014	SQL Server 2012 on Azure IaaS SQL Server 2014 on Azure IaaS SQL Server 2012 on Azure PaaS SQL Server 2014 on Azure PaaS SQL Server 2016 on Azure PaaS
SQL Server 2016	SQL Server 2012 on Azure IaaS SQL Server 2014 on Azure IaaS SQL Server 2012 on Azure PaaS SQL Server 2014 on Azure PaaS SQL Server 2016 on Azure PaaS

Cloud SQL Server source to all cloud target types

This section shows the supported targets when the source SQL Server database is **cloud-based** and the target is **cloud-based**.

Source SQL Server version	Supported cloud targets
SQL Server 2012 on Azure IaaS	SQL Server 2012 on Azure IaaS SQL Server 2014 on Azure IaaS SQL Server 2012 on Azure PaaS SQL Server 2014 on Azure PaaS SQL Server 2016 on Azure PaaS
SQL Server 2014 on Azure IaaS	SQL Server 2012 on Azure IaaS SQL Server 2014 on Azure IaaS SQL Server 2012 on Azure PaaS SQL Server 2014 on Azure PaaS SQL Server 2016 on Azure PaaS

Conditions of support

This section shows the limitations of capturing from a SQL Server source.

- SQL Server native replication and SharePlex replication cannot be used at the same time for the same database.
- All SQL Server tables in the SharePlex replication configuration must be defined in the database with a primary key. A SharePlex key definition is not sufficient, nor can SharePlex use all columns as a key. You can use SharePlex filtering features in the configuration file to omit tables that do not have a primary key.
- A transaction must COMMIT to the SQL Server source before it can be captured by SharePlex.
- Replication between tables in same database is not supported. Replication between tables in different databases in the same or another SQL Server instance is supported.

Supported operations

- SharePlex supports only DML replication from SQL Server to any of the supported targets.
- SharePlex does not support the replication of SQL Server DDL operations.

The DML operations supported are: INSERT, UPDATE, and DELETE operations between SQL Server source tables and supported target tables that contain supported column types.

NOTES:

SharePlex only replicates DML changes that are made to SQL Server source objects that are listed in the configuration file.

SharePlex requires SQL Server Replication to be installed on the source system. See the *SharePlex Installation Guide* for more information.

Supported data types — SQL Server to SQL Server

SharePlex supports DML operations between SQL Server databases that contain the following SQL Server data types:

- bigint
- binary
- bit
- char
- date
- datetime
- datetime2
- datetimeoffset

- decimal
- float
- image
- integer
- money
- nchar
- ntext
- numeric
- nvarchar
- nvarchar(max)
- real
- smalldatetime
- smallint
- text
- time
- tinyint
- uniqueidentifier
- varbinary
- varbinary(max)
- varchar
- varchar(max)

Supported data types — SQL Server to Oracle and Open Target

The following tables show the supported data type mappings between a SQL Server source and a given Oracle or Open Target database. These are the mappings that have been tested and are known to be successful across a variety of use cases. Mapping to unsupported data types may succeed but have not been tested.

SQL Server to MySQL supported data types

SQL Server	MySQL	Notes
binary (length <= 255)	binary	
binary (length > 255)	longblob	
char (length <= 255)	char	

SQL Server	MySQL	Notes
char (length > 255)	longtext	
date	date	
datetime	datetime	
datetime2	datetime	MySQL datetime supports up to 6 fractional digits, whereas SQL Server datetime2 supports up to 7 fractional digits.
decimal	decimal	
float	float	
image	longblob	
integer	integer	
nchar (length <= 255)	nchar	
nchar (length > 255)	longtext (utf8)	The character set of the target longtext column must be utf8.
ntext	longtext (utf8)	The character set of the target longtext column must be utf8.
numeric	numeric	
nvarchar	varchar	The effective maximum length of a VARCHAR is subject to the maximum row size (65,535 bytes, which is shared among all columns) and the character set used.
nvarchar(max)	longtext (utf8)	The character set of the target longtext column must be utf8.
real	float	
smalldatetime	datetime	
smallint	smallint	
text	longtext	
time	time	
tinyint	tinyint	
varbinary	varbinary	

SQL Server	MySQL	Notes
varbinary(max)	longblob	
varchar	varchar	The effective maximum length of a VARCHAR is subject to the maximum row size (65,535 bytes, which is shared among all columns) and the character set used.
varchar(max)	longtext	

SQL Server to Oracle supported data types

SQL Server	Oracle
bigint	NUMBER
binary	RAW
char	CHAR
date	DATE
datetime	TIMESTAMP
datetime2	TIMESTAMP
datetimeoffset	TIMESTAMP WITH TIME ZONE
decimal	NUMBER
float	NUMBER
image	BLOB
integer	NUMBER
nchar	NCHAR
ntext	NCLOB
numeric	NUMBER
nvarchar	NVARCHAR
nvarchar(max)	NCLOB
real	NUMBER
smalldatetime	DATE
smallint	NUMBER
text	CLOB

SQL Server	Oracle
sql_variant	ANYDATA
float	BINARY_DOUBLE
text	CLOB
time	DATE
tinyint	NUMBER
varbinary	RAW
varbinary(max)	BLOB
varchar	VARCHAR
varchar(max)	CLOB

SQL Server to PostgreSQL supported data types

SQL Server*	PostgreSQL
bit	boolean
char	char
date	date
datetime	timestamp(3)
datetime2	timestamp
datetimeoffset	timestamp with timezone
float	double precision
image	bytea
integer	integer
nchar	char
ntext	text
numeric	numeric
nvarchar	varchar
nvarchar(max)	text
real	real
smalldatetime	timestamp(0)
smallint	smallint
text	text
time	time

SQL Server*	PostgreSQL
tinyint	smallint
varbinary	bytea
varbinary(max)	bytea
varchar	varchar
varchar(max)	text

*NOTE: EDB Postgres Advanced Server version 10 is not supported by SharePlex.

SQL Server to XML supported data types

SQL Server	XML
int	decimal
numeric	decimal
smallint	decimal
tinyint	decimal
float	decimal
real	decimal
date	dateTime
datetime	dateTime
datetime2	dateTime
datetimeoffset	dateTimeStamp
smalldatetime	dateTime
time	dateTime
char	string
nchar	string
text	coob
ntext	clob
nvarchar	string
binary	base64Binary

SQL Server	XML
varbinary	base64Binary
image long	base64Binary

Supported SharePlex features

The following table shows whether or not SharePlex features are supported when SQL Server is the source.

SharePlex feature	Supported with Oracle Target	Supported with OpenTarget
Replication from SQL Server tables to any target where the version of SharePlex is earlier than 9.0.	No	No
Auto-add of new tables that satisfy a wildcard	No	No
flush command	Yes	Yes
reconcile command (target instantiation)	Yes	Yes
compare/compare using and repair/repair using commands	No	No
copy/copy using and append/append using commands	No	No
Hash horizontally partitioned replication	No	No
Column-based horizontally partitioned replication	No	No
Vertically partitioned replication	Yes	Yes
Column mapping	Yes	Yes
Key definition (SQL Server tables must have a defined key)	No	No
Build configuration with scripts	No	No
Named queues	Yes	Yes
Post Enhanced Performance	No	No
Transformation	Yes	No
Conflict resolution	Yes	No
Peer-to-peer replication (bi-directional)	No	No
Consolidated replication (many to one)	Yes	Yes
Broadcast replication (one to many)	Yes	Yes
Cascading replication (source-intermediary-target)	Yes	Yes

SharePlex feature	Supported with Oracle Target	Supported with OpenTarget
High availability replication (active/passive bi-directional)	No	No
Change history target (CDC)	Yes	No
Data encryption	No	No
Data compression	No	No
SSH	No	No
auth_hosts file	No	No
Monitoring scripts	No	No
SNMP monitoring	No	No
Continue posting on error (SP_OPX_CONT_ON_ERR)	Yes	Yes
Suspend on out of sync errors (SP_OPX_OUT_OF_SYNC_SUSPEND)	Yes	Yes
Reduced key (SP_OPX_REDUCED_KEY)	Yes	Yes
sp_ctrl commands	See the command documentation in the Reference Guide.	See the command documentation in the Reference Guide.
Logical Transaction Rollback on out-of-sync transactions	Yes	No

SharePlex preinstallation checklist

Contents

- [Overview](#)
- [Network checklist](#)
- [Installer checklist](#)
- [Windows system checklist](#)
- [SQL Server source checklist](#)
- [Open Target checklist](#)

Overview

Review and satisfy all of the requirements in this checklist before installing SharePlex.

NOTE: The requirements in this checklist apply to all source and target systems where SharePlex will be installed unless otherwise noted.

Network checklist

Requirement	Completed? (Y/N)
<p>Add SharePlex users and groups to the nameserver.</p> <p>If you are installing SharePlex in a network managed by a nameserver such as NIS or NISPLUS, do the following before you install SharePlex:</p> <ul style="list-style-type: none"> • Add SharePlex users to the nameserver. • Add the SharePlex groups to the nameserver. <p>The SharePlex security groups spadmin (administrator), spopr (operator), and spview</p>	

Requirement	Completed? (Y/N)
<p>(viewer) control access to SharePlex processes. Add each SharePlex user to one of these groups on the nameserver. For more information, see Assign SharePlex users to security groups on page 46.</p> <p>To add the user groups:</p> <ol style="list-style-type: none"> 1. For NIS add the groups to the group.byname and group.bygid maps. For NISPLUS, add them to the group.org_dir table. 2. Add the SharePlex Administrator user to the spadmin group on the nameserver. 3. Create the spadmin group in the /etc/group file (on Unix and Linux) or the User Accounts control panel (Windows), and then add the SharePlex Administrator user to the group. <p>To add SharePlex groups to the local system after you install SharePlex, see Assign SharePlex users to security groups on page 46.</p>	
<p>Ensure that SharePlex can resolve host names.</p> <p>If you find that SharePlex cannot connect to a host, try mapping the host name to an alphanumeric alias in the following locations:</p> <ul style="list-style-type: none"> • Network: The NIS and DNS servers • Unix and Linux: Local /etc/hosts file • Windows: Local hosts file <p>In these files, put each entry on an individual line. The following is an example, where sysA and sysB are aliases:</p> <pre>111.22.33.44 sysA.company.com sysA # source system 55.66.77.88 sysB.company.com sysB # target system</pre>	
<p>Resolve to the right network card</p> <p>If you have multiple network cards on Windows, you must set the SP_SYS_HOST_NAME environment variable to the network card name or IP address that you want SharePlex to use as the local host. Otherwise, if the server reboots after the SharePlex configuration is activated, the IP address may bind to a different network card from the one that is bound in the replication configuration.</p> <p>For more information about SharePlex environment parameters and how to set them, see the SharePlex Reference Guide.</p>	
<p>Verify the SharePlex port number.</p> <p>By default SharePlex uses the port number 2100 (hex equivalent is 834) for both TCP/IP and UDP. If port 2100 is available to SharePlex, no further action is needed. You will need to enter the SharePlex port number during the installation procedure, at which time you can specify a different port number if needed.</p> <p>IMPORTANT! The SharePlex port number must be the same one on all machines in the replication configuration so that they can communicate through TCP/IP connections. Make certain the SharePlex port number is open for both TCP/IP and UDP on the firewall.</p>	

Requirement	Completed? (Y/N)
<p>VerifyTCP protocol</p> <p>SharePlex has been tested on IP v6 networks, but it is impossible to test all scenarios. See the System Requirements in the SharePlex Release Notes for more informationon.</p>	

Installer checklist

Requirement	Completed? (Y/N)
<p>Assign a directory to store the downloaded SharePlex installation package.</p> <p>This directory requires approximately the following disk space:</p> <ul style="list-style-type: none"> • Unix and Linux: 200 MB • Windows: 60 MB plus 400 MB of temporary disk space <p>It can be removed after SharePlex is installed.</p>	
<p>Plan the SharePlex product directory.</p> <p>You can create a directory for the SharePlex software files or let the SharePlex installer create it. This directory requires approximately the following disk space:</p> <ul style="list-style-type: none"> • Unix and Linux: 120 MB • Windows: 600 MB plus 20 MB for the MKS Toolkit® <p>Install this directory on the following:</p> <ul style="list-style-type: none"> • Unix and Linux: a separate filesystem from the one that contains the database. • Windows: a separate internal hard drive or partition from the one that contains the database. <p>Do not install SharePlex on a raw device.</p>	
<p>Plan the SharePlex variable-data (working) directory.</p> <p>This directory is installed by the SharePlex installer with a name of your choosing. It contains the working data and varies greatly in size in correlation to the volume of data being generated. Install this directory on a separate filesystem from the one that contains the database, but not on a raw device.</p> <p>To replicate data from more than one database on a system, use a variable-data directory for each one. Ideally they should be on different filesystems.</p> <p>Do not install the variable-data directory within the SharePlex product directory. Both directories contain identically named files, and SharePlex utilities that clean up the environment (if this becomes necessary) could remove the wrong files. You can install both directories under one parent directory if desired.</p>	

Requirement	Completed? (Y/N)
<p>NOTE: Always monitor disk usage when there is an active SharePlex configuration, especially when there are peaks in transaction activity.</p>	
<p>Create the SharePlex security groups.</p> <p>SharePlex provides three security groups to enable access control through sp_ctrl. On Unix and Linux, unless you install SharePlex as a root user, the SharePlex Administrator user and the SharePlex admin group must exist prior to installation. For more information, see Assign SharePlex users to security groups on page 46.</p> <p>NOTE: If you install as root, you are prompted by the installer to create these groups.</p>	
<p>Get a valid SharePlex license key.</p> <p>You must have a valid permanent or trial license key from Quest to run SharePlex. The installer prompts for the license key and the text string in the Site Message that Quest Software provides with the license.</p> <p>The current license model for SharePlex is to license for a specific host, which depending on edition is licensed by core(s) or socket(s) and specific message repository (i.e. database, JMS/text files) etc. Specifics of license terms should be obtained from your account manager.</p>	

Windows system checklist

Requirement	Completed? (Y/N)
<p>Confirm that the Windows version is supported.</p> <p>Review the SharePlex Release Notes to make certain your operating system is supported.</p>	
<p>Be prepared to restart the system.</p> <p>On the Windows platform, SharePlex installs the MKS Toolkit® operating environment from Parametric Technology Corporation (PTC). The default folder for the MKS Toolkit® is C:\Program Files\MKS Toolkit.</p> <p>Set system permissions so that the MKS Toolkit files cannot be moved or removed after they are installed.</p> <p>If this is a first-time MKS Toolkit installation, you will be prompted to restart the system.</p>	
<p>Adjust the page size.</p> <p>SharePlex needs an additional 200 MB of page file size if more than 80 percent of the current total page file size is being used. Greater page size enables SharePlex to process large transactions more quickly.</p>	

Requirement	Completed? (Y/N)
<p>Assign a user who will own the SharePlex directories.</p> <p>Assign a member of the Windows Administrator group to own the SharePlex installation and variable-data directories. This user must exist before you run the SharePlex installer and must have system privileges to read the Oracle redo logs.</p>	
<p>(Oracle) Verify the Oracle Registry entries.</p> <p>(Test machines only) On machines where Oracle has been installed and uninstalled many times, the Oracle entries in the Registry may be corrupted. Before you install SharePlex on a test machine, uninstall all Oracle software and delete all Oracle Registry entries. Then, re-install Oracle by using the Oracle installation program, which creates Registry entries correctly. SharePlex relies on these entries to obtain database environment information.</p>	
<p>(Oracle) Set ORACLE_HOME as the first entry in the PATH variable.</p> <p>SharePlex expects the path to the Oracle binaries to be the first entry in the Windows PATH system variable. Change the variable, if needed, and verify that the path is correct.</p>	

SQL Server source checklist

Requirement	Completed? (Y/N)
<p>Confirm database ownership</p> <p>Make certain that the SQL Server source databases that you will be replicating are owned by a SQL Server user, not a domain user.</p>	
<p>Configure a Data Source Name</p> <p>Create a System (not User) Data Source Name (DSN) for the SQL Server database on the Windows system. The DSN can use either Windows NT authentication or SQL Server authentication. If you configure the DSN to use NT authentication and are using SQL Server 2012 or later, grant the NT Authority\SYSTEM user the sysadmin fixed server role. (For earlier versions of SQL Server, sysadmin is granted to the NT Authority\SYSTEM user by default.)</p> <p>Test connection to the database through this DSN.</p> <p>If you are using SharePlex for multiple databases that are in different SQL Server instances on the same system, each of those databases must have a unique name. Because SharePlex identifies a database by its name, if two or more databases have the same name, SharePlex will only connect to one of them.</p> <p>If databases in different SQL Server instances have the same name but you are only using SharePlex for one of them, the names can remain identical without causing connection conflicts.</p>	

Requirement	Completed? (Y/N)
<p>Install SQL Server Replication Components</p> <p>SharePlex Capture makes use of the underlying components of the native SQL Server replication components. SQL Server Replication must be installed before you install and set up SharePlex, and then the SharePlex Database Setup utility must be run to configure a local Distribution Agent. This utility is typically run as part of the SharePlex installation procedure.</p>	
<p>Be ready to quiesce the source database</p> <p>Before you activate a configuration to capture from a SQL Server source, you must quiesce the tables that are in that configuration. The tables must remain quiesced for the entire activation process.</p>	
<p>Ensure that tables have a primary key</p> <ul style="list-style-type: none"> • To capture from a source SQL Server database, all SQL Server source tables must have a primary key. This is a requirement of the native SQL Server replication, which is used in part by SharePlex for data capture. • All target tables must have corresponding keys. 	
<p>Satisfy requirements for database naming</p> <p>These use cases can cause connection problems for SharePlex unless resolved as recommended.</p> <p>Names of replication databases are identical among local instances</p> <p>If you are using SharePlex for multiple databases that are in different SQL Server instances on the same system, each of those databases must have a unique name. Because SharePlex identifies a database by its name, if two or more databases have the same name, SharePlex will only connect to one of them.</p> <p>If databases in different SQL Server instances have the same name but you are only using SharePlex for one of them, the names can remain identical without causing connection conflicts.</p> <p>Source and target have the same name, but different case</p> <p>SQL Server source and target databases that have the same name but different case collation can cause connection problems for SharePlex. If you cannot either change the case of the databases to be identical, or rename one of them to a different name, you can use the following workaround after you install SharePlex and run the <code>mss_setup</code> utility.</p> <p><i>Workaround:</i></p> <ol style="list-style-type: none"> 1. On the target system, open the <code>connections.yaml</code> file in the SharePlex variable-data directory. 2. Copy the entire set of parameters, but do not delete the original set. 3. Paste the copied set of parameters after the original set. 4. In the pasted set of parameters, change the case of the <code>r.database</code> parameter to match the case of the source database. <p>Example:</p>	

