



IBM VisualAge TeamConnection Enterprise Server

User's Guide

Version 3.0



IBM VisualAge TeamConnection Enterprise Server

User's Guide

Version 3.0

Fifth Edition (October 1998)

Note

Before using this document, read the general information under "Notices" on page xiii.

This edition applies to Version 3.0 of the licensed program IBM TeamConnection and to all subsequent releases and modifications until otherwise indicated in new editions. Make sure you are using the correct edition for the level of the product.

Order publications by phone or fax. The IBM Software Manufacturing Company takes publication orders between 8:30 a.m. and 7:00 p.m. eastern standard time (EST). The phone number is (800) 879-2755. The fax number is (800) 284-4721.

You can also order publications through your IBM representative or the IBM branch office serving your locality. Publications are not stocked at the address below.

A form for comments appears at the back of this publication. If the form has been removed, address your comments to:

IBM Corporation
Attn: Information Development
Department T99B/Building 062
P.O. Box 12195
Research Triangle Park, NC, USA 27709-2195

You can fax comments to (919) 254-0206.

If you have comments about the product, address them to:

IBM Corporation
Attn: Department TH0/Building 062
P.O. Box 12195
Research Triangle Park, NC, USA 27709-2195

You can fax comments to (919) 254-4914.

When you send information to IBM, you grant IBM a nonexclusive right to use or distribute the information in any way it believes appropriate without incurring any obligation to you.

© **Copyright International Business Machines Corporation 1992, 1995, 1996, 1997, 1998. All rights reserved.**
Note to U.S. Government Users — Documentation related to restricted rights — Use, duplication or disclosure is subject to restrictions set forth in GSA ADP Schedule Contract with IBM Corp.

Contents

Figures	xi
Notices	xiii
Trademarks	xv
About this book	xvii
How this book is organized	xvii
Conventions	xvii
Tell us what you think	xviii

Part 1. Introducing TeamConnection 1

Chapter 1. An introduction to TeamConnection	3
TeamConnection definitions	4
TeamConnection's client/server architecture	4
TeamConnection database	5
Interfaces	5
Families	6
Users and host lists	6
Parts	6
Components	7
Releases	8
Work areas	8
Drivers	9
Defects and features	10
Processes	10
Build	11
Packaging	12
Roles people play	12

Part 2. Developing a product using TeamConnection 15

Chapter 2. Getting familiar with the TeamConnection client interfaces	17
Using the GUI	17
Starting the GUI	18
Stopping the GUI	19
Performing tasks with the GUI	19
Using the Settings notebook	20
Online help information	21
Using the command line interface	22
Using the TeamConnection web client	23
Chapter 3. The basics of using TeamConnection	25
Laying the groundwork	25

Authority to perform tasks	26
Finding objects within TeamConnection	27
Finding parts	27
Using work areas	28
Naming your work areas	29
Creating parts.	29
Naming your parts	30
Preparing to build your parts	30
Working with parts	31
Working in serial or concurrent development mode	31
Working with common parts	32
Getting parts from TeamConnection	33
Checking parts in to TeamConnection	34
Finding different versions of TeamConnection objects	35
Versioning releases	35
Versioning work areas	36
Versioning drivers	37
Versioning parts	38
Working with defects and features	38
Testing and verifying part changes	39
Chapter 4. The states of TeamConnection objects.	41
Defects and features	41
The states of work areas	45
The states of drivers	47
Verification and test records	49
Chapter 5. Working with no component or release processes	51
Working in serial development	51
Accepting a defect	52
Creating a work area	53
Checking out a part	54
Searching for a part.	55
Checking in a part	58
Verifying and testing part updates	60
Freezing the work area.	65
Refreshing the work area	66
Building the application	67
Integrating the work area	68
Closing a defect	69
Working in concurrent development	70
Refreshing the work area from the driver.	71
Integrating the work area	72
Reconciling differences.	73
Chapter 6. Working with component and release processes	77
Moving through design, size, and review	78
Changing defect ownership	79
Accepting a defect	80