Requirement	Completed? (Y/N)
<pre> r.my_database: database: MY_DATABASE dsn: my_sp1 dstype: sqlserver password: 558ec793ac plugin: sqlserver plugin_direction: target plugin_version: 1 user: qarun r.MY_DATABASE: database: MY_DATABASE dsn: my_sp1 dstype: sqlserver password: 558ec793ac plugin: sqlserver plugin_direction: target plugin_version: 1 user: qarun </pre>	
<p>Considerations if replicating to an Oracle target</p> <p>These use cases apply only if you are replicating from SQL Server to an Oracle target.</p> <p>Character column definitions:</p> <p>SQL Server defines CHAR and VARCHAR data in bytes, whereas Oracle can define it in bytes or characters depending on the semantics definition of the database or the specific table. Additionally, SQL Server allows larger maximum column sizes than Oracle. To allow for these differences in column length, adjustments must be made to the Oracle target table definitions as follows to ensure that the target columns can fit all of the data:</p> <ul style="list-style-type: none"> • For SQL Server char and varchar columns less than or equal to 1000 bytes in length, define the Oracle columns as CHAR and VARCHAR, and specify the length (semantics) as character. • For SQL Server char and varchar columns greater than 1000 bytes in length, define the Oracle columns as CLOB. • For SQL Server nchar columns less than or equal to 1000 characters in length, define the Oracle columns as NCHAR equal in size or greater than the SQL Server ones. • For SQL Server nchar columns greater than 1000 characters in length, define the Oracle columns as NCLOB. • For SQL Server nvarchar columns less than or equal to 2000 characters in length, define the Oracle columns as NVARCHAR equal in size or greater than the SQL Server ones. • For SQL Server nvarchar columns greater than 2000 characters in length, define the Oracle columns as NCLOB. • For SQL Server binary and varbinary columns less than or equal to 2000 bytes in length, define the Oracle columns as RAW equal or greater than the SQL Server ones. • For SQL Server binary and varbinary columns greater than 2000 bytes in length, 	

Requirement	Completed? (Y/N)
<p>define the Oracle columns as BLOB.</p> <p>To view a table representing these relationships, see Set up replication from SQL Server to Oracle in the SharePlex Installation and Setup Guide.</p> <p>Letter Case:</p> <p>To support replication between a source of one database type and a target of another type, the letter case of the names of the source and target columns must be the same, for example the column names on both sides in lower case or both sides in upper case. If the case differs between the source and target column names, use the column mapping feature to map the column names in the configuration file.</p>	
<p>Address varchar(max) issues</p> <p>The following issues are known to exist when capturing SQL Server varchar(max) data types:</p> <p>Length compatibility:</p> <p>If replicating varchar(max) data to a SQL Server target, make certain that the data size does not exceed 1GB in length. Although SQL Server supports varchar(max) data of up to 2 GB in length, the Windows ODBC driver supports VARCHAR(max) data of up to 1 GB in length. Sample ODBC error message, which is included in the Post error message: [ODBC SQL Server Driver]Invalid precision value.</p> <p>Bulk inserts:</p> <p>Bulk inserts that contain tables with large-value varchar(max) columns causes the SQL Server replication components (used by SharePlex) to truncate the values. The truncation is from size 7900 onward. For example, a column value of 8001 bytes is replicated with a size of 101, and a column value of 11000 is replicated with a size of 3100. To work around this issue, configure the table to store large-value data types out-of-row by issuing the following command:</p> <pre>exec sp_tableoption 'owner.table','large value types out of row', 'on'</pre>	

Open Target checklist

All Open Targets

Requirement	Completed? (Y/N)
<p>Perform any required database upgrades.</p> <p>Perform any required database upgrades before you install SharePlex. This ensures that SharePlex gets the most current object definitions when you run Database setup during the installation and setup steps.</p>	
<p>Confirm the database release version.</p>	

Requirement	Completed? (Y/N)
Verify that the release version of the database is supported by SharePlex.	
<p>Consider Case</p> <p>To support replication between a source of one database type and a target of another type, the letter case of the names of the source and target columns must be the same, for example the column names on both sides in lower case or both sides in upper case. If the case differs between the source and target column names, use the column mapping feature to map the column names in the configuration file.</p>	
<p>Install the appropriate ODBC driver.</p> <p>Verify that the appropriate ODBC driver is installed for your target, and install one if there is not one present.</p> <p>For Microsoft SQL Server, make certain to do the following:</p> <p>Install the Microsoft SQL Server ODBC Driver. It must be that driver, <i>not</i> the Microsoft SQL Server Native Client, or SharePlex will return an error when you run the Database Setup utility to configure Post to connect to the database. To tell the difference between the two drivers:</p> <ul style="list-style-type: none"> • The Microsoft SQL Server <u>ODBC Driver</u> has versions such as 06.02.9200. • The Microsoft SQL Server <u>Native Client</u> has versions such as 11.00.3513. <p>To test a driver with SharePlex you can use the OTS utility. You can use a standalone version of this utility before you install SharePlex, or you can run the utility from the installation directory after you install SharePlex. See the OTS documentation in the SharePlex Reference Guide for more information.</p>	
<p>Enable case sensitivity on key columns</p> <p>Enable case-sensitivity for the data of any character-based primary key columns or columns that form a unique index. This ensures that Post compares the correct source and target key values so that it updates the correct target row and prevents unique constraint errors. Unless the key values are case sensitive, cases like the following can happen:</p> <pre>Create table Sales (CustName varchar(20) primary key); insert into Sales values ('abc company');</pre> <p><i>(Succeeds)</i></p> <pre>insert into Sales values ('ABC Company');</pre> <p><i>(Fails with unique constraint violation error)</i></p>	
<p>Disable triggers, cascade deletes, and foreign keys on the target tables.</p> <p>Triggers, cascaded DELETES, and foreign keys must be disabled on Open Target tables. DML changes resulting from triggers, cascaded DELETES, and foreign keys on the source system enter the transaction log and are replicated to the target database by SharePlex. If the same mechanisms are allowed to occur on the target parent table, they initiate changes</p>	

Requirement	Completed? (Y/N)
<p>to the child tables that are duplicated through replication. These duplicate operations cause out-of-sync errors.</p> <p>All tables with foreign keys to one another must all be included in the replication configuration for accurate replication of the source foreign key results. All tables with referential constraints must exist in the target database. If you leave one or more out, the referential integrity could become corrupted.</p>	

SQL Server Target

Requirement	Completed? (Y/N)
<p>Configure a System Data Source Name</p> <p>Create a System (not User) Data Source Name (DSN) for the SQL Server database on the Windows system. The DSN can use either Windows NT authentication or SQL Server authentication. If you configure the DSN to use NT authentication and are using SQL Server 2012 or later, grant the NTAuthority\SYSTEM user the sysadmin fixed server role. (For earlier versions of SQL Server, sysadmin is granted to the NT Authority\SYSTEM user by default.)</p> <p>Test connection to the database through this DSN.</p> <p>IMPORTANT!</p> <p>If you are using SharePlex for multiple databases that are in different SQL Server instances on the same system, each of those databases must have a unique name. Because SharePlex identifies a database by its name, if two or more databases have the same name, SharePlex will only connect to one of them.</p> <p>If databases in different SQL Server instances have the same name but you are only using SharePlex for one of them, the names can remain identical without causing connection conflicts.</p>	

PostgreSQL Target

Requirement	Completed? (Y/N)
<p>Configure a System Data Source Name (Windows installations)</p> <p>If you are replicating data to a PostgreSQL target on Windows, you must use a DSN (data source name) as the connection method. This DSN must exist before you install SharePlex.</p>	
<p>Add a source entry to the PostgreSQL pg_hba.conf file</p> <p>For a PostgreSQL target (non-cloud), make certain that the PostgreSQL pg_hba.conf file has an entry for the SharePlex source server IP address, which gives that server access to</p>	

Requirement	Completed? (Y/N)
PostgreSQL target server and its databases. Example entry: host all all sourceip md5	

Download the SharePlex installer

Contents

- [About the SharePlex installers](#)
- [Where to get SharePlex installers](#)

About the SharePlex installers

There are different installers for SharePlex based on the platform type and datastore type. This topic helps you understand the differences and the naming conventions used.

Linux and Unix

The SharePlex installer on Linux and Unix is a self-extracting installation file with the extension **.tpm**.

Oracle:

There is a separate SharePlex installer build for each supported Oracle database and platform.

SharePlex-release#-build#-DatabaseVersion-platform-version-chipset.tpm

Note: If the operating system version that you have is not listed, choose the highest number that is *below* your version.

Open Target:

The SharePlex Open Target installer supports all Open Target targets on all supported Linux platforms.

SPX-release#-build#-rh-40-amd64-m64.tpm

The installer creates a temporary target directory, within the current directory, for extraction. This temporary target directory is removed upon installation completion. You can extract the files to a file system that is separate from the SharePlex installation location by using the **-t** option when running the **.tpm** file. For additional options, see [Advanced SharePlex installer options](#) on page 110.

Windows

On Windows, the SharePlex installer is named **sp_setup_version.exe**. It is a bundle that contains SharePlex binaries for all of the supported databases and versions, including components for capture from a SQL Server source database.

The installer installs the following items:

- The SharePlex binaries and files
- The **SharePlex for *database*** program
- The **MKS Platform Components** program from Parametric Technology Corporation (default **C:\Program Files\MKS Toolkit**)
- Windows registry entries under **\HKEY_LOCAL_MACHINE\SOFTWAREWow6432node**.
- One or more **SharePlex *port_number*** Windows services (depending on the installed configuration)

Do not remove or modify any of these components while SharePlex is in production, including **SharePlex Installer**. These components all support SharePlex operation or upgrade.

Where to get SharePlex installers

Download the SharePlex installation package that matches the database version and operating system you are using.

Additionally, download any SharePlex patches, so that you can install them after you install the base software.

1. Go to the Quest Software Support page: <http://support.quest.com/>
2. Click **Download Software**.
3. In the search box, type **SharePlex** and press **Go**.
4. Click the arrow in the **Download** column for the version you need. You can also click the file name for access to more information and to download the file(s).
5. Transfer the file to system where you are installing SharePlex.
6. You are ready to begin the installation process. Be sure to thoroughly read the version specific Release Notes *prior* to running the installer.

Install SharePlex on Linux and UNIX

Contents

- [Install SharePlex on Linux/Unix for Oracle Database](#)
- [Install SharePlex on Linux/Unix for Open Target Databases](#)

Install SharePlex on Linux/Unix for Oracle Database

IMPORTANT: These instructions are for installing SharePlex on an Oracle target system to support replication from a SQL Server source to a target Oracle database. To install SharePlex to capture from an Oracle source database, see the SharePlex Installation and Setup Guide for Oracle.

Read this before you begin:

- **These instructions assume that you understand and satisfied all requirements in the [SharePlex preinstallation checklist](#) on page 23.**
- Perform the installation steps on all Unix and Linux machines involved in SharePlex replication. In a cluster, install on the primary node, which is the one to which the shared disk is mounted.
- These instructions assume installation as non-root. To install as root, see [Install SharePlex as Root](#) on page 112.
- The SharePlex security groups and SharePlex Administrator must exist on the system prior to installation. See [Installer checklist](#) on page 25 for more information.
- You can run the installer in either of the following ways:
 - [Run the installer in interactive mode](#) on page 37
 - [Run the installer in unattended mode](#) on page 113

Run the installer in interactive mode

In interactive mode, you are prompted for each part of the installation information.

1. Log in to the system as the user that will be named as the SharePlex Administrator during this installation. This user will own the installation files and binaries.
2. If `sp_cop` is running, shut it down.

```
sp_ctrl> shutdown
```
3. Copy the installation file to a temporary directory where you have write permissions.
4. Grant executable permissions to the file.

```
#chmod 555installation_file
```
5. Run the `.tpm` file. If installing SharePlex in a cluster, run the installer from the primary node (the one to which the shared disk is mounted)

```
# .installation_file
```
6. Verify that the information shown on the first screen corresponds to the Oracle version and platform you are upgrading.
7. You are prompted for the following:

Prompt	Input
Installation type	Select <New Installation> .
Product directory location (path)	Enter the path to the SharePlex installation directory. If the specified directory does not exist, the installer creates it. If the directory exists, it must be empty. The installer quits if the directory contains prior SharePlex installations or other files. In a cluster, install on the shared disk. For more information, see Installation and setup for Oracle cluster on page 1.
Variable data directory location	Specify an empty directory. The installer creates the specified directory if it does not exist. IMPORTANT! Do not install this directory into the SharePlex product directory. In a cluster, install the variable-data directory on the shared disk. For more information, see Installation and setup for Oracle cluster on page 1.
SharePlex Admin group	Enter the DBA-privileged group to which the SharePlex Administrator user belongs, which will own the SharePlex binaries. If the default group of the SharePlex Administrator is <code>oinstall</code> , select any option, and make certain this user is listed under <code>oinstall</code> in the <code>etc/group</code> file. For more information, see Installer

Prompt	Input
	checklist on page 25.
ORACLE_SID of the database	<p>Enter the Oracle SID of the database for which you are installing SharePlex.</p> <p>If you want to use one set of SharePlex binaries to replicate data on a system that contains multiple minor or patch release versions of Oracle Database, select the lowest patch release version that you want to include in replication. Alternatively, you can install separate instances of SharePlex for each Oracle version.</p> <p>In an Oracle cluster, the ORACLE_SID must be an alias in the <code>tnsnames.ora</code> file that is used for all cluster nodes. For more information, see Set up SharePlex in an Oracle cluster on page 1.</p>
ORACLE_HOME	Enter the path to the Oracle HOME directory of the selected Oracle SID.
TCP/IP port for SharePlex	Enter the port number to use for SharePlex TCP/IP communications.
License key (do you have?)	<p>Press Enter to accept the default of Y (yes). If you do not have a license, enter no.</p> <p>For licensing on a cluster, see Installation and setup for Oracle cluster on page 1</p> <p>At any point before you run SharePlex you can add the license key with the <code>splex_add_key</code> utility. For more information, see SharePlex license utilities on page 106.</p>
License key	Enter the license key you received from Quest.
Customer name	Enter the SiteMessage text string provided by Quest with the license key.

The installer displays the location of the install log file and then quits.

See [Next steps](#).

Next steps

Task	Description
Patch SharePlex	If you downloaded patches for this version of SharePlex, apply them now.
Run Database Setup	Run the Database Setup utility for Oracle (<code>ora_setup</code>) to establish a database account and connection information for SharePlex. For more information, see Database Setup for

Task	Description
	Oracle on page 90.
(Heterogeneous replication)	Install SharePlex on the Open Target system and perform database setup operations. See: <ul style="list-style-type: none"> • Install SharePlex on Linux/Unix for Open Target Databases on page 39 • Install SharePlex on Windows on page 42 • Set up replication from Oracle to a different database type on page 1
Set up security	For more information, see Assign SharePlex users to security groups on page 46.
Multi-instance configurations	To install multiple instances of SharePlex on this system, such as to support consolidated replication, see Run SharePlex in the SharePlex Administrator's Guide .

Install SharePlex on Linux/Unix for Open Target Databases

Read this before you begin:

- These instructions assume that you understand and satisfied all preinstallation requirements that apply to your environment. See the [SharePlex preinstallation checklist on page 23.](#)
- For preinstallation information about installing SharePlex for target databases hosted on Amazon cloud services, see [Installation and setup for cloud-hosted databases](#)
- These instructions assume installation as non-root. To install as root, see [Install SharePlex as Root on page 112.](#)
- The SharePlex security groups and SharePlex Administrator must exist on the system prior to installation. See [Installer checklist on page 25](#) for more information.
- You can run the installer in either of the following ways:
 - [Run the installer in interactive mode on page 39](#)
 - [Run the installer in unattended mode on page 113](#)

Run the installer in interactive mode

In interactive mode, you are prompted for each part of the installation information.

1. Log in to the system as the user that will be named as the SharePlex Administrator during this installation. This user will own the installation files and binaries.
2. (Reinstallations) If `sp_cop` is running, shut it down.
3. Copy the installation file to a temporary directory where you have write permissions.
4. Grant executable permissions to the file.

```
# chmod 555 installation_file
```

- Run the .tpm file. If installing SharePlex in a cluster, run the installer from the primary node (the one to which the shared disk is mounted)

```
# .Installation_file
```

- You are prompted for the following:

Prompt for:	Input
Installation type	Select <New Installation>.
Product directory location (path)	Enter the path to the SharePlex installation directory. If the specified directory does not exist, the installer creates it. If the directory exists, it must be empty. The installer quits if the directory contains prior SharePlex installations or other files.
Variable data directory location	Specify an empty directory. The installer creates the specified directory if it does not exist. IMPORTANT! Do not install this directory into the SharePlex product directory.
SharePlex Admin group	Enter the DBA-privileged group to which the SharePlex Administrator user belongs, which will own the SharePlex binaries.
TCP/IP port for SharePlex	Enter the port number to use for SharePlex TCP/IP communications.
License key (do you have?)	Press Enter to accept the default of Y (yes). If you do not have a license, enter no . At any point before you run SharePlex you can add the license key with the splex_add_key utility. For more information, see SharePlex license utilities on page 106.
License key	Enter the license key you received from Quest.
Customer name	Enter the SiteMessage text string provided by Quest with the license key.

The installer displays the location of the install log file and then quits.

See [Next steps](#).

Next steps

Task	Description
Patch SharePlex	If you downloaded patches for this version of SharePlex, apply them now.
Set up security	For more information, see Assign SharePlex users to security groups on page 46.

Task	Description
Perform Database Setup and other setup tasks	See: Set up replication from SQL Server to a different target type on page 52
Repeat	Repeat all of the installation procedures for all Unix and Linux machines that will be involved in SharePlex replication.
Multi-instance configurations	To install multiple instances of SharePlex on this system, such as to support consolidated replication, see the SharePlex Administrator's Guide for the correct setup.

Install SharePlex on Windows

Contents

[Install SharePlex on Windows](#)

[Change global settings in the MKS toolkit](#)

Install SharePlex on Windows

Read this before you begin:

- **These instructions assume that you understand and satisfied all requirements in the [SharePlex SharePlex preinstallation checklist on page 23](#).**
- Perform the installation steps on all Windows machines involved in SharePlex replication. On Windows, install SharePlex on **all nodes** of each cluster, on the same port number, with identical path names. This is required to make the binaries and the required MKS Toolkit components available to all nodes, and to establish Registry entries.

To run the installer

1. Log into Windows as the SharePlex Administrator.
2. (Reinstallation only)
 - a. Run **SpUtils** from the shortcut on the Windows desktop. If the system is running Windows 2008 or higher, run **SpUtils** as an Administrator.
 - b. Select the **SharePlex Services** tab.
 - c. Select the correct port, and then stop the SharePlex service.
 - d. Close the utility.
3. Run the **sp_setup** installation program and follow the prompts:

Prompt	Input
Destination folder	SharePlex installation directory (known as the <i>product</i> directory). Specify an empty existing directory or a new one to be created.
Select components	Specify the database for which you are installing SharePlex.
Port number	The port number on which this instance of SharePlex will run. The default is 2100. Do not use port 2101. That port is reserved for the SpRemote utility.
Variable Data directory	The directory where SharePlex stores its working files and data. Specify an empty existing directory or a new one to be created. To avoid permissions errors during runtime, do not install the variable-data directory in the Program Files directory.
Program Manager group	The folder in the Programs menu where SharePlex programs will be launched.
MKS Platform Components	Appears if SharePlex was never installed on the system or if this is a reinstallation. Accept the default Program Files location or specify a different one. If prompted to restart your system, you can postpone the restart until after you finish this installation.
Confirm installation	Confirm the installation information.
SharePlex license	Click Add License . Type or paste the following exactly as received from Quest Software, including any spaces: <ul style="list-style-type: none"> a. License Key: The license key, which is case sensitive. b. Customer Name: The Site Message text that was included with the license. The name is case-sensitive. c. Click OK. <p>If you do not have a license, click Cancel. The installation succeeds without a license, but SharePlex cannot be started. You can add the license later from the License Key tab of the SpUtils utility, which is installed on the Windows desktop.</p>
Finish	If you were prompted to restart the system after you

Prompt**Input**

installed the MKS Toolkit files, you may do so after exiting the installer.

Proceed to **Next Steps**.

Next steps:

The following setup tasks must be finished before you start up replication.

Task	Description
Install the SharePlex service	For more information, see Install the SharePlex service on page 109.
Patch SharePlex	If you downloaded patches for this version of SharePlex, apply them now.
Perform Database Setup and other setup tasks	If installing SharePlex for a SQL Server source database, run the Database Setup utility. For more information, see Database Setup for SQL Server on page 101. If installing SharePlex for an Open Target database, perform required additional setup: Set up replication from SQL Server to a different target type on page 52
Set up security	For more information, see Assign SharePlex users to security groups on page 46.
Configure the MKS Toolkit	For more information, see Change global settings in the MKS toolkit on page 44.
Repeat	Repeat all of the installation procedures for all Windows machines that will be involved in SharePlex replication.
Multi-instance configurations	To install multiple instances of SharePlex on this system, such as to support consolidated replication, see Run multiple instances of SharePlex in the SharePlex Administrator's Guide . Each instance of SharePlex must be installed on a different port number.

Change global settings in the MKS toolkit

After you install SharePlex on a Windows system, you need to change the default setting for Global Resources memory in the MKS Toolkit.

1. From the Windows Control Panel, select **Configure MKS Toolkit**.
2. Select the **Runtime Settings** tab.
3. Select **Miscellaneous Settings** from the **Categories** pull-down menu.
4. Under **Global Settings** in the **Max Memory for Global Resources** text box, change the value by typing **67108864** in the box.

5. **Click Apply**, then **OK** to close the dialog box.
6. Restart the system.

Assign SharePlex users to security groups

Contents

- [Overview](#)
- [About the SharePlex security groups](#)
- [Create and populate SharePlex groups on Unix and Linux](#)
- [Create and populate SharePlex groups on Windows](#)

Overview

The SharePlex security groups provide access control to the SharePlex command and control system. Without proper configuration of these groups, anyone with permissions on the system can use the commands that view, configure, and control data replication.

About the SharePlex security groups

To monitor, control, or change SharePlex replication, a person must be assigned to one of the SharePlex security groups on the systems where he or she will be issuing commands. Each group corresponds to an authorization level, which determines which SharePlex commands a person can issue. To execute a command, a user must have that command's authorization level or higher.

Use the **authlevel** command to determine your authorization level for issuing SharePlex commands on a system.

Description of the SharePlex security groups

Refer to the following table to determine the group and authorization level that you want to grant each SharePlex user.

User Authorization Levels and Roles

Auth level	User type	User group	User roles
1	Administration	spadmin*	<p>You need at least one user with Administrator rights on each source and target system.</p> <p>Can issue all SharePlex commands. Commands that can <i>only</i> be issued by a SharePlex Administrator are:</p> <ul style="list-style-type: none">• startup, shutdown• all configuration commands relating to an active configuration• all parameter commands except list param• start capture• stop capture• abort capture• truncate log
2	Operator	spopr	Can issue all SharePlex commands except those listed above.
3	Viewer	spview	Can view lists, status screens, and logs to monitor replication only.

NOTE: The default name for the SharePlex administrator group is **spadmin**, but you can designate any group or specify any name for that group during installation.

Create and populate SharePlex groups on Unix and Linux

Where and when to create the SharePlex groups on Unix and Linux depends on whether you install SharePlex as a root or non-root user.

- If you install as non-root, create the groups in the **/etc/group** file before you run the SharePlex installer. In a cluster, create them on all nodes.*
- If you install SharePlex as a root user, you can direct the installer to create the groups in the **/etc/group** file. If you install in a cluster, the installer creates the groups on the primary node, but you must create them yourself on the other nodes.

* The groups must exist because the installer adds the SharePlex Administrator user to the **spadmin** group during the installation process. In a cluster, this user is only added to the primary node. You must add the SharePlex Administrator user to the other nodes.

To create the groups in **/etc/group**

```
# groupadd spadmin
# groupadd spopr
# groupadd spview
```

To assign a user to a group

1. Open the `/etc/group` file.
2. Add the Unix or Linux user name to the appropriate group. To assign a list of user names to a group, use a comma-separated list (see the following example).

```
spadmin:*:102:spadmin,root,jim,jane,joyce,jerry
```

If the password field is null, no password is associated with the group. In the example, the asterisk (*) represents the password, "102" represents the numerical group ID, and **spadmin** is the group. The group ID must be unique.

3. Save the file.

Users can verify their authorization levels by issuing the **authlevel** command in **sp_ctrl**.

Create and populate SharePlex groups on Windows

On Windows, the SharePlex groups are created in the Windows **User Accounts** control panel by the SharePlex installer. To assign users to these groups, use that control panel after you install SharePlex.

Users can verify their authorization levels by issuing the **authlevel** command in **sp_ctrl**.

Set up replication between SQL Server databases

Contents

- [About these instructions](#)
- [Configure SharePlex on the source](#)
- [Configure SharePlex on the target](#)

About these instructions

SharePlex can post replicated SQL Server source data to a Microsoft SQL Server target database through an Open Database Connectivity (ODBC) interface. These instructions contain the basic setup steps to take to support this target.

For additional information, see the following:

- For the datatypes and operations that are supported when using SharePlex to replicate from and to a SQL Server database, see [System Requirements — SQL Server Capture](#) on page 12 .
- For additional configuration options, activation steps, and monitoring information, see the [SharePlex Administration Guide](#).
- For reference documentation on SharePlex commands, parameters and utilities, see the [SharePlex Reference Guide](#).

Configure SharePlex on the source

Configure SharePlex and the database on the source system as follows.

Run Database Setup

If you did not run Database Setup for SQL Server during the installation of SharePlex, run the utility now to establish a database account and connection information for SharePlex. For more information, see [Database Setup for SQL Server](#) on page 1.

Ensure that all tables have a primary key

To replicate from a source SQL Server database to a target SQL Server database, all SQL Server source tables must have a primary key. This is a requirement of the native SQL Server replication, which is used in part by SharePlex for data capture. See the [SQL Server source checklist](#) on page 27 for more information about these and other pre-configuration requirements for a SQL Server source.

All Oracle targets must have corresponding keys.

Ensure varchar(max) length compatibility

If replicating varchar(max) data to a SQL Server target, make certain that the data size does not exceed 1GB in length. Although SQL Server supports varchar(max) data of up to 2 GB in length, the Windows ODBC driver supports VARCHAR(max) data of up to 1 GB in length. Sample ODBC error message, which is included in the Post error message: [ODBC SQL Server Driver]Invalid precision value.

Configure replication

To configure replication from a SQL Server source to a SQL Server target, use the following syntax in the configuration file on the source system.

NOTE: See [Configure SharePlex to replicate data](#) in the [SharePlex Administration Guide](#) for additional information about creating a configuration file.

```
Datasource:r.database_name
```

```
src_owner.table
```

```
tgt_owner.table
```

```
host@r.database_name
```

where:

- *r*. identifies the database as non-Oracle.
- *database_name* is the name of the SQL Server database. **IMPORTANT!** *database_name* must be the *actual name* of the database, not a data source name (DSN).
- *src_owner.table* is the owner and name of the source table.
- *tgt_owner.table* is the owner and name of the target table.*
- *host* is the name of the target system.

IMPORTANT!

- If a database is case-sensitive, enclose the case-sensitive object names in quotes.
- If the letter case of the column names on the source is different from the letter case of the target columns, for example the source is all capitals while the target is lower case, use the column mapping feature to map the column names in the configuration file. For more information, see [Map source and target columns](#) in the SharePlex [Administration Guide](#).

Source configuration example

The following configuration file replicates table `HR.EMP` from the source to target table `Region1.Emp` on target system. The target table is case-sensitive.

```
Datasource:r.mssql  
HR.EMP          "Region1"."Emp"          sysprod@r.mydb
```

Configure SharePlex on the target

If you did not run Database Setup for SQL Server during the installation of SharePlex, run the utility now to establish a database account and connection information for SharePlex. For more information, see [Database Setup for SQL Server](#) on page 101.

Set up replication from SQL Server to a different target type

Contents

- [About these instructions](#)
- [Set up replication from SQL Server to Oracle](#)
- [Set up replication from SQL Server to MySQL](#)
- [Set up replication from SQL Server to PostgreSQL](#)
- [Set up replication from SQL Server to Kafka](#)

About these instructions

This chapter contains instructions for configuring SharePlex to replicate from SQL Server to another type of target, for example SQL Server to PostgreSQL. This is known as heterogeneous replication.

These instructions highlight specific tasks that are pertinent to the flow of data between source and target. Refer to other topics in the SharePlex documentation as needed to complete the configuration, deploy any optional features that apply, and monitor and maintain the environment.

For additional information, see:

- For the SharePlex-supported datastores, datatypes and operations that are supported by SharePlex, see [System Requirements — SQL Server Capture](#) on page 12.
- For additional configuration options, activation steps, and monitoring information, see the [SharePlex Administration Guide](#).
- For reference documentation on SharePlex commands, parameters and utilities, see the [SharePlex Reference Guide](#).

Set up replication from SQL Server to Oracle

SharePlex can capture from a SQL Server source database and replicate the data to an Oracle target. These instructions provide an overview of the steps required to support replication between these database types.

For the versions, datatypes and operations that are supported when using SharePlex to replicate to this target, see the SharePlex [Release Notes](#).

Review column names

To support replication between a source of one database type and a target of another type, the letter case of the names of the source and target columns must be the same, for example the column names on both sides in lower case or both sides in upper case. If the case differs between the source and target column names, use the column mapping feature to map the column names in the configuration file.

See the SharePlex [Administration Guide](#) for more information about column mapping with SharePlex.

Ensure column length compatibility

SQL Server defines CHAR and VARCHAR data in bytes, whereas Oracle can define it in bytes or characters depending on the semantics definition of the database or the specific table. Additionally, SQL Server allows larger maximum column sizes than Oracle. To allow for these differences in column length, adjustments must be made to the Oracle target table definitions as follows to ensure that the target columns can fit all of the data:

- For SQL Server char and varchar columns **less than or equal to 1000 bytes** in length, define the Oracle columns as CHAR and VARCHAR, and specify the length (semantics) as **character**.
- For SQL Server char and varchar columns **greater than 1000 bytes** in length, define the Oracle columns as CLOB.
- For SQL Server nchar columns **less than or equal to 1000 characters** in length, define the Oracle columns as NCHAR equal in size or greater than the SQL Server ones.
- For SQL Server nchar columns **greater than 1000 characters** in length, define the Oracle columns as NCLOB.
- For SQL Server nvarchar columns **less than or equal to 2000 characters** in length, define the Oracle columns as NVARCHAR equal in size or greater than the SQL Server ones.
- For SQL Server nvarchar columns **greater than 2000 characters** in length, define the Oracle columns as NCLOB.
- For SQL Server binary and varbinary columns **less than or equal to 2000 bytes** in length, define the Oracle columns as RAW equal or greater than the SQL Server ones.
- For SQL Server binary and varbinary columns **greater than 2000 bytes** in length, define the Oracle columns as BLOB.

The following chart represents these relationships:

SQL Server Source column definition	Length (bytes)	Required Oracle column definition
char <i>length</i>	<=1000	CHAR(<i>length</i> char)
	>1000	CLOB
varchar <i>length</i>	<=1000	VARCHAR(<i>length</i> char)
	>1000	CLOB
nchar <i>length</i>	<=1000	NCHAR(<i>length</i>)
	>1000	NCLOB
nvarchar <i>length</i>	<=2000	NVARCHAR(<i>length</i>)
	>2000	NCLOB
binary	<=2000	RAW(<i>length</i>)
	>2000	BLOB

Configure SharePlex on the source

Configure SharePlex and the database on the source system as follows.

Run Database Setup

Run Database Setup for SQL Server to establish a database account and connection information for SharePlex. See [Database Setup for SQL Server](#) on page 101.

Ensure that all tables have a primary key

To replicate from a source SQL Server database to a target Oracle database, all SQL Server source tables must have a primary key. This is a requirement of the native SQL Server replication, which is used in part by SharePlex for data capture. See the [SQL Server source checklist](#) on page 27 for more information about these and other pre-configuration requirements for a SQL Server source.

All Oracle target tables must have corresponding keys.

Configure replication

To configure replication from SQL Server to Oracle, use the following syntax in the configuration file on the source system.

NOTE: See [Configure SharePlex to replicate data](#) in the [SharePlex Administration Guide](#) for additional information about creating a configuration file.

```
Datasource:r.database_name
```

```
src_owner.table
```

```
tgt_owner.table
```

```
host@o.SID
```

where:

- `r`. identifies the source database as non-Oracle, in this case SQL Server.
- `database_name` is the name of the SQL Server database. **IMPORTANT!** `database_name` must be the *actual name* of the database, not a data source name (DSN).
- `src_owner.table` is the owner and name of the source table.
- `tgt_owner.table` is the owner and name of the target table.*
- `host` is the name of the target system.
- `o`. identifies the target database as Oracle.
- `SID` is the ORACLE_SID of the target Oracle database.

IMPORTANT!

- If a database is case-sensitive, enclose the case-sensitive object names in quotes.
- If the letter case of the column names on the source is different from the letter case of the target columns, for example the source is all capitals while the target is lower case, use the column mapping feature to map the column names in the configuration file. See [Map source and target columns](#) in the [SharePlex Administration Guide](#) for more information.

Source configuration example

The following configuration file replicates table `HR.EMP` from the source to target table `Region1.Emp` on target system. The target table is case-sensitive.

```
Datasource:r.mssl
HR.EMP          "Region1"."Emp"          sysprod@o.Oracle
```

Configure SharePlex on the target

1. Make certain that the database setup meets all of the requirements in [Set up Oracle database objects for replication](#) on page 1.
2. Run Database Setup for Oracle to establish a database account and connection information for SharePlex. See [Database Setup for Oracle](#) on page 90.
3. See [Set up an Oracle environment for replication](#) on page 1 for additional Oracle setup instructions.

Set up replication from SQL Server to MySQL

SharePlex can capture from a SQL Server source database and replicate the data to a MySQL target. These instructions provide an overview of the steps required to support replication between these database types.

For the versions, datatypes and operations that are supported when using SharePlex to replicate to this target, see the SharePlex [Release Notes](#).

Review column names

To support replication between a source of one database type and a target of another type, the letter case of the names of the source and target columns must be the same, for example the column names on both sides in lower case or both sides in upper case. If the case differs between the source and target column names, use the column mapping feature to map the column names in the configuration file.

See the SharePlex [Administration Guide](#) for more information about column mapping with SharePlex.

Configure SharePlex on the source

Configure SharePlex and the database on the source system as follows.

Run Database Setup

Run Database Setup for SQL Server to establish a database account and connection information for SharePlex. See [Database Setup for SQL Server](#) on page 101.

Ensure that all tables have a primary key

- To replicate from a source SQL Server database to a target MySQL database, all SQL Server source tables must have a primary key. This is a requirement of the native SQL Server replication, which is used in part by SharePlex for data capture. See the [SQL Server source checklist](#) on page 27 for more information about these and other pre-configuration requirements for a SQL Server source.
- All MySQL target tables must have corresponding keys.

Configure replication

To configure replication from SQL Server to MySQL, use the following syntax in the configuration file on the source system.

NOTE: See [Configure data replication](#) in the SharePlex [Administration Guide](#) for additional information about creating a configuration file.

```
Datasource:r.database_name
```

```
src_owner.table
```

```
tgt_owner.table
```

```
host@r.database_name
```

where:

- *database_name* is the name of the SQL Server database. **IMPORTANT!** *database_name* must be the *actual name* of the database, not a data source name (DSN).
- *src_owner.table* is the owner and name of the source table.
- *tgt_owner.table* is the owner and name of the target table.*
- *host* is the name of the target system.
- *database_name* is the name of the target database.

IMPORTANT!

- If a database is case-sensitive, enclose the case-sensitive object names in quotes.
- If the letter case of the column names on the source is different from the letter case of the target columns, for example the source is all capitals while the target is lower case, use the column mapping feature to map the column names in the configuration file. See [Map source and target columns](#) in the [SharePlex Administration Guide](#) for more information.

Source configuration example

The following configuration file replicates table `HR.EMP` from the source to target table `Region1.Emp` on target system. The target table is case-sensitive.

```
Datasource:r.mssl  
hr.emp          region1.emp          sysprod@r.mydb
```

Configure SharePlex on the target

1. Make certain that the database setup meets all of the requirements in [Open Target checklist](#) on page 30.
2. Run Database Setup for MySQL ([mysql_setup](#)) to establish a database account and connection information for SharePlex. For more information, see [Database Setup for MySQL](#) on page 88.

Set up replication from SQL Server to PostgreSQL

SharePlex can capture from a SQL Server source database and replicate the data to a PostgreSQL target. These instructions provide an overview of the steps required to support replication between these database types.

For the versions, datatypes and operations that are supported when using SharePlex to replicate to this target, see the SharePlex [Release Notes](#).

Review column names

To support replication between a source of one database type and a target of another type, the letter case of the names of the source and target columns must be the same, for example the column names on both sides in lower case or both sides in upper case. If the case differs between the source and target column names, use the column mapping feature to map the column names in the configuration file.

See the SharePlex [Administration Guide](#) for more information about column mapping with SharePlex.

Configure SharePlex on the source

Configure SharePlex and the database on the source system as follows.

Run Database Setup

Run Database Setup for SQL Server to establish a database account and connection information for SharePlex. See [Database Setup for SQL Server](#) on page 101.

Ensure that all tables have a primary key

- To replicate from a source SQL Server database to a target PostgreSQL database, all SQL Server source tables must have a primary key. This is a requirement of the native SQL Server replication, which is used in part by SharePlex for data capture. See the [SQL Server source checklist](#) on page 27 for more information about these and other pre-configuration requirements for a SQL Server source.
- All PostgreSQL target tables must have corresponding keys.

Configure replication

To configure replication from SQL Server to PostgreSQL, use the following syntax in the configuration file on the source system.

NOTE: See [Configure SharePlex to replicate data](#) in the [SharePlex Administration Guide](#) for additional information about creating a configuration file.

```
Datasource:r.database_name
```

```
src_owner.table
```

```
tgt_owner.table
```

```
host@r.database_name
```

where:

- *database_name* is the name of the SQL Server database. **IMPORTANT!** *database_name* must be the *actual name* of the database, not a data source name (DSN).
- *src_owner.table* is the owner and name of the source table.
- *tgt_owner.table* is the owner and name of the target table.*
- *host* is the name of the target system.
- *database_name* is the name of the target database.

IMPORTANT!

- If a database is case-sensitive, enclose the case-sensitive object names in quotes.
- If the letter case of the column names on the source is different from the letter case of the target columns, for example the source is all capitals while the target is lower case, use the column mapping feature to map the column names in the configuration file. See [Map source and target columns](#) in the [SharePlex Administration Guide](#) for more information.

Source configuration example

The following configuration file replicates table `HR.EMP` from the source to target table `Region1.Emp` on target system. The target table is case-sensitive.

```
Datasource:r.mssql
```

```
hr.emp
```

```
region1.emp
```

```
sysprod@r.mydb
```

Configure SharePlex on the target

1. Make certain that the database setup meets all of the requirements in [Open Target checklist](#) on page 30 .
2. Run Database Setup for PostgreSQL (**pg_setup**) to establish a database account and connection information for SharePlex. For more information, see [Database Setup for PostgreSQL](#) on page 98.

Set up replication from SQL Server to Kafka

Overview

The SharePlex Post process can connect and write to a Kafka broker. The data can be written in JSON or XML output as a sequential series of operations as they occurred on the source, which can then be consumed by a Kafka consumer.

These instructions contain setup instructions that are specific to this target. Install SharePlex on the source and target according to the appropriate directions in this manual before performing these setup steps.

For the versions, datatypes and operations that are supported when using SharePlex to replicate to this target, see the SharePlex [Release Notes](#).

Guidelines for posting to Kafka

- A SharePlex Post process acts as a Kafka producer. A SharePlex Post process can write to one or more topics that have one or more partitions.
- The SharePlex Post process does not create a topic itself, but you can configure the Kafka broker to auto-create topics.

Configure SharePlex on the source

When replicating data to Kafka, configure the source database and SharePlex on the source system as follows.

Run Database Setup

Run Database Setup for SQL Server to establish a database account and connection information for SharePlex. See [Database Setup for SQL Server](#) on page 101.

Ensure that all tables have a primary key

- To replicate from a source SQL Server database to any target, all SQL Server source tables must have a primary key. This is a requirement of the native SQL Server replication, which is used in part by SharePlex for data capture. See the [SQL Server source checklist](#) on page 27 for more information about

these and other pre-configuration requirements for a SQL Server source.

- All target tables must have corresponding keys.

Configure replication

On the source, create a SharePlex configuration file that specifies capture and routing information. The structure that is required in a configuration file varies, depending on your replication strategy, but this shows you the required syntax for routing data to Kafka.

```
Datasource:r.database_name
```

```
src_owner.table
```

```
!kafka[:tgt_owner.table]
```

```
host
```

where:

- *database_name* is the name of the source SQL Server database.
- *src_owner.table* is the owner and name of the source table.
- **!kafka** is a required keyword indicating SharePlex is posting to Kafka.
- *:tgt_owner.table* is optional and specifies the owner and name of a target table. Use this feature if you want the data to appear as if it came from a table other than the source table. Allow no spaces between **!kafka** and *:tgt_owner.table*. Type case-sensitive names in the correct case and enclose them within double quotes, as in "MySource"."MyTable."
- *host* is the name of the target system.

NOTE: See [Configure SharePlex to replicate data](#) in the [SharePlex Administration Guide](#) for additional information about creating a configuration file.

Source configuration example

```
Datasource:r.mssl
```

```
my_source.my_table !kafka sysprod
```

Configure SharePlex on the target

These instructions configure the SharePlex Post process to connect to Kafka. You must have a running Kafka broker.

To configure post to Kafka

1. Create a Kafka topic.
2. Start **sp_cop**. (Do not activate the configuration yet.)
3. Run **sp_ctrl**.
4. Issue the **target** command to configure posting to a Kafka broker and topic. The following are example commands.

```
sp_ctrl> target x.kafka set kafka broker=host1:9092,host2:9092,host3:9092
```

```
sp_ctrl> target x.kafka set kafka topic=shareplex
```

See [View and change Kafka settings](#) for command explanations and options.

NOTE: Specify more than one broker so that SharePlex will attempt to connect to the other brokers in the list if any one of them is down.

Set the Kafka record format

SharePlex can output to either XML or JSON format as input to Kafka. XML is the default. To set the input format and specify format options, use one of the following **target** commands:

```
target x.kafka set format record=json
```

or:

```
target x.kafka set format record=xml
```

To view samples of these formats, see the **format** category of the **target** command documentation in the SharePlex [Reference Guide](#).

View and change Kafka settings

To view current property settings for output to Kafka, use the following **target** command:

```
target x.kafka show
```

To change a property setting, use the following command.

```
target x.kafka [queue queuename] set kafka property=value
```

where:

- queue *queuename* is the name of a Post queue. Use this option if there are multiple Post processes.
- *property* and *value* are shown in the following table.

Table 1: Kafka target properties

Property	Input Value	Default
broker = <i>broker</i>	Required. The host and port number of the Kafka broker, or a comma delimited list of multiple brokers. This list is the bootstrap into the Kafka cluster. So long as Post can connect to one of these brokers, it will discover any other brokers in the cluster.	localhost:9092
client_id = <i>ID</i>	Optional. A user-defined string that Post will send in each request to help trace calls.	None
compression.code ={ <i>none</i> , <i>gzip</i> , <i>snappy</i> }	Optional. Controls whether data is compressed in Kafka. Options are none , gzip or snappy .	None
partition ={ <i>number</i> <i>rotate</i> <i>rotate trans</i> }	Required. One of the following: <ul style="list-style-type: none">• A fixed partition number: Directs Post to post messages only to the specified partition number. For example, setting it to 0 directs Post to post only to partition 0. This option is suitable for use in testing or if the target has multiple channels of	0

Property	Input Value	Default
	<p>data posting to the same Kafka topic.</p> <ul style="list-style-type: none"> The keyword rotate: Directs Post to apply messages to all of the partitions of a topic in a round-robin fashion. The partition changes with each new message. For example if a topic has three partitions, the messages are posted to partitions 0,1,2,0,1,2, and so on in that order. The keyword rotate trans: This is similar to the rotate option, except that the partition is incremented with each transaction rather than with each message. For example, if a topic has three partitions, the messages are posted to partition 0 until the commit, then to partition 1 until the commit, and so on in that order. This option is suitable if you are replicating multiple tables to a single topic. It allows you to distribute data across several partitions, while still preserving all of the operations of a transaction together in a single partition. This enables a consumer that reads from a single partition to receive a stream of complete transactions. 	
<code>request.required.acks=value</code>	<p>Optional. This is a Kafka client parameter. By default it is set to a value of -1, which means all. Consult the Kafka documentation about this subject, because all really means <i>all in-sync replicas</i>. This parameter can be used in conjunction with the min.insync.replicas broker parameter to tune behavior between availability and data consistency. IMPORTANT: It is possible for data to be lost between a Kafka producer (SharePlex in this case) and a Kafka cluster, depending on these settings.</p>	-1
<code>topic=topic_name</code>	<p>Required. The name of the target Kafka topic.</p> <p>This string may contain the special sequences %o or %t. The %o sequence is replaced by the owner name of the table that is being replicated. The %t sequence is replaced by the table name of the table that is being replicated. This feature may be used in conjunction with a Kafka server setting of auto.create.topics.enabled set to 'true'. Also view your server settings for default.replication.factor and num.partitions because these are used as defaults when topics are auto created. IMPORTANT! If using multiple topics, you must also set the following properties with the target command:</p> <ul style="list-style-type: none"> The output must be in JSON. Set the record property of the format category to json: target x.kafka set format record=json 	shareplex

Property	Input Value	Default
	<ul style="list-style-type: none"> Commits must be disabled. Set the commit property of the json category to no: target x.kafka set json commit=no 	

* To avoid latency, if Post detects no more incoming messages, it sends the packet to Kafka immediately without waiting for the threshold to be satisfied.

Set recovery options

If the Kafka process aborts suddenly, or if the machine that it is running on aborts, row changes may be written twice to the Kafka topic. The consumer must manage this by detecting and discarding duplicates.

Every record of every row-change operation in a transaction has the same transaction ID and is also marked with a sequence ID. These attributes are **id** and **msgIdx**, respectively, under the **txn** element in the XML output (see [Set up replication from SQL Server to Kafka](#)). These two values are guaranteed to be the same if they are re-written to the Kafka topic in a recovery situation.

If desired, you can configure Post to include additional metadata with every row-change record by using the following command:

```
target x.kafka [queue queuename] set metadata property [, property]
```

Table 2: Optional metadata properties

Property	Description
time	The time the operation was applied on the source.
userid	The ID of the database user that performed the operation.
trans	The ID of the transaction that included the operation.
size	The number of operations in the transaction.

Example

```
target x.kafka set metadata time, userid, trans, size
```

To reset the metadata

```
target x.kafka [queue queuename] reset metadata
```

To view the metadata

```
target x.kafka [queue queuename] show metadata
```

Set up replication to a cloud-hosted Open Target database

There are some differences in the way that SharePlex installs in an IaaS cloud environment and in a PaaS cloud environment. These differences are in the installation and configuration of SharePlex. Once installed and configured, SharePlex operates in the cloud the same way that it operates in on-premise installations.

For a list of supported cloud configurations, see the SharePlex Release Notes.

Installation in an IaaS (accessible) environment

If your cloud database service is a true IaaS virtual computing environment, you can install and run a custom application environment, access the operating system, and manage access permissions and storage. In this environment, SharePlex is installed directly on the cloud server just as you would install it locally, without any special setup requirements.

See the regular preinstallation and installation instructions in this manual. To determine the instructions that apply to your environment, review the table of contents or bookmarks as appropriate for the documentation delivery platform that you are using.

Installation in a PaaS (non-accessible) environment

You can install SharePlex to post to a PaaS target in one of the following ways:

- You can use your on-premises production source server to run all of the SharePlex replication components. In this setup, both source *and* target replication processes (and their queues) are installed on one server. The SharePlex Post process connects through a remote connection to the target cloud database.

NOTE: In a high-volume transactional environment, the buildup of data in the post queues and the presence of multiple Post processes may generate unacceptable overhead for a production system. In that case, you should use an intermediary server.

- You can use an on-premises intermediary server to run the Import and Post components (and the post queues). Post connects to the cloud target through a remote connection. This method removes most of the replication overhead from the source server.

NOTE: Because SharePlex connects remotely to a PaaS environment, use the correct database driver for the local machine.

To post to a PaaS target from the source server

All steps are performed on the **source server**.

1. Complete the [SharePlex preinstallation checklist](#) on page 23.
2. Install SharePlex. See:
 - [Install SharePlex on Windows](#) on page 42
3. Run the appropriate database setup utility for the **source database**. See:
 - [Database Setup for SQL Server](#) on page 101

IMPORTANT:

Reply **Y** when asked if the database will be used as a source.

4. Install the appropriate ODBC client of the target cloud database.
5. Run the appropriate database setup utility, this time for the **target cloud database**. See [Database Setup Utilities](#) on page 88.

IMPORTANT:

- If the target is Aurora, use the **mysql_setup** utility.
 - If the target is SQL Server cloud, when asked if the database will be used as a source, enter **N**.
 - (All cloud targets) Specify the full target database name when prompted for the connection string.
6. Specify the following in the routing map of the SharePlex configuration file:
 - a. the name of the **source** server as the target host.
 - b. the name of the **cloud database** as the target database.

In the following example using a MySQL target, **source3** is the source system and **sptest3** is the target database.

```
datasource:r.mysource
#source tables          target tables          routing map
HR.EMP                 "sptest3"."emp"          source3@r.sptest3
```

To post to a PaaS target from an intermediary server

1. Complete the [SharePlex preinstallation checklist](#) on page 23.
2. On the **source** server, install SharePlex for the source database. See:

- [Install SharePlex on Windows](#) on page 42
3. On the **source** server, run the appropriate database setup utility for the source database. See:
 - [Database Setup for SQL Server](#) on page 101

IMPORTANT:

- (Oracle setup only) When asked whether this is a bequeath connection, enter **Y** to use bequeath, unless this system is RAC.
 - Reply **Y** when asked if the database will be used as a source.
4. On the **intermediary** server, install the appropriate ODBC client of the target cloud database.
 5. On the **intermediary** server, install SharePlex for the target cloud database. See:
 - [Install SharePlex on Linux/Unix for Oracle Database](#) on page 36
 - [Install SharePlex on Windows](#) on page 42

6. On the **intermediary** server, run the appropriate database setup utility for the target cloud database. See [Database Setup Utilities](#) on page 88.

IMPORTANT:

- If the target is Aurora, use the **mysql_setup** utility.
 - If the target is SQL Server, reply **N** when asked if this database will be used as a source.
 - (All targets) Specify the full target database name when prompted for the connection string.
7. Specify the following in the routing map of the SharePlex configuration file:
 - a. the name of the intermediary server as the target host.
 - b. the name of the cloud database as the target database.

In the following example using a MySQL target, **intermediary3** is the intermediary system and **sptest3** is the target cloud database.

```
datasource:r.mysource
#source tables          target tables          routing map
HR.EMP                  "sptest3"."emp"          intermediary3@r.sptest3
```

Basic SharePlex demonstration - all platforms

Contents

- [Overview](#)
- [Pework for the demonstrations](#)
- [Create and activate a configuration](#)
- [Demonstration of replication](#)
- [Demonstration of data compare and repair](#)
- [Demonstration of named post queues](#)

Overview

This chapter demonstrates the basics of SharePlex replication. This demonstration can be run on Unix, Linux, or Windows systems for any of the supported SharePlex source and target databases.

NOTES:

- These demonstrations are for use with databases. They do not support replication to a file or a messaging container.
- These are only demonstrations. Do not use them as the basis for deployment in a production environment. To properly implement replication in your environment, follow the instructions in the SharePlex [Installation and Setup Guide](#) and the SharePlex [Administration Guide](#).
- For more information about the commands used in the demonstrations, see the SharePlex [Reference Guide](#).
- The demonstrations assume that SharePlex is fully installed on a source system and one target system, and that any pre- and post-installation setup steps were performed.

What you will learn

- How to activate a configuration
- How SharePlex replicates smoothly from source to target systems
- How SharePlex quickly and accurately replicates large transactions
- How SharePlex queues the data if the target system is unavailable
- How SharePlex resumes from its stopping point when the target system is recovered
- How SharePlex recovers after a primary instance interruption
- How SharePlex replicates an Oracle TRUNCATE command
- How SharePlex verifies synchronization and repairs out-of-sync rows
- How to use named queues to spread the processing of different tables across parallel Post processes

Pework for the demonstrations

Before you run the basic demonstrations, have the following items available.

Tables used in the demonstrations

You will replicate **splex.demo_src** from the source system to **splex.demo_dest** on the target system. These tables are installed by default into the SharePlex schema, which in these demonstrations is "splex." Your SharePlex schema may be different. Verify that these tables exist.

Description of the demo tables.

Column Name	Data Type	Null?
NAME	varchar2(30)	
ADDRESS	verchar2(60)	
PHONE	varchar2(12)	

INSERT scripts

- Create a SQL script named **insert_demo_src** that inserts and commits 500 rows into the **splex.demo_src** table. You will run this script during some of the demonstrations.
- If you will be using the demonstration of named post queues, create a SQL script named **insert_demo_dest** that inserts and commits 500 rows into the **splex.demo_dest** table. You will run this script during some of the demonstrations.

Create and activate a configuration

SharePlex gets its replication instructions from a configuration file, which defines the objects that are to be replicated. The file specifies following:

- The *datasource* (source database) — the identifier of the source database.
- The *source* objects — the names of the objects that contain the data to be replicated.
- The *target* objects — the names of the target objects that will receive the replicated data.
- The *routing map* — the name of the target system and, if the target is a database, its identifier.

NOTE: This demonstration demonstrates replication from one database to another. It does not cover replication to a file or a messaging container.

Create a configuration file

Perform these steps on the **source** system. The demonstration objects are assumed to be in the schema named **splex**.

1. Run **sp_ctrl**.
2. Issue the following command to create a configuration file named **sample_config** in the default text editor.
`sp_ctrl(source)> create config sample_config`
3. In the text editor, build your configuration file based on the appropriate template, as shown in the [Configuration templates](#). Allow no spaces between the characters in a component (source specification, target specification, routing map), and place at least one space between each component.
4. Save the file, then exit the editor. SharePlex automatically saves the file in the **config** sub-directory of the variable-data directory.
5. In **sp_ctrl**, verify that the configuration file will activate successfully.
`sp_ctrl(source)>verify config sample_config`

Configuration templates

SQL Server source to SQL Server or other Open Target target

```
datasource:r.source_database_name
splex.demo_src      splex.demo_dest      target_system@r.database_name
```

where:

- *source_database_name* is the name of the source SQL Server database.
- *target_system* is the name or IP address of the target system.
- *database_name* is the name of the target database.

SQL Server source to Oracle target

```
datasource:r.source_database_name
splex.demo_src      splex.demo_dest      target_system@o.target_SID
```

where:

- *source_database_name* is the name of the source SQL Server database.
- *target_system* is the name or IP address of the target system.
- *target_SID* is the ORACLE_SID of the target database.

Activate the configuration

Perform these steps on the **source** system. When you activate a configuration, SharePlex is ready to capture transactional changes that are made to the specified source data.

1. Activate the configuration.

```
sp_ctrl(source)> activate config sample_config
```

NOTE: Configuration names are case-sensitive.

2. Confirm that the configuration activated successfully. The name **sample_config** should appear under **File Name**, and the word **Active** should appear under **State**.

```
sp_ctrl(source)> list config
```

Troubleshooting Tips

If your configuration activation fails, issue the **view config sample_config** command to view the file. Compare it to the template and make sure all of the information you entered is correct. Make certain you specified the correct database identifier. Check your syntax for extra spaces or missing components.

Because the configuration file is not active, you can edit it directly with the following command:

```
sp_ctrl(source)> edit config sample_config
```

Save the changes, then re-try the activation.

NOTE: To change an active configuration, you must copy it to a new file first, and then edit and activate the copy. For more information, see [Add or change table specifications in an active configuration](#) in the SharePlex Installation and Setup Guide.

Demonstration of replication

This section demonstrates the speed and accuracy of SharePlex replication. It can also be used to verify that SharePlex was installed and configured correctly.

Verify replication startup

This test verifies that replication is working properly.

Perform these steps in the native SQL interface of the database.

1. On the **source**, TRUNCATE **splex.demo_src** to make certain it is empty.
truncate table splex.demo_src;
2. TRUNCATE the target **splex.demo_dest** table.

3. Insert and commit a record into **splex.demo_src**:

```
insert into splex.demo_src values ('Jim', '123 Main Street', '123-456-7890');
commit;
```
4. Verify that the record exists in **splex.demo_dest** on the target system:

```
select * from splex.demo_dest;
```

The query should show the replicated record, and only that record.

Verify replication of large data volumes

This test verifies that SharePlex replicates large volumes of data quickly and accurately. Perform these steps in the native SQL interface of the database.

1. On the **source**, TRUNCATE **splex.demo_src** to make certain it is empty.

```
truncate table splex.demo_src;
```
2. TRUNCATE the target **splex.demo_dest** table.
3. Run the **insert_demo_src** script to insert 500 rows into the **splex.demo_src** table.
4. Verify that all of the record exist in **splex.demo_dest** on the target system:

```
select count (*) from splex.demo_dest;
```

The count should match the number of records inserted by the **Inserts** script.

Verify queuing and continuity of replication

This test shows you how:

- SharePlex queues replicated data on the source system if the target system is unavailable.
 - SharePlex resumes replication from where it left off when the target becomes available.
1. On the **source**, TRUNCATE **splex.demo_src** to make certain it is empty.

```
truncate table splex.demo_src;
```
 2. TRUNCATE the target **splex.demo_dest** table.
 3. On the **target**, shut down SharePlex to simulate that this system is unavailable.

```
sp_ctrl(source)> shutdown
```
 4. On the **source**, run the **insert_demo_src** script to insert records into **splex.demo_src**.
 5. On the **source**, issue the following command to show the status of the local SharePlex queues.

```
sp_ctrl(source)> qstatus
```

The output should show a backlog of messages.

NOTES:

- If SharePlex on the target were running instead of stopped, there would be no messages in the queue. Replication of 500 rows (and the clearing of the queues) would typically happen faster than the time it took to run the script and then issue the **qstatus** command. To confirm that the queues are storing the records, continue to issue the command. The backlog value should be the

same each time.

- If your environment permits it, you can perform the same test with the same results by unplugging the network connection to the source system before you run the script.
6. On the **target**, start **sp_cop** to allow replication to resume.
 7. On the **target**, run **sp_ctrl**.
 8. On the **target**, verify that the SharePlex processes started.

```
sp_ctrl(target)> status
```
 9. On the **target**, verify that the records inserted by the script on the source system now exist in the target database.

```
select count (*) from splex.demo_dest;
```

The count should match the number of records inserted by the **Inserts** script.

Verify SharePlex capture recovery

This test shows how SharePlex recovers after an interruption to data capture.

1. On the **source**, TRUNCATE **splex.demo_src** to make certain it is empty.

```
truncate table splex.demo_src;
```
2. TRUNCATE the target **splex.demo_dest** table.
3. On the **source**, stop the Capture process.

```
sp_ctrl(source)> stop capture
```
4. On the **source**, run the **Insert** script to generate a transaction to **splex.demo_src**.
5. On the **source**, start the Capture process.

```
sp_ctrl(source)> start capture
```
6. On the **source**, verify that Capture is running and that it processed the data. Look at **Capture state** and **Operations captured**.

```
sp_ctrl(source)> status
```
7. On the **target**, use the native SQL interface to verify that all records inserted by the script on the source now exist in the target database.

```
select count (*) from splex.demo_dest;
```

Demonstration of data compare and repair

This demonstration shows you how to use the SharePlex **compare** command to compare the source and target data, and then repair rows that are out of synchronization.

1. On the **source**, TRUNCATE **splex.demo_src** to make certain it is empty.
`truncate table splex.demo_src;`
2. TRUNCATE the target **splex.demo_dest** table.
3. On the **source**, use the **insert_demo_src** script to insert rows into **splex.demo_src**.
4. On the **target**, verify that all of the data posted to **splex.demo_dest**. If the command shows that data is still in the queue, issue the command again.
`sp_ctrl(target)>qstatus`
5. On the **source**, issue the **compare** command in **sp_ctrl**.
`sp_ctrl(source)>compare splex.demo_src`
Note: The command determines the correct target table from the configuration file.
6. On the **source**, view the results of the comparison. There should be no rows out of synchronization.
`sp_ctrl(source)>compare status`
7. On the **target**, UPDATE **splex.demo_dest** to change the values of the NAME column in two or more rows. The UPDATE causes the source and target tables to be out of synchronization.
8. On the **source**, issue the **compare** command again. It should show that the rows you updated in **splex.demo_dest** are out of synchronization.
`sp_ctrl(source)>compare splex.demo_src`
9. On the **source**, issue the **repair** command to repair the rows that are out of synchronization.
`sp_ctrl(source)>repair splex.demo_src`
10. On the **source**, verify that the repair was performed.
`sp_ctrl(source)>repair status`
11. On the **target**, you can manually verify that the repair was accurate by using a SELECT statement to view all rows in both tables.
`select * from splex.demo_src;`
`select * from splex.demo_dest;`

Demonstration of named post queues

This demonstration shows you how to use named post queues to process different tables through different, parallel Post processes to improve performance.

Clean up the replication environment

NOTE: The demonstration objects are assumed to be in the schema named **splex**.

1. If you ran previous demonstrations, do the following:
 - a. On the **source** and **target**, run **sp_ctrl** and issue the following command to shut down **sp_cop**.
`sp_ctrl(source) shutdown`

```
sp_ctrl(target) shutdown
```

- b. On the **source** and **target**, run the **ora_cleansp** or **mss_cleansp** utility according to the instructions in [SharePlex utilities](#) on page 1. This removes the queues from the previous demonstrations and deactivates the previous configuration.
2. On the **source** and **target**, TRUNCATE the **splex.demo_src** and **splex.demo_dest** tables.

```
truncate table splex.demo_src;  
truncate table splex.demo_dest;
```

Create a configuration file

Perform these steps on the **source** system. The demonstration objects are assumed to be in the schema named **splex**.

You will configure the following:

- **demo_src** on the source replicates through a post queue named **q1** to target **demo_dest**.
 - **demo_dest** on the source replicates through a post queue named **q2** to target **demo_src**.
1. Run **sp_ctrl**.
 2. Issue the following command to create a configuration file named **postq_config** in the default text editor.

```
sp_ctrl(source)>create config postq_config
```
 3. In the text editor, build your configuration file based on the appropriate template, as shown in the [Configuration templates](#). Allow no spaces between the characters in a component (source specification, target specification, routing map), and place at least one space between each component.
 4. Save the file, then exit the editor. SharePlex automatically saves the file in the **config** sub-directory of the variable-data directory.
 5. In **sp_ctrl**, verify that the configuration file will activate successfully.

```
sp_ctrl(source)>verify config postq_config
```

Configuration templates

SQL Server source to Oracle target

```
datasource:r.source_database_name
```

```
splex.demo_src          splex.demo_dest          target_system:q1@o.target_SID
```

```
splex.demo_dest        splex.demo_src          target_system:q2@o.target_SID
```

where:

- *source_database_name* is the name of the source SQL Server database.
- *target_system* is the name or IP address of the target system.
- *target_SID* is the ORACLE_SID of the target database.

SQL Server source to SQL Server or other Open Target target

```
datasource:r.source_database_name
```

```
splex.demo_src          splex.demo_dest          target_system:q1@r.database_name
```

```
splex.demo_dest        splex.demo_src          target_system:q2@r.database_name
```

where:

- *source_database_name* is the name of the source SQL Server database.
- *target_system* is the name or IP address of the target system.
- *database_name* is the name of the target SQL Server database.

Activate the configuration

IMPORTANT! Make certain you created the scripts described in [Prewrite for the demonstrations](#) on page 68. Perform these steps on the **source** system. When you activate a configuration, SharePlex is ready to capture transactional changes that are made to the specified source data.

1. Activate the configuration.

```
sp_ctrl (source) > activate config postq_config
```

NOTE: Configuration names are case-sensitive.

2. Confirm that the configuration activated successfully. The name **postq_config** should appear under **File Name**, and the word **Active** should appear under **State**.

```
sp_ctrl (source) > list config
```

Generate data

On the **source**, run the following scripts to insert data into the two local demonstration tables (in this demonstration, the **demo_dest** table is used as a source):

- Run the **insert_demo_src** script to insert 500 rows into the **splex.demo_src** table.
- Run the **insert_demo_dest** script to insert 500 rows into the **splex.demo_dest** table.

View the post queues

1. On the **target**, run **sp_ctrl**.
2. On the **target**, issue the show post command with the detail option. This command shows statistics about the number of messages processed and the queues that were used.

```
sp_ctrl (target) > show post detail
```

In the output, the **Queue** field shows the name of each post queue, in this case **q1** and **q2**, and the **Operations Posted** field shows the number of operations that were posted. Each queue should show 500 operations posted.

Solve installation problems

Contents

- Overview
- Solve license utility problems
- Solve installer problems
- Solve Database Setup problems
- Solve database connection problems
- Solve SharePlex startup problems
- Solve sp_ctrl problems
- Solve host connection problems
- How to find the ORACLE_SID and ORACLE_HOME

Overview

This chapter reviews some common problems that you could experience when installing or running SharePlex for the first time after installation.

Sometimes there are special installation instructions that supersede or supplement certain instructions in this manual. In addition, there can be known issues for this version that you should be aware of during or after installation. Please read the Release Notes for the version of SharePlex that you are installing before you begin the installation process.

Solve license utility problems

Are all machines connected to the network?

The inability of SharePlex components to perform initial TCP operations can sometimes appear to be license key or license utility errors. If you know you entered the correct key and machine IDs, verify that all systems on which you are loading SharePlex are connected to the network. The network node name and IP address of

each system must be established sufficiently to allow SharePlex to perform TCP operations, even though the machines themselves are not yet configured. Also check to make sure that nobody has renamed the `/etc/resolv.conf` file (if using a DNS nameserver).

Did you enter the correct key and/or machine ID number?

If you received this error message: “Cannot add license: License key is illegal,” it could be that you entered an invalid license key. Assuming that you retyped the key correctly and still received an error, it probably means that the license key, though valid, is not the correct key for this system. Except for trial keys, which are generic, license keys are assigned to a specific machine according to the machine’s identification number (such as “host ID” on Sun systems, “machine ID” on HP systems, etc.).

You probably received at least two license keys from Quest — one for a source system and one for a target system — or if you are installing on multiple machines in a cluster, you should have a key for each one. Verify that the key you entered is the one that was issued for this system by comparing it to the machine identification number for which it was issued.

To view the machine ID and add a key on a Windows system

1. Run the **SpUtils** utility from the SharePlex product directory, then click the **License** tab.
2. Enter the license key and the SiteMessage code from the email that your company received from the Quest licensing team.

To view the machine ID and add a key on a Unix or Linux system

Run the `splex_uname` application from the `install` sub-directory of the SharePlex product directory on the machine whose ID number you want to confirm. It displays the ID number for the local machine, as shown in the example below.

```
$ /splex/proddir/install/splex_uname
```

```
Host ID = 2198894273 (831076C1 HEX)
```

Run the `splex_add_key` utility from the SharePlex product directory and add the license key and SiteMessage code from the email that your company received from the Quest licensing team.

Solve installer problems

Is `sp_cop` shut down?

If you installed SharePlex on this system before, and you are re-installing it, the installation will return errors if SharePlex is running on this system. Shut down SharePlex using the `shutdown` command in `sp_ctrl`, or you can shut down the SharePlex service if this is a Windows system. If you are unable to run `sp_ctrl`, or if any SharePlex processes will not die, locate the process (using `ps -ef | grep sp_` on Unix and Linux systems or the **Taskmgr** tab available from the **SpUtils** application provided for Windows systems) and kill it. When all SharePlex processes have been killed, run the installation program again.

Are all systems connected to the network?

Check to see that all systems on which you are loading SharePlex are connected to the network. The network node name and IP address of each system must be established sufficiently to allow SharePlex to perform TCP operations, even though the target machines themselves are not yet configured.

Note: These failures may appear to be license utility errors, but it is usually the inability of the license utilities and other components of SharePlex to perform initial TCP operations. Also check to make sure that nobody has renamed the `/etc/resolv.conf` file (if using a DNS nameserver).

Did you enter the SharePlex groups in the name service?

If your environment uses a name service such as NIS or NISPLUS, you need to add the SharePlex groups and services to the nameserver before you run the SharePlex installation program, and the SharePlex Administrator must be named in the **SharePlex Admin group** on the nameserver before you install SharePlex. Instructions are on page 31. If these procedures are not performed, the installation will generate an error at the point in which it attempts to verify that the groups exist.

Is the database open?

The database must be open while you are installing SharePlex.

Did you specify a valid ORACLE_SID and ORACLE_HOME?

If you specify an invalid ORACLE_SID or ORACLE_HOME value for an Oracle instance, the installation script is unable to locate the correct Oracle libraries to link to, and it will fail with an error such as this:

```
"Cannot find shared library usr1/oracle/ 8.1.6/lib/libclntsh.so; Exiting."
```

Re-run the installation script again to provide the correct values for ORACLE_SID and ORACLE_HOME.

Solve Database Setup problems

This section helps you diagnose problems that are associated with the SharePlex database account and connection information that was created with the Database Setup utility when SharePlex was installed on the system.

NOTE: For more information about Database Setup, see [Database Setup Utilities](#) in the SharePlex Reference Guide.

If the issue you are experiencing is not listed in this documentation, search the SharePlex Knowledge Base at: <https://support.quest.com>.

The Knowledge Base provides filtering options and links to other resources that can help you use and troubleshoot SharePlex.

Oracle setup issues

Problem	Description	Solution
Incorrect ORACLE_SID and/or ORACLE_HOME	If SharePlex cannot interact with Oracle, it might be using the wrong ORACLE_SID and/or ORACLE_HOME.	<ol style="list-style-type: none">1. See How to find the ORACLE_SID and ORACLE_HOME on page 82 to determine the Oracle values.2. Rerun the Database Setup utility. For more

Problem	Description	Solution
		information, see Database Setup Utilities in the SharePlex Reference Guide .
Insufficient database privileges	If the Database Setup utility fails, the person who runs it may not have the correct privileges	For more information, see Database Setup Utilities on page 88.
Asterisk as the ORACLE_SID entry	Sometimes, the oratab file has an * (asterisk) symbol instead of a value for the ORACLE_SID.	Ensure that a valid ORACLE_SID is in the oratab file, and then try running the database setup again.
More than one oratab file (Sun Solaris)	On Solaris systems, the oratab file is typically located in the /var/opt/oracle directory, but because other platforms store the oratab file in the /etc directory, there could be a second oratab in the /etc directory.	Either move, rename or delete the secondary oratab file, and then try running the database setup again.
/etc/loggingroups file exists (HP-UX)	Look for an /etc/loggingroups file on the system. This file existed on HP-UX systems prior to the adaptation of POSIX standards. To allow backward compatibility, HP-UX gives priority to /etc/loggingroups , and uses the /etc/group file only if /etc/loggingroups does not exist.	Edit the /etc/group file to make its entries identical to those in the /etc/loggingroups file, then delete the etc/loggingroups file.
Oracle not running	Oracle must be running and the instance must be open while you run the Database Setup utility. The utility accesses Oracle to establish SharePlex as a user and install its internal tables.	Start Oracle and open the instance.
sp_cop is running	The SharePlex sp_cop process cannot be running while you are running the Database Setup utility.	If it is running, shut it down using the shutdown command in sp_ctrl . Run sp_ctrl from the bin sub-directory in the SharePlex product directory.
Oracle library location not correct	On Unix and Linux systems, SharePlex expects the Oracle library to be in the \$ORACLE_HOME/lib or \$ORACLE_HOME/lib32 directory. In some environments, the Oracle library has a different name than what SharePlex expects it to be, or it is installed in a different location than expected (or both). In that case, you will see an error message when you attempt to run the Database Setup utility.	Install the appropriate library from Oracle and then re-start SharePlex (if it is stopped). SharePlex will link to the correct library from that point forward.
Id.so.1: sqlplus: fatal: libsunmath.so.1: can't open file: errno=2" error	On Unix and Linux systems, this error indicates that SharePlex cannot find the libsunmath and libshareplex libraries, even though the link exists in the proper place.	You can use either of these solutions: <ul style="list-style-type: none"> • Create a softlink for \$ORACLE_HOME/lib/libsunmath.so.1 in the /usr/lib directory. or...

Problem	Description	Solution
Wrong user-id	To run Database Setup on Unix and Linux systems, the set-user-id for the Oracle software need to be rwsr-s-x . Those permissions allow non-Oracle users to log into SQL*Plus.	<ul style="list-style-type: none"> In the ECXpert/config/bdg.ini file in the [DB_ENV] section add the following line: <code>LD_LIBRARYPATH=full oracle home path/lib</code> Set the correct values for set-user-id .

SQL Server setup issues

Problem	Description	Solution
The Database Setup utility fails with cursor or connection error	If the Database Setup utility cannot interact with the SQL Server database, you might have specified the wrong name for the database.	When prompted for the database name in the setup interface, you must specify the actual name of the database, not the SQL Server instance name or the Data Source Name(DSN). Find out the name of the database, then use it when you run the setup again.
Insufficient database privileges	If the Database Setup utility fails, the person who runs it may not have the correct privileges	For more information, see Database Setup Utilities on page 88.
No replication components	SharePlex Capture makes use of the underlying components of the native SQL Server replication components. SQL Server Replication must be installed before you install and set up SharePlex, and then the SharePlex Database Setup utility must be run to configure a local Distribution Agent. This utility is typically run as part of the SharePlex installation procedure.	Install SQL Server replication, then run the Database Setup utility again. For more information, see Database Setup Utilities in the SharePlex Reference Guide .

Solve database connection problems

Did you verify the credentials?

If SharePlex cannot connect to a source or target database, you can view the login credentials that are being used for that database by using the **connection** command with the **show** option. For example:


```
sp_ctrl1> connection r.mydb show
or...
sp_ctrl1> connection o.mydb show
```

You can also view connection settings in the **connections.yaml** file, which is stored in the **data** sub-directory of the SharePlex variable-data directory. If there are no settings in this file, it means that the SharePlex database setup procedure was not performed on this database.

To view the correct database setup procedure for the database, see [SharePlex utilities](#) on page 1.

You can use the **connection** command to update connection properties. For more information, see the SharePlex Reference Guide.

Did you assign a DBA role to the SharePlex Oracle user?

The SharePlex Oracle user requires a DBA role with unlimited privileges. The SharePlex user is created with the default Oracle profile under the assumption that the profile has the unlimited resource privileges assigned by Oracle as the default. If SharePlex is unable to interact with Oracle, check to see if the default was changed. If so, assign SharePlex a DBA role with unlimited privileges for all definitions.

Solve SharePlex startup problems

Was the user an authorized SharePlex user?

Only a member of the SharePlex administrator group (default name is **spadmin**) can start **sp_cop**. A root user that is not a member of this group can start **sp_cop**, but no users (including root) will be able to connect through **sp_ctrl** to issue commands. For more information, see [Assign SharePlex users to security groups](#) on page 46.

Is this a cluster environment?

In order for the SharePlex processes to issue name lookups and migrate properly in a clustered environment (where a package name supersedes the local system name), the **SP_SYS_HOST_NAME** parameter must be set to the correct package name. In addition, the host name set by this parameter must be the same on all members of the cluster so that the name can bind to a socket and the **/etc/hosts** file or nameserver can correctly map the parameter's value to the correct IP address.

The **sp_cop** program should only be started through the cluster management software.

Was the filesystem mounted as nosuid?

On Unix and Linux systems, if the filesystem is mounted as nosuid, SharePlex must be started by the installation owner. In this case, members of the SharePlex administrator group (**spadmin** by default), other than the installation owner, will not be able to run SharePlex.

Solve sp_ctrl problems

Did you assign the users to the SharePlex groups?

Only one SharePlex user, the Administrator who owns the SharePlex binaries and files, is created during SharePlex installation. Other users must be assigned to the appropriate SharePlex user groups. These groups

control the authorization levels for various SharePlex functions.

To issue a specific command (such as **activate config** or **stop export**), a user must have that command's authorization level or higher. For example, a SharePlex Administrator (authorization level 1) can issue any command, but a member of the **spview** group can only issue status commands and a few other commands that do not directly affect the replication processes.

For more information, see [Assign SharePlex users to security groups on page 46](#).

Solve host connection problems

If SharePlex cannot resolve a host name, try creating an alias for it using a simple alphanumeric name. Map the name to the alias in the following locations:

- Network: The NIS and DNS servers
- UNIX: Local **/etc/hosts** file
- Windows: Local **hosts** file

In the hosts files, put each entry on an individual line. The following is an example, where **sysA** and **sysB** are aliases:

```
111.22.33.44 sysA.company.com sysA # source system
55.66.77.88 sysB.company.com sysB # target system
```

How to find the ORACLE_SID and ORACLE_HOME

When setting up SharePlex to work with an Oracle database, you provide the ORACLE_SID and then SharePlex gets the ORACLE_HOME from the Windows Registry or the **oratab** file on Unix/Linux. Both values are stored in the SharePlex environment. SharePlex uses the Oracle libraries that are in the location specified with ORACLE_HOME.

To determine the ORACLE_SID and ORACLE_HOME being used by SharePlex

Issue the **orainfo** command in **sp_ctrl**.

```
sp_ctrl (mysys111:2101)> orainfo
Oracle instance #1:
  Oracle SID ora12
  Oracle HOME /oracle/products/12
  Oracle Version 12
Oracle instance #2:
  Oracle SID ora12
  Oracle HOME /oracle/products/12
  Oracle Version 12
```

To determine the default ORACLE_SID and ORACLE_HOME on Windows

View the Oracle entry in the Registry at `\HKEY_LOCAL_MACHINE\SOFTWARE\ORACLE`.

To determine the default ORACLE_SID and ORACLE_HOME on UNIX and Linux

On most Unix and Linux systems the **oratab** file is under **/etc/oratab**. On Oracle Solaris systems, it is under **/var/opt/oracle**, but sometimes there is an **oratab** file in the **/etc** directory as well.

The entry in the file looks like the following example:

```
qa12:/qa/oracle/ora12/app/oracle/product/12.0
```

In the example, **qa12** is the **ORACLE_SID** and **/qa/oracle/ora12/app/oracle/product/12.0** is the **ORACLE_HOME**.

Remove SharePlex from a system

Contents

- [Overview](#)
- [Remove SharePlex from Unix/Linux](#)
- [Remove SharePlex from Windows](#)

Overview

This section contains instructions for using the SharePlex uninstallation program to remove SharePlex from a system. The uninstaller permanently removes the replication environment from the system.

To preserve the replication environment, including the queues that store the data, you can install a SharePlex upgrade or reinstall the current version, rather than uninstall SharePlex. Before you upgrade or reinstall SharePlex, see the Release Notes for the version you are installing to determine if there are any special upgrade or installation requirements.

Remove SharePlex from Unix/Linux

1. Log on as a user who has privileges to remove the SharePlex directories.
2. Run `sp_ctrl`.
3. Issue the `shutdown` command to shut down SharePlex.

```
sp_ctrl> shutdown
```
4. Exit `sp_ctrl`.
5. Use the following command to make certain that no SharePlex processes are running.

```
# ps -ef | grep sp_
```
6. Kill any processes that are still running.

7. Remove the SharePlex product directory, variable-data directory, and the hidden .shareplex directory from the system.
8. Drop the SharePlex database user from the database.

NOTE:

If desired, remove any objects that SharePlex installs in the database. These objects begin with **SHAREPLEX_**. You can remove any such objects using standard SQL commands.

Remove SharePlex from Windows

Files removed by the uninstall program

When SharePlex was installed, a file named **install.log** was installed in the SharePlex product directory. This file contains a record of all files created on the system by the installation program. When you run the uninstall program to remove SharePlex, it reads **install.log** to determine which files to remove.

If you select the **Automatic** uninstall option, it removes the following components for *all instances of SharePlex* (all SharePlex services on all ports):

- The SharePlex desktop icons.
- The SharePlex menu items (except for the top-level Program Manager Group folder).
- The SharePlex product directory. If the variable-data directory is installed under the product directory but contains no files that were added after installation, it is removed as well.

Use the **Custom** uninstall option to remove specific files while leaving the others intact.

Files not removed by the uninstall program

The uninstall program does not remove the following components.

- The SharePlex and MKS Toolkit entries in the Registry.
- Files that already existed in the product directory when the current version of SharePlex was installed.
- Files created by SharePlex or a user in the product or variable-data directory after the current version of SharePlex was installed. Such files can be removed manually after SharePlex is removed. This is standard procedure for most Windows applications.
- The SharePlex variable-data directory. To remove this directory, delete it through the operating system. The uninstall program does not remove this directory because there could be user-created files, such as configuration files and custom parameter settings, that you want to keep for a future installation of SharePlex.
- The MKS Components programs and files. These can be removed separately using the [Remove the MKS Toolkit operating environment](#) procedure.

Remove SharePlex from the system

Perform these tasks in the order shown.

Remove the SharePlex service

Before you remove the SharePlex software from the system, follow these steps to stop and remove each SharePlex service on the system. If you remove SharePlex first, you cannot use **SpUtils** to stop and remove the service and must do so from the **Services** utility of the Windows Control Panel.

1. Double-click the **SpUtils** desktop icon.
2. Click the **SharePlex Services** tab.
3. Select the port number for the SharePlex instance that you want to remove. The **Current State** field displays the status of SharePlex on that port.
4. If the service is running, click **Stop**. Make certain the status shows *Stopped*.
5. Click **Remove**.
6. Stop and remove other SharePlex services as needed.
7. Close the **SharePlex Services** dialog box.

Remove the SharePlex software

1. From Start menu, click **Programs**, then navigate to the SharePlex program folder and click **Uninstall**. Alternatively, you can use the Windows Control Panel as you normally would to remove software.
2. In the **Select Uninstall Method** dialog box, select an uninstall option.
 - Select **Automatic** to remove everything listed in [Files removed by the uninstall program](#). This is the recommended procedure because it is the cleanest way to remove SharePlex from the system. Click **Finish** to execute the uninstallation.
 - Select **Custom** to selectively remove files. Use this option only if you must retain some files while deleting others. You receive a series of prompts to remove different sets of files.
 - Select **Repair** to re-install files or update Registry entries. Click **Finish** to execute the repair.

Remove the MKS Toolkit operating environment

Changing and removing Windows Registry components may be harmful to the system. Contact your System Administrator for assistance if needed.

1. Remove the **MKS Platform Components** software using the standard Programs uninstall option in the Windows Control Panel.
2. Remove the **NuTCRACKER** service from the **Services** utility of the Windows Control Panel.

Remove Registry entries

The following Registry entries exist and can be removed by a System Administrator.

```
HKEY_LOCAL_MACHINE\SOFTWARE\Wow6432\Node\Datafocus  
HKEY_LOCAL_MACHINE\SOFTWARE\Wow6432\Node\Mortice Kern Systems  
HKEY_LOCAL_MACHINE\SOFTWARE\Wow6432\Node\Quest Software
```

Remove the SharePlex user and database objects

This uninstaller does not remove the SharePlex database objects. All SharePlex objects are installed in the SharePlex schema or database that was specified when the Database Setup portion of the installation was performed.

IMPORTANT! If you intend to reinstall SharePlex or MKS Toolkit, they will not reinstall unless these Registry entries are removed.

Database Setup Utilities

The Database Setup utilities automatically configure a source or target database to allow SharePlex connections and to establish required database components that support replication.

Contents

- [Database Setup for MySQL](#)
- [Database Setup for Oracle](#)
- [Database Setup for PostgreSQL](#)
- [Database Setup for SQL Server](#)
- [Database Setup for Tiberio](#)

Database Setup for MySQL

Overview

Run the Database Setup utility for MySQL (**mysql_setup**) on a MySQL system to establish SharePlex as a MySQL database user. This utility creates the following:

- A SharePlex user account with full DBA privileges
- Tables and indexes for use by SharePlex and owned by the SharePlex user in a database of your choosing
- A default database connection.

Supported databases

MySQL on Linux. For supported Linux platforms and versions, see the SharePlex Release Notes.

Guidelines for use

- Run the setup utility on all MySQL instances in the SharePlex replication configuration.
- Within a cluster, run the setup utility on the node to which the shared disk that contains the SharePlex variable-data directory is mounted.
- For consolidated replication, run the setup utility for each variable-data directory.

Required privileges

Review the following requirements to ensure that the setup succeeds.

- The setup utility must be run as a MySQL Administrator that retains all of that user's default privileges. Both local and cloud MySQL Administrators should have the required privileges by default. The Administrator user is able to grant SharePlex the required privileges to operate on the database and to create the SharePlex database account and objects.
- (Cloud installations) Common restrictions on privileges in cloud-hosted database services make it difficult for the setup utility to succeed in every possible type of scenario. To ensure that the database setup succeeds, *only* use the setup utility for the following purposes: To do a *first-time* database setup with a *new* SharePlex user, or, to *modify* an existing SharePlex user that either owns the database or has access to it.

Run Database Setup for MySQL

1. Shut down any running SharePlex processes and **sp_cop** on the MySQL system.
2. Run the **mysql_setup** program from the **bin** subdirectory of the SharePlex product directory.
IMPORTANT! If you installed the SharePlex instance on any port other than the default of 2100, use the **-p** option to specify the port number. For example, in the following command the port number is 9400.

```
C:\users\splex\bin> mysql_setup -p9400
```

Table 3: Setup prompts and responses

Prompt	Response
Enter the MySQL connection string [] :	Enter a connection string that connects to the MySQL database. Do not use a DSN. If you are replicating data larger than 500 MB to MySQL Aurora on Amazon RDS, include the MySQL parameter max_allowed_packet in the connection string and set its value to the maximum size of the data. See the example. Example connection string DRIVER=/usr/lib64/libmyodbc5.so;socket=/var/lib/mysql/mysql.sock;character-set-server=utf8;collation-server=utf8_general_ci;max_allowed_packet=2G;wait_timeout= 6000;Server=servername.amazonaws.com
Enter the MySQL Administrator name :	Enter the name of the MySQL Administrator. This user will perform the setup work on the SharePlex account and schema.

Prompt	Response
Enter the password for the Administrator account :	Enter the password of the Administrator.
Enter the replication target database name :	Enter the name of the MySQL database where you want to install the SharePlex objects.
Database name <i>database</i> does not exist. Would you like to create it? [y] :	If this prompt is displayed, the specified database does not exist. Press Enter to have the setup utility create it for you.
Would you like to create a new SharePlex user [y]:	Press Enter to accept the default to create a new SharePlex database user account, or enter n to use an existing account as the SharePlex database user.
Enter the name of the new SharePlex user: Enter the name of the existing SharePlex user:	One of these prompts is displayed depending on whether you elected to create a new user or use an existing user. Enter the name of the SharePlex user.
Enter the password for the SharePlex user :	Enter the password of the SharePlex user account.
Re-enter the password for the SharePlex user :	Enter the SharePlex password again.

A successful setup terminates with a message similar to the following:

```
Completed SharePlex for MySQL database configuration
SharePlex User name: mysql29
Database name: mysql29
Target specification in SharePlex configuration: r.mysql29
```

Database Setup for Oracle

Overview

Use the Database Setup utility for Oracle (**ora_setup**) to establish SharePlex as an Oracle user and create the required SharePlex database objects. This setup utility creates the following:

- A SharePlex account
- Tables and other objects for use by SharePlex and owned by the SharePlex account
- Default connection for the SharePlex user

It is recommended that you review all of the content in this topic before running the setup utility.

Supported databases

Oracle source or target on supported platforms

When to run Oracle Setup

Whether or not to run this utility at the time of SharePlex installation depends on whether the database is a source, intermediary, or target database, and on how you intend to synchronize the data. To view the initial synchronization procedures, see the SharePlex [Administration Guide](#).

System Type	When to run Oracle Setup
Source system	During installation of SharePlex
Intermediary system	An intermediary system is used in a cascading configuration, where SharePlex replicates data to a remote system (or systems) and then sends that data from the intermediary system to the final target. If you intend to configure SharePlex to post data to a database on an intermediary system, and you intend to use a hot backup to establish the data on that system and the target, do not run the Database Setup utility on the intermediary or target systems. You will run it when you perform the initial synchronization procedure.
Target system	Depends on the method that you will use to synchronize the source and target data when you are ready to activate replication: <ul style="list-style-type: none"> ◦ If you intend to use transportable tablespaces or a cold copy (such as export/import, store/restore from tape, FTP), run the Database Setup utility during SharePlex installation. ◦ If you intend to use a hot backup to establish the target data, <i>do not</i> run the Database Setup utility. You will run it when you perform the initial synchronization procedure. <p>NOTE: If you run the Database Setup utility before the backup and recovery, the setup gets overwritten, and you will need to re-run it again after the backup and recovery.</p>

Supported Oracle Connections

The setup utility can configure any of the following connections for the SharePlex user to use when connecting to the database.

Database type	Connection
Database with or without ASM	Bequeath
Database with or without ASM	TNS alias (A TNS login is specified for both the database and the ASM instance.)
PDB with ASM	TNS alias for the PDB and either TNS or bequeath for the ASM instance.
Amazon RDS	TNS alias

Required run privileges

The user who runs the Database Setup utility must have the following privileges:

Amazon RDS source or target

The user who runs the setup utility must be the master user that was created when the Oracle RDS instance was created. You are prompted for this user during the setup.

Non-multitenant (standard) database

The user who runs the setup utility must have DBA privileges, but if support for TDE is required, then this user must have SYSDBA privileges.

Multitenant database

The user who runs the setup utility should have SYSDBA privileges (recommended), but at minimum the user should be a DBA user with privileges for **sys.users\$** and **sys.enc\$**. The minimum following grants are required for the SharePlex user:

```
create user c##sp_admin identified by sp_admin;
grant dba to c##sp_admin container=ALL;
grant select on sys.user$ to c##sp_admin with grant option container=ALL;
```

If TDE support is required for the CDB, then the following *additional* privilege is required:

```
grant select on sys.enc$ to c##sp_admin with grant option container=ALL;
```

Privileges granted to SharePlex

The Database Setup utility grants to the SharePlex database user the following privileges.

Privilege granted	Description
DBA role	The Database Setup utility grants DBA role and unlimited resource privileges, tablespace privileges, and read privileges to the redo logs.
Default	By default this profile has the unlimited resource privileges originally assigned by Oracle.

Privilege granted	Description
Oracle profile	
Grants	<p>The following grants are issued to SharePlex:</p> <ul style="list-style-type: none"> To access the data dictionary (outside the DBA roles) if <code>O7_DICTIONARY_ACCESSIBILITY</code> is set to <code>FALSE</code>: grant select any dictionary to <i>SharePlexUser</i>; To replicate DDL: grant select any table to <i>SharePlexUser</i> with admin option; grant create any view to <i>SharePlexUser</i> with admin option;

Requirements

- Install the database client on the system where you are running Oracle Setup. Consult the Oracle documentation for the appropriate client version to use with the database.
- Run the Database Setup utility for all source and target Oracle instances in the SharePlex replication configuration.
- Within a cluster, run the Database Setup utility on all nodes of the cluster. This ensures that the SharePlex settings in the Windows Registry include the correct `ORACLE_SID`.
- For a consolidated replication topography, or other topology with multiple variable-data directories, run the Database Setup utility for each variable-data directory.
- SharePlex supports local BEQUEATH connections or remote connections using a TNS alias. Be prepared to supply Oracle Setup the needed connection values for whichever connection you want to use. If using TNS, the `tnsnames.ora` file must be configured prior to running setup.
- If the Oracle database is a multitenant container database, run the Database Setup utility for each pluggable database involved in a replication scenario. A SharePlex user and schema objects must exist in each PDB.
- If you run the Database Setup utility when there is an active configuration, the DDL that the setup performs to install or update the SharePlex internal tables will be replicated to the target. To work around this issue, set the `SP_OCT_REPLICATE_ALL_DDL` parameter to `0` before running the utility, then return the parameter to its previous setting after the setup is complete. This parameter takes effect immediately.

SharePlex schema storage requirements

The Database Setup utility for Oracle installs some database objects for use by SharePlex. The storage requirements for these objects should be satisfied before running Oracle Setup. See the following table.

Storage	Description
SharePlex objects tablespace	The Database Setup utility installs some tables into a tablespace of your choosing. All but the <code>SHAREPLEX_LOBMAP</code> table use the default storage settings of the tablespace.

Storage	Description
	<p>The SHAREPLEX_LOBMAP table contains entries for LOBs stored out-of-row. It is created with a 1 MB INITIAL extent, 1 MB NEXT extent, and PCTINCREASE of 10. The MAXEXTENTS is 120, allowing the table to grow to 120 MB.</p> <p>Preferred action: If you enable supplemental logging for primary and unique keys, you can set the SP_OCT_ENABLE_LOBMAP parameter to 0, and nothing will be stored in the SHAREPLEX_LOBMAP table. In this case, you do not have to consider its size growth. It is recommended that you enable supplemental logging for primary and unique keys to maximize the performance of the Read process.</p> <p>Alternate action: The default storage usually is sufficient for SHAREPLEX_LOBMAP, permitting more than 4 million LOB entries. If the Oracle tables to be replicated have numerous LOB columns that are inserted or updated frequently, consider increasing the size the SharePlex tablespace accordingly. Take into account that this table shares the tablespace with other SharePlex tables.</p> <p>If the database uses the cost-based optimizer (CBO) and the tables that SharePlex processes include numerous LOBs, incorporate the SHAREPLEX_LOBMAP table into the analysis schedule.</p> <p>NOTE: A new installation of SharePlex does not change storage parameters from a previous installation.</p>
SharePlex temporary tablespace	<p>The Database Setup utility prompts for a temporary tablespace for SharePlex to use for sorts and other operations, including sorts performed by the compare commands. The default temporary tablespace is the one where the SharePlex objects are installed. If you plan to use the compare commands to compare large tables, especially those without a primary or unique key, specify a dedicated temporary tablespace for SharePlex.</p>
SharePlex index tablespace	<p>The Database Setup utility prompts for a tablespace to store the indexes for the SharePlex tables. The default index tablespace is the one where the SharePlex objects are installed. To minimize I/O contention, specify a different index tablespace from the one where the tables are installed.</p> <p>NOTE: If indexes from a previous version of SharePlex are installed in the SharePlex objects tablespace, you can move them to a different tablespace and then specify that tablespace when you run the setup utility.</p>

Run Database Setup for Oracle

IMPORTANT! The Oracle instance must be open before this procedure is performed.

1. (Unix and Linux only) If you are using multiple variable-data directories, export the environment variable that points to the variable-data directory for the SharePlex instance for which you are running Database Setup.

ksh shell:

```
export SP_SYS_VARDIR=/full_path_of_variable-data_directory
```

csh shell:

```
setenv SP_SYS_VARDIR /full_path_of_variable-data_directory
```

2. Shut down any SharePlex processes that are running, including **sp_cop**.

- Run the Database Setup program from the command prompt of the operating system, using the full path from the SharePlex **bin** subdirectory.

IMPORTANT! On Windows, if you installed SharePlex on any port other than the default of 2100, use the **-p** option to specify the port number. For example, in the following command the port number is 9400:

```
C:\users\splex\bin>ora_setup -p9400
```

- Specify whether the system is a source system, a target system, or both a source and target system in the SharePlex configuration.

NOTE: This prompt only appears the first time that you run setup for this database.

- For connection type, select **Oracle**.
- Refer to the following table for the prompts and responses to configure SharePlex correctly for the desired connection type.

Table 4: Setup prompts and response

Prompt	Response
Will SharePlex install be using a BEQUEATH connection? (Entering 'n' implies a SQL*net connection):	<p>Press Y to use a local BEQUEATH connection, or press N to use a TNS alias connection.</p> <p>NOTE: Press N to use a TNS alias if:</p> <ul style="list-style-type: none"> the database is a multitenant database SharePlex is capturing from, or posting to, a remote database, such as one on Amazon RDS. the database is in a cluster (such as Oracle RAC)
Are you configuring SharePlex for an AWS RDS database?	<p>Press N if you are not configuring SharePlex for an Oracle database on RDS.</p> <p>Press Y if you are configuring SharePlex for an Amazon AWS RDS database.</p>
<p>One of the following prompts is shown:</p> <ul style="list-style-type: none"> If you selected BEQUEATH= Y: Enter the Oracle SID for which SharePlex should be installed: If you selected BEQUEATH = N: Enter the TNS alias for which SharePlex should be installed: 	<p>Non-multitenant database: Accept the default or type the correct SID or TNS alias. On RAC, the TNS alias must be a global alias.</p> <p>Multitenant database: Type the TNS alias of the PDB.</p> <p>Amazon RDS: Type the TNS alias of the RDS database.</p>
One of the following prompts is shown:	Non-multitenant database: Type the name of a

Prompt	Response
<ul style="list-style-type: none"> • If the database is not on RDS: Enter a DBA user for <i>SID</i>: • If the database is on RDS: In order to create the SharePlex tables and user account, we must connect to the RDS database using the RDS Master user. 	<p>database user that has DBA privileges.</p> <p>Multitenant database: Type the name of a common user who has the required privileges to install the account and objects.</p> <p>Amazon RDS database: Type the name of the RDS master user.</p>
<p>One of the following prompts is shown:</p> <ul style="list-style-type: none"> • If the database is not on RDS: Enter password for the DBA account, which will not echo: • If the database is on RDS: Enter the password for the RDS master user, which will not echo. 	<p>Non-multitenant database: Type the password of the DBA user.</p> <p>Multitenant database: Type the password of the common user. Omit the @ and the rest of the connect string. SharePlex constructs the connect string in the proper format.</p> <p>Amazon RDS database: Type the password of the RDS master user.</p>
<p>Current SharePlex user is <i>user</i>. Would you like to create a new SharePlex user?</p>	<p>Press N to update an existing SharePlex account or Y to create a new SharePlex account. Type the credentials when prompted.</p> <p>You are allowed five attempts to type a valid password for an existing SharePlex user. Passwords are obfuscated.</p> <p>IMPORTANT! If there is an active configuration and you changed the SharePlex schema, copy the SharePlex objects from the old schema to the new one to preserve the replication environment.</p>
<p>Do you want to enable replication of tables with TDE?</p>	<p>Press Y to be prompted for the path name of the TDE wallet file. Type the fully qualified path for the TDE wallet file, including the wallet file name.</p> <p>Press N if not replicating TDE tables.</p>
<p>Enter the default tablespace for use by SharePlex:</p>	<p>Press Enter to accept the default or type the name of a different tablespace.</p>
<p>Enter the temporary tablespace for use by Shareplex:</p>	<p>Press Enter to accept the default or type the name of a different tablespace.</p>

Prompt**Response**

Enter the index tablespace for use by SharePlex:

Press **Enter** to accept the default or type the name of a different tablespace.

Will the current setup for sid: *SID* be used as a source (including cases as source for failover or master-master setups)?

Press **Y** if this is a source system or press **N** if this is a target system. **IMPORTANT:** All systems in a master-master configuration (peer-to-peer) and in a high-availability configuration are considered to be source systems due to the bidirectional nature of the replication.

NOTE:

- The following prompts are only shown if the database is a source on ASM.
- If this is an Oracle target, the setup is now complete.

ASM detected. Do you wish to connect to ASM using BEQUEATH connection?

Press **Y** for SharePlex to use a BEQUEATH connection to connect to the ASM instance, or press **N** to use a TNS alias.

IMPORTANT! If the database uses ASM *and* the database TNS alias is configured with a SCAN IP, then you must specify connection through an ASM TNS alias in order for SharePlex to connect to the ASM instance.

The following prompt is displayed if you selected a BEQUEATH connection:

Press **Y** to use the same user and password as the login user, or press **N** to be prompted for a different user and password.

Do you wish to keep connecting using the same user/password?

NOTE: If you selected to use a BEQUEATH connection to connect to ASM, the database setup is complete. Continue to [Note about the tnsnames file](#).

If you selected **N**, you need to supply a TNS alias, and the prompts continue.

Enter the ASM tns alias to be used by SharePlex:

Type the name of the TNS alias.

Enter an ASM admin (has both sysdba and sysasm privilege) username for *alias*:

Type the name of a user with sysasm and sysdba privileges to the ASM instance.

Prompt	Response
Enter user password for user:	Type the password of the user.
<p>SharePlex installs internal objects that include a package to support the SDO_GEOMETRY datatype of the Oracle Spatial and Graph option. If this option is not installed in the database, the following prompt is shown:</p> <p>The SharePlex object that supports replication of SDO_GEOMETRY cannot be installed because the Oracle Spatial and Graph feature is not installed. Do you want to continue with the setup without support for SDO_GEOMETRY? [n]:</p>	Press Y to continue the database setup without support for SDO_GEOMETRY, or press N to terminate ora_setup .

Note about the tnsnames file

When you set up SharePlex for database connection through a TNS alias and ASM connection locally through a BEQUEATH connection (through OS authentication), it is important to set up the **tnsnames.ora** file correctly on each node. Assuming a SharePlex database account exists on the primary node, SharePlex will always connect to the primary ASM_SID automatically because it was provided when SharePlex was installed. However, upon failover, SharePlex must query the local **v\$asm_client** view to get the correct ASM_SID of the failover instance. Therefore, ensure that the IP address of a given node is always listed first in the local **tnsnames.ora** file on that node.

Database Setup for PostgreSQL

Overview

Run the Database Setup utility for PostgreSQL (**pg_setup**) to establish a user account, schema, and tables for use by SharePlex.

Supported databases

All implementations of the PostgreSQL open-source database on supported platforms

Guidelines for use

- Run the Database Setup utility on all target PostgreSQL instances in the SharePlex replication configuration.
- Within a server cluster, run the Database Setup utility on the node to which the shared disk that contains the SharePlex variable-data directory is mounted.
- For consolidated replication, run the Database Setup utility for each variable-data directory.
- On **Windows**, you must supply a DSN (data source name) as the connection method.
- On **Linux**, you can provide a connection string or a DSN (data source name). Note the following when deciding on the connection method on Linux:

Connection Type on Linux	What to do
--------------------------	------------

Connection string	You do not need to provide a user, password, or default database in a connection string. The connection string must have the port, server and driver defined. This is an example from EDB Postgres Plus Advanced Server; your connection string will probably be different based on your own database setup: Port=5444;server=localhost;driver=/u01/PostgresPlus/connectors/odbc/lib/edb-odbc.so;database=edb;
-------------------	---

DSN	If you have a DSN defined, and you want to use it for the SharePlex connection, copy or link the ODBC files in which that DSN is defined (odbc.ini and odbcinst.ini) to the odbc subdirectory of the SharePlex variable-data directory. This prevents connection errors when the SharePlex processes connect to the database. If you do not have a DSN defined but want to use one, you can create it in the template files provided in the odbc subdirectory.
-----	---

Required privileges

Review the following requirements to ensure that the setup succeeds.

- The Database Setup utility must be run as a PostgreSQL Administrator in order to grant SharePlex the required privileges to operate on the database and to create the SharePlex database account.
- (Symfaware only) If the person who is running the setup utility is not a Fujitsu Enterprise Postgres owner, set the environment variable LD_LIBRARY_PATH to include the path to the **lib** subdirectory in the Fujitsu Enterprise Postgres installation directory. The LD_LIBRARY_PATH is set in the **.bash_profile** file of the Fujitsu Enterprise Postgres owner.

Example:

```
export LD_LIBRARY_PATH= /opt/symfoserver64/lib:$LD_LIBRARY_PATH
```

If you do not set this path, the following error occurs:

```
symbol lookup error: /opt/fsepv95client64/odbc/lib/psqlodbca.so: undefined  
symbol: PQconnectdbParams
```

- Cloud installations:
 - Common restrictions on privileges in cloud-hosted database services make it difficult for the setup utility to succeed in every possible type of scenario. To ensure that the database setup succeeds, *only* use the setup utility for the following purposes: To do a *first-time* database setup with a *new* SharePlex user, or, to *modify* an existing SharePlex user that either owns the database or has access to it.
 - On Amazon RDS, you might need to grant usage/privileges on the target schema and tables to the SharePlex user, as shown in the following example:
Log in as the schema owner, then issue the following commands:
grant usage on schema *schema_name* to *user_name*;
grant all privileges on all tables in schema *schema_name* to *user_name*;

Run Database Setup for PostgreSQL

1. Shut down any running SharePlex processes and **sp_cop** on the target system.
2. Run the **pg_setup** program from the **bin** subdirectory of the SharePlex product directory.
IMPORTANT! If you installed the SharePlex instance on any port other than the default of 2100, use the **-p** option to specify the port number. For example, in the following command the port number is 9400.
\$ /users/splex/bin> pg_setup -p9400

Table 5: Setup prompts and responses

Prompt	Response
<p>(Linux)</p> <p>Enter the PostgreSQL DSN name or connection string [] :</p>	Enter a DSN or a connection string if the system is Linux, or enter a DSN if the system is Windows.
<p>(Windows)</p> <p>Enter the PostgreSQL DSN name [] :</p>	
Enter the PostgreSQL Administrator name :	Enter the name of the PostgreSQL Administrator. This user will perform the work on the SharePlex account.
Enter the password for the Administrator account :	Enter the password of the Administrator.
Enter the replication target database name:	Enter the name of the database that you want to contain the SharePlex tables and other objects for use by SharePlex. You can enter the name of a new or existing database.
Database name <i>database</i> does not exist. Would you like to create it? [y] :	If this prompt is displayed, the specified database does not exist. Press Enter to have the setup utility create it for you.

Prompt	Response
Would you like to create a new SharePlex user [y]:	Press Enter to accept the default to create a new SharePlex database user account and schema of the same name in the specified database, or enter n to use an existing SharePlex account.
Enter the name of the new SharePlex user: Enter the name of the existing SharePlex user:	One of these prompts is displayed depending on whether you elected to create a new user or use an existing user. Enter the name of the SharePlex user.
Enter the password of the SharePlex user :	Enter the password of the SharePlex user account.
Re-enter the password for the SharePlex user :	This prompt is only shown if you created a new user. Enter the SharePlex password again.

A successful setup terminates with a message similar to the following:

```
Completed SharePlex for PostgreSQL database configuration
SharePlex User name: splex
Database name: ndb5
Target specification in SharePlex configuration: r.ndb5
```

Database Setup for SQL Server

Overview

Run the Database Setup utility for SQL Server (**mss_setup**) on a Microsoft SQL Server system to establish SharePlex as a SQL Server database user. This utility creates the following:

- A SharePlex user account with full DBA privileges
- Tables and indexes for use by SharePlex and owned by the SharePlex user in a database of your choosing
- A default database connection.

Supported databases

Microsoft SQL Server on Windows

Guidelines for use

- A *system* DSN (data source name) must exist for the SQL Server database. SharePlex Post uses the DSN to connect to the database through ODBC.
- Run the Database Setup utility on all SQL Server instances in the SharePlex replication configuration.
- Within a cluster, run the Database Setup utility on the node to which the shared disk that contains the variable-data directory is mounted.
- For consolidated replication, run the Database Setup utility for each variable-data directory.

Required privileges

Review the following requirements to ensure that the setup succeeds.

- The Database Setup utility must be run as a SQL Server System Administrator in order to grant SharePlex the required privileges to operate on the database and to create the SharePlex database account and objects.
- (Cloud installations) Common restrictions on privileges in cloud-hosted database services make it difficult for the setup utility to succeed in every possible type of scenario. To ensure that the database setup succeeds, *only* use the setup utility for the following purposes: To do a *first-time* database setup with a *new* SharePlex user, or, to *modify* an existing SharePlex user that either owns the database or has access to it.

Run Database Setup for SQL Server

1. Shut down any running SharePlex processes and **sp_cop** on the SQL Server system.
2. Run the **mss_setup** program from the **bin** subdirectory of the SharePlex product directory.
IMPORTANT! If you installed the SharePlex instance on any port other than the default of 2100, use the **-p** option to specify the port number. For example, in the following command the port number is 9400.

```
C:\users\splex\bin> mss_setup -p9400
```

Table 6: Setup prompts and responses

Prompt	Response
Enter the Microsoft SQL Server DSN name [] : :	Enter the data source name (DSN) that connects to SQL Server. Make certain the DSN is a system DSN, not a user DSN.
Enter the Microsoft SQL Server Administrator name : :	Enter the name of the SQL Server Administrator. This user will perform the setup work on the SharePlex account and schema.
Enter the password for the Administrator account : :	Enter the password of the Administrator.

Prompt	Response
Enter the database name:	Enter the name of the database where you want to install the SharePlex objects.
Database name <i>database</i> does not exist. Would you like to create it? [y] :	If this prompt is displayed, the specified database does not exist. Press Enter to have the setup utility create it for you.
Would you like to create a new SharePlex login [y]:	Press Enter to accept the default to create a new SharePlex database user account, or enter n to use an existing account as the SharePlex database user.
Enter the name of the existing SharePlex login: Enter the name of the new SharePlex login:	One of these prompts is displayed depending on whether you elected to create a new user or use an existing user. Enter the name of the SharePlex user.
Enter the password for <i>login</i> :	Enter the password of the SharePlex user account.
Re-enter the password for <i>login</i> :	Enter the SharePlex password again.
Will this database be used as a source?	Accept the default of N if the database will only be a target. Enter Y if this database will be a source database for SharePlex. A response of Y prompts the setup to prepare the database for data capture and installs the SharePlex account and objects.

A successful setup terminates with a message similar to the following:

```
Completed SharePlex for Microsoft SQL Server database configuration
```

```
SharePlex User name: splx
```

```
Database name: db1
```

```
Target specification in SharePlex configuration: r.db1
```

Database Setup for Tibero

Overview

Run the Database Setup utility for Tibero (**tb_setup**) on a Tibero system to establish SharePlex as a Tibero database user. This utility creates the following:

- A SharePlex user account with full DBA privileges
- Tables and indexes for use by SharePlex and owned by the SharePlex user in a database of your choosing
- A default database connection.

Supported databases

For supported Linux platforms and versions, see the SharePlex Release Notes.

Guidelines for use

- Run the setup utility on all Tibero instances in the SharePlex replication configuration.
- Within a cluster, run the setup utility on the node to which the shared disk that contains the SharePlex variable-data directory is mounted.
- For consolidated replication, run the setup utility for each variable-data directory.

Required privileges

The setup utility must be run as a Tibero Administrator that retains all of that user's default privileges. The Administrator user is able to grant SharePlex the required privileges to operate on the database and to create the SharePlex database account and objects.

Run Database Setup for Tibero

1. Shut down any running SharePlex processes and **sp_cop** on the Tibero system.
2. Run the **tb_setup** program from the **bin** subdirectory of the SharePlex product directory.
IMPORTANT! If you installed the SharePlex instance on any port other than the default of 2100, use the **-p** option to specify the port number. For example, in the following command the port number is 9400.

```
C:\users\splex\bin> tb_setup -p9400
```


Table 7: Setup prompts and responses

Prompt	Response
Enter the Tibero DSN name or connection string [] :	Enter a connection string that connects to the Tibero database.
Enter the Tibero Administrator name :	Enter the name of the Tibero Administrator. This user will perform the setup work on the SharePlex account and schema.
Enter the password for the Administrator account :	Enter the password of the Administrator.
NOTE:User 'shareplex' does not have the superuser role.	
Enter the replication target database name :	Enter the name of the Tibero database where you want to install the SharePlex objects.
Database name <i>database</i> does not exist. Would you like to create it? [y] :	If this prompt is displayed, the specified database does not exist. Press Enter to have the setup utility create it for you.
Would you like to create a new SharePlex user [y]:	Press Enter to accept the default to create a new SharePlex database user account, or enter n to use an existing account as the SharePlex database user.
Enter the name of the new SharePlex user:	One of these prompts is displayed depending on whether you elected to create a new user or use an existing user. Enter the name of the SharePlex user.
Enter the name of the existing SharePlex user:	
Enter the password for the SharePlex user :	Enter the password of the SharePlex user account.
Re-enter the password for the SharePlex user :	Enter the SharePlex password again.

A successful setup terminates with a message similar to the following:

```
Completed SharePlex for Tibero database configuration
SharePlex User name: splx
Database name: tbl29
Target specification in SharePlex configuration: r.tbl29
```

General SharePlex Utilities

The SharePlex utilities help you configure, test, and manage the SharePlex environment.

Contents

- [SharePlex license utilities](#)
- [Install the SharePlex service](#)

SharePlex license utilities

Description

Use the SharePlex license utilities to view and add license keys to hosts in the SharePlex replication environment. Each installation of SharePlex requires a valid license key. There are three types of SharePlex license keys:

- Temporary license keys (demonstration licenses)
- Permanent license keys
- Site license keys

The current license model for SharePlex is to license for a specific host, which depending on edition is licensed by core(s) or socket(s) and specific message repository (i.e. database, JMS/text files) etc. Specifics of license terms should be obtained from your account manager.

If you do not have a valid license key, you may obtain one from Quest Technical Support or your Quest sales representative. Use the appropriate procedure in this documentation to obtain the necessary information to support your license request.

License utilities on Unix and Linux

On Unix and Linux systems, separate utilities are used to:

- [Gather information to request a license](#)
- [Add a license key](#)
- [View a license key](#)

Gather information to request a license

Use the **splex_uname** utility to collect information about the local host that is required for Quest to generate a license key. Run this utility from the **util** sub-directory of the SharePlex product directory.

```
$ /proddir/util/splex_uname
```

The information is similar to the following example:

```
Local system info:
Host Name           = irvspxu09
Host ID             = 2198894273 (831076C1 HEX)
Operating System    = SunOS
Machine model       = SUNW,UltraAX-i2
CPU Type            = sparc9
CPU Speed           = 500
Number of CPUs      = 1
```

Add a license key

Use the **splex_add_key** utility to add a license key to a machine during the installation of SharePlex or afterward to convert from one type of license to another.

You can use the **splex_add_key** utility as follows:

- You can use **splex_add_key** on the primary node of a cluster to install licenses for all secondary nodes in the cluster, because they all share one variable-data directory.
- You cannot use **splex_add_key** to add licenses for non-clustered machines from one machine. It must be run on each non-clustered replication system so that the license information is stored in the variable-data directory on each system.

To run **splex_add_key**

1. Log on to the system as the SharePlex Administrator.
2. Run **sp_ctrl** on the machine where you want to install a license key.
3. If SharePlex is running, shut it down.

```
sp_ctrl> shutdown
```
4. Run **splex_add_key** from the **install** sub-directory of the SharePlex product directory.

```
$ lproddir/install/splex_add_key
```
5. Choose the appropriate option for how to install the key:

```
SharePlex License Utility
1) Read license key from file
2) Enter license key manually
3) Add license key for alternate host
q) Quit License Utility
Enter option:
```

NOTE: Use option **3** when installing on secondary nodes in a cluster. To install SharePlex in a cluster, see [Preinstallation instructions for Oracle cluster](#) .

6. If you are entering the key manually, type the license key exactly as you received it from Quest, including the **SPLEXKEY=** string if one prepends the key, any spaces, and any capitalization. Press **Enter** when finished typing the key.
7. Enter the **SiteMessage** text string provided by Quest with the license key.
NOTE: If you selected option **3** to install in a cluster, the utility continually prompts for another machine ID to license, until you exit the license utility.

8. Enter **q** to exit the utility.
9. Start SharePlex when you are ready for replication to resume.

View a license key

Use the `splex_get_key` utility to view the SharePlex license key, license type, and expiration date on a machine. Run this utility from the `install` sub-directory of the SharePlex product directory.

```
$ /proddir/install/splex_get_key
```

The information is similar to the following example:

```
$ /splex/proddir/install/splex_get_key
Customer Name = Quest
License Key = EANLCF3P56F8M1JB1K6RBJCBB6TBEP0
Product Name = SharePlex Oracle
License type = demo key
Expired on = 01/01/2008 12:00 AM
```

License utility on Windows

On Windows platforms, all licensing functions are performed through the **SpUtils** utility, including viewing the machine ID for a system. Launch the **SpUtils** utility from the shortcut on the Windows desktop.

Gather information to request a license

Select the **Information** tab to view information about the local host that is required for Quest to generate a license key. In a cluster, run **SpUtils** on each node.

Add a license key

Use **SpUtils** to add a license key to a machine during the installation of SharePlex or afterward to convert from one type of license to another. In a cluster, run **SpUtils** on the primary node.

1. Log on to Windows as the SharePlex Administrator.
2. If SharePlex is running, do the following:
 - a. Select the **SharePlex Services** tab.
 - b. Select the port number of the SharePlex instance for which you want to add a license key.
 - c. Click **Stop**.
 - d. Wait for **Current State** to display a message that the service stopped.
3. Select the **License Keys** tab.
4. Select the SharePlex port number from the **Port** list.
5. Click **Add License**, then type or paste the information exactly as you received it from Quest, as follows:
 - a. **License Key**: The license key, including any spaces. The key is case-sensitive.
 - b. **Customer Name**: The text string that was included with the license. The name is case-sensitive.
6. (Windows Cluster) To add a license for another node in a cluster, select the **Add Alternate Host Id** check

box, then select the machine ID from the CPU ID list box. Repeat this step to add a license to all secondary nodes in the cluster.

7. Click **OK to close the utility**.
8. To start SharePlex, do the following:
 - a. Select the **SharePlex Services** tab.
 - b. Select the port number of the SharePlex instance that you licensed.
 - c. Click **Start**.
 - d. Wait for **Current State** to display a message that the service started.

View a license key

Select the **License Key** tab, then select the port number (if running multiple SharePlex instances) for which you want to view license information.

- Under **License Key**, view the actual license key and SiteMessage text string.
- Under **Status**, view the type of SharePlex installation, the license type, and the expiration date.

In a cluster, run **SpUtils** on each node to view the local license.

Install the SharePlex service

SharePlex runs as a service on the Windows platform. The service name is **SharePlex port_number**, where *port_number* is the port number associated with that SharePlex instance.

SharePlex is not installed as a Windows service during the initial installation. You must add and start the service through the **SpUtils** utility.

To add and start SharePlex as a service

1. Run the **SpUtils** utility from the SharePlex entry in the Programs menu.
2. Select the **Service** tab.
3. Select the SharePlex port number for which you are installing the service.
4. Click **Install**. (A "Service Stopped" message indicates that the service is installed.)
5. (Optional) Click **Start** to start the service.

The service is installed in auto-startup mode (start when the system starts) so that replication begins as soon as possible. To change startup status, use the **Services** applet of the **Administrative Tools** in the Windows Control Panel.

Advanced SharePlex installer options

The use of additional command line options when installing SharePlex is usually not necessary. These options are typically employed when working with Support to resolve specific issues.

The installer command line options and their descriptions follow:

USAGE

```
tpm [<options>] [ [<package> | <location>] ... ]
```

OPTIONS

```
-v, --verbose Turns verbose mode on
```

```
-h, -?, --help Prints out this message
```

```
--debug Starts the interactive debugger
```

```
--info Print information about installed
```

packages

```
--install Perform product installation
```

```
--remove Perform product deinstallation
```

```
--commit Commit last installation
```

```
--revert Revert last installation
```

```
-t, --tmp <directory> Temporary directory location
```

```
-d, --directory <directory> Working directory
```

```
-f, --force Unconditionally update existing files
```

```
--no-cleanup Do not perform cleanup on failure
```

```
--nocleanup Same as --no-cleanup, for compatibility
```

--list List the content of the archive
--extract Extract the archive into the current directory
-r, --responses <yaml file> Use the responses from a specified file
-D, --defaults Accept default answers
-l, --log Leave the installation log file

DESCRIPTION

Provides package management facilities. Packages can be installed, removed, reverted or committed. The utility also figures out its role based on the command name of its invocation path. For example, "tpm-install" is treated as "tpm --install", "tpm-remove" as "tpm --remove", etc.

It can also be invoked as part of a self extracting package invocation, in which case it is treated as "tpm --install".

Note: All command line options for the .tpm file are preceded by two dashes.

Install SharePlex as Root

You can install SharePlex as a root user. When you install as a root user, the installer prompts you to select whether or not to create the SharePlex user groups. When the installer creates the groups, it adds the SharePlex Administrator user to the **spadmin** group. For more information about these groups, see [Assign SharePlex users to security groups](#) on page 46.

In a cluster, the installer adds the SharePlex groups to the primary node, but you must add them to the other nodes yourself.

Additionally, see [Network checklist](#) on page 23 for instructions on adding the groups to a nameserver.

To install as root

1. Log in to the system as a root user.
2. Copy the SharePlex installer file to a temporary directory where you have write permissions. The installer file has a naming format of:

```
SharePlex-release#-DBVersion-platform.tpm
```

3. Change the permissions of the file as follows:

```
# chmod 555 SharePlex-release#-DBVersion-platform.tpm
```

4. Run the installer as directed in [Install SharePlex on Linux and UNIX](#) on page 36.

Run the installer in unattended mode

SharePlex can be installed unattended through the use of a response file. This installation method speeds the installation of multiple SharePlex instances. The file supplies responses to the standard installer prompts, while providing on screen status information.

NOTE: When running in unattended mode, the installation process does not call the system password utility. If you create a new SharePlex user during the installation, that user will remain locked until the password is set manually.

Response files that you can edit are located in the **install** subdirectory of the SharePlex product (installation) directory:

```
/productdir/install
```

Enter responses in the file

IMPORTANT! The response file contains two sections. **Only the top section is user configurable.** Do not edit the bottom section. The bottom section begins with the line "Do not change settings that appear below."

Edit the top section of the response file to provide the responses for the installation. **Only edit the values to the right of the colon,** and make certain there is a space between the colon and the response.

The following example is for non-root installation:

```
# To install SharePlex with the unattended option please
# modify the settings below. You may safely modify only the values
# to the right of the colon, and the colon must be immediately
# followed by a space. Editing the values to the left of the colon
# may impact the unattended install causing the process to become
# interactive.
#
SharePlex Admin group: spadmin
```

```
product directory location: /home/splex/proddir

variable data directory location: /home/splex/var_dir

# not required for Open Target installations # ORACLE_SID that corresponds to
this installation: oracledb

# not required for Open Target installations# ORACLE_HOME directory that
corresponds to this ORACLE_SID: /home/oracle/products/version

TCP/IP port number for SharePlex communications: 2100

the License key: XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX

the customer name associated with this license key: SharePlex_Trial
```

To run the response file

From the command shell of the operating system, run the **.tpm** installation file with the **-r** option followed by the full path to the response file.

```
# ./installation_file -r /users/shareplex/product.rsp
```

Next Steps

For additional setup requirements, see the **Next Steps** section in one of the following depending on the platform:

[Install SharePlex on Linux/Unix for Oracle Database on page 36](#)

[Install SharePlex on Linux/Unix for Open Target Databases on page 39](#)

SharePlex installed items

The following describes the items that are installed during the SharePlex setup process.

SharePlex Directories

SharePlex uses two main directories:

The product directory: This is the SharePlex installation directory, where the SharePlex programs and libraries are stored.

The variable-data directory: This is the SharePlex working directory, where the queue files, log files and other components that comprise the current replication environment are stored.

NOTE: These directories are often referred to as *productdir* and *vardir*, respectively.

Do not remove, rename or edit any files or directories installed by SharePlex. Some directories contain hidden files that are essential for replication. Some files appear empty but must exist under their original names because they are referenced by one or more SharePlex processes. Some items in the directories are for use only under the supervision of Quest Technical Support.

Programs meant for general use in a production environment are documented in the published SharePlex documentation. If you do not find documentation for a program in a SharePlex directory, do not attempt to run it. Contact Quest Technical Support first.

The directory structure and files within the two main SharePlex directories differs slightly between the UNIX and Windows platforms. Files and directories can vary from version to version of SharePlex, but the basic structure appears as follows.

SharePlex product directory

Sub-directory	Contents
BACKUP	Uninstall information
bin	SharePlex executable files
config	Internally used content.

Sub-directory	Contents
data	Default parameter settings
doc	Catalog of exception messages
install	(Unix and Linux only) Scripts related to installation, licensing and upgrades
lib	SharePlex shared libraries
log	SharePlex log files
mks_oe	Runtime installation files for third-party software used by SharePlex.
util	SharePlex utilities
.app-modules	(Unix and Linux only) Hidden internal directory that contains raw executables. Do not use the contents of this directory to launch processes.
.meta-inf	(Unix and Linux only) Hidden internal directory that contains meta information used during the installation process.

SharePlex variable-data directory

Sub-directory	Contents
config	Configuration files for this installation of SharePlex.
data	Status Database, configuration activation information, user-defined parameter settings, and other user-defined files that direct replication activities.
db	Configuration internal database for each activation of a configuration file.
downgrd	Information about SharePlex targets that are a lower version than the source.
dump	Core files (if a process fails)
log	SharePlex log files
rim	Queue files (working data files)
save	Information about active and inactive configurations.
state	Information about the current state of SharePlex when a configuration is active, such as the object and sequence caches.
temp	Used by the copy and append features and other SharePlex sync-related processes.
oos	Stores the transactions that contain out-of-sync operations when the SP_OPO_SAVE_OOS_TRANSACTION parameter is enabled.

SharePlex Installed Objects

Much of the replication process is controlled and tracked through a series of internal objects that are installed into the source or target database during the installation of SharePlex. They are essential for SharePlex to operate, so do not alter them in any way.

NOTE: Not all objects are used for all databases. Most are used for Oracle databases. If you do not see an object in your database, it is not relevant to the database, or the information is stored internally within the SharePlex configuration. If you see an object that is in your database but not in this list, it is not being used in the current release.

Table	Object type	Description
DEMO_SRC	Table	Used as the source table for the SharePlex demonstrations.
DEMO_DEST	Table	Used as the target table for the SharePlex demonstrations.
SHAREPLEX_ACTID	Table	Used by Capture to checkpoint its state.
SHAREPLEX_ANALYZE	Table	Used by the analyze command.
SHAREPLEX_CHANGE_OBJECT	Table	Used by users to stop and resume replication for an object.
SHAREPLEX_COMMAND	Table	Used for the flush , abort and purge commands.
SHAREPLEX_CONFIG	Table	Used by the activation and Capture processes to mark the start of a new activation.
SHAREPLEX_DATA	Table	Used by the SharePlex wallet for Oracle TDE replication.
SHAREPLEX_DATAEQUATOR	Table	Used by the compare and repair commands and the Post process to synchronize their operations.
SHAREPLEX_DATAEQUATOR_INSERT_TEMP	Table	Used as a temporary table by the compare and repair commands.
SHAREPLEX_DATAEQUATOR_UPDATE_TEMP	Table	Used as a temporary table by the compare and repair commands.
SHAREPLEX_DATAEQUATOR_DELETE_TEMP	Table	Used as a temporary table by the compare and repair commands.
SHAREPLEX_DDL_CONTROL	Table	Used to refine control of DDL that is enabled for replication by the SP_OCT_REPLICATE_ALL_DDL parameter.
SHAREPLEX_JOBID	Sequence	Used by the sp_cnc process and the compare , repair , and copy commands to provide a unique job ID.
SHAREPLEX_JOBS	Table	Used by the sp_cnc process and the compare , repair , and copy commands to store information about a job.
SHAREPLEX_JOB_STATS	Table	Used by the sp_cnc process and the compare , repair , and copy commands to store information about a job.
SHAREPLEX_JOBS_CONFIG	Table	Used by the disable jobs and enable jobs commands.
SHAREPLEX_LOB_CACHE	Table	Used by the Capture process when processing VARRAYs stored as LOB.

Table	Object type	Description
SHAREPLEX_LOBMAP	Table	Used by the Capture process to map LOBIDs and rows when a table with LOB columns does not have PK/UK logging enabled.
SHAREPLEX_LOGLIST	Table	Used by the Capture process to track inactive RAC instances.
SHAREPLEX_MARKER	Table	Used by the Read process when PK/UK logging is not enabled.
SHAREPLEX_OBJMAP	Table	Used by the activation and Capture processes to define the objects in replication.
SHAREPLEX_PARTITION_CACHE	Table	Used by the Capture process to map Oracle partition IDs to tables in replication.
SHAREPLEX_SYNC_MARKER	Table	Used by the copy command and the Read and Post processes to sync their operations.
SHAREPLEX_TRANS or SHAREPLEX_OPEN_TRANS	Table	Used by the Post process to store checkpoints and to mark transactions that were applied in a master-to-master configuration.

MKS Toolkit

On the Windows platform, SharePlex installs the MKS Toolkit® operating environment from Parametric Technology Corporation (PTC), formerly known as Mortice Kern Systems NuTCRACKER. This software enables SharePlex to be ported to all supported platforms in a uniform manner.

SharePlex users have no interaction with the MKS Toolkit software, other than to ensure that its service is running when the SharePlex service is running. The MKS Toolkit is installed in automatic startup mode by default when it is installed during SharePlex installation.

The default folder for MKS Toolkit is C:\Program Files\MKS Toolkit, but a different folder might have been selected during installation. The MKS Toolkit files must not be moved after they are installed.

SharePlex Registry Entries

On Windows systems, SharePlex installs the following Registry entries.

- **\HKEY_LOCAL_MACHINE\Software\Wow6432node\Quest Software\SharePlex:** This contains the SharePlex environment information and is where any user-created variables must be entered, such as SP_SYS_VARDIR. Each instance of SharePlex is defined in the Registry as a port number, and each has its own set of environment variables.
- **\HKEY_LOCAL_MACHINE\Software\Wow6432node\Mortice Kern Systems:** This contains Registry entries for the MKS Toolkit® operating environment.
- **\HKEY_LOCAL_MACHINE\Software\Wow6432node>Datafocus:** This contains Registry entries for the MKS Toolkit® operating environment.

About us

We are more than just a name

We are on a quest to make your information technology work harder for you. That is why we build community-driven software solutions that help you spend less time on IT administration and more time on business innovation. We help you modernize your data center, get you to the cloud quicker and provide the expertise, security and accessibility you need to grow your data-driven business. Combined with Quest's invitation to the global community to be a part of its innovation, and our firm commitment to ensuring customer satisfaction, we continue to deliver solutions that have a real impact on our customers today and leave a legacy we are proud of. We are challenging the status quo by transforming into a new software company. And as your partner, we work tirelessly to make sure your information technology is designed for you and by you. This is our mission, and we are in this together. Welcome to a new Quest. You are invited to Join the Innovation™.

Our brand, our vision. Together.

Our logo reflects our story: innovation, community and support. An important part of this story begins with the letter Q. It is a perfect circle, representing our commitment to technological precision and strength. The space in the Q itself symbolizes our need to add the missing piece — you — to the community, to the new Quest.

Contacting Quest

For sales or other inquiries, visit www.quest.com/contact.

Technical support resources

Technical support is available to Quest customers with a valid maintenance contract and customers who have trial versions. You can access the Quest Support Portal at <https://support.quest.com>.

The Support Portal provides self-help tools you can use to solve problems quickly and independently, 24 hours a day, 365 days a year. The Support Portal enables you to:

- [Submit and manage a Service Request](#)
- [View Knowledge Base articles](#)
- [Sign up for product notifications](#)
- [Download software and technical documentation](#)
- [View how-to-videos](#)
- [Engage in community discussions](#)
- [Chat with support engineers online](#)
- [View services to assist you with your product](#)