Approving the fix	81
Checking out a part	82
Checking in the changes	83
Freezing the work area.	85
Building the application	86
Accepting fix records	87
Integrating changed parts into a release	88
Adding a driver member	88
Reconciling the differences	89
Refreshing the driver	91
Building the driver	92
Restricting the driver	93
Integrating the parts.	94
Completing the driver	95
Testing the built application	95
Using a configured process	96
Retrieving a past version of a part	97

Part 3. Using TeamConnection Notes Integrated Databases 101

Chapter 7. Introduction to TeamConnection Integrated Notes Databases	103
Getting started	103
Prerequisites and dependencies	104
Using TeamConnection with Lotus Notes	104
Sources of user information	104
Database types	105
Forms and subforms	106
Views	107
Reviews	109
Document archiving.	109
Chapter 8. Creating and Maintaining Integrated Notes Databases	111
Initializing the original template and creating a database.	112
Creating customized production databases	116
Performing reconciliation	118
Database maintenance: refreshing design from a template	119
Chapter 9. Database Design Strategies and Advanced Customization.	121
Rules of thumb and general advice.	121
Using the Customization setup facility	122
Notes Database Customization	123
Modify TeamConnection Access.	124
Reconciliation of Notes and TeamConnection Data	124
Advanced customization	125

Part 4. Using TeamConnection to build applications. 127

Chapter 10. Basic build concepts	129
---	-----

The physical structure of the build function	129
The build object model	130
Parent-child relationships in a build tree	131
Working with a build tree	133
Putting the pieces together	135
Chapter 11. Installing, starting, and stopping build servers	137
Installing the build function	137
Creating a build server on MVS	137
Creating a build server on MVS/OE	140
Starting build servers using teamcbld	140
Starting an MVS build server	141
Starting the MVS/OE build server	143
Creating build startup files (for non-MVS environments)	144
Stopping the build servers	145
Stopping an MVS build server	145
Chapter 12. Working with build scripts and builders	147
Creating a builder	147
Writing a build script	150
Passing parameters to a build script	151
Writing a simple build script	153
Writing an executable file for a build script	153
Testing a build script	154
Modifying the contents of a build script	155
Putting a builder to work	155
Removing a builder from a part	156
Working with VisualAge C++ and Templates	157
Chapter 13. Working with MVS build scripts and builders	159
Creating a builder for MVS builds	159
Writing an MVS build script	163
File name conversions for MVS	164
Passing parameters to an MVS build script	164
TeamConnection syntax for MVS build scripts	165
Supported JCL syntax	166
Example of a build script for a C compile	168
Example of a build script for a COBOL compile	171
Example of a build script for a link	172
Chapter 14. Working with parsers	175
Creating a parser	175
Writing a parser command file	177
Putting a parser to work	178
Removing a parser from a part	179
Chapter 15. Building an application: an example	181
Starting the build servers	182
Creating builders and parsers	183

Creating the build tree for the application	184
Starting the build on the client	189
Putting the build scripts to work	191
Finishing the job and reporting the results to the user	192
Monitoring the progress of a build	192
Running a build in spite of errors	193
Building all parts, regardless of build times	193
Finding out which parts will be built	194
Canceling a build	194
More sample build trees	195
Defining multiple outputs from a single build event	195
Synchronizing the build of unrelated parts	196

Part 5. Using TeamConnection to package products. 199

Chapter 16. Using TeamConnection to package a product	201
Setting up your build tree for packaging	201
Setting up a build tree for the gather tool.	202
Chapter 17. Using the Gather tool	205
Using the teamcpak command for the Gather tool.	205
Command line flags.	206
Examples of the teamcpak gather command	208
Writing a package file for the Gather tool.	208
Syntax rules for a Gather package file.	208
Chapter 18. Using the Tivoli Software Distribution packaging tool	215
Using the teamcpak command with Tivoli Software Distribution	215
Command line flags.	216
Example of the teamcpak softdist command.	217
Writing a package file for Tivoli Software Distribution	217
Syntax rules for a Tivoli Software Distribution package file	217
Keywords for a Tivoli Software Distribution package file	218
Problem determination for the Tivoli Software Distribution tool.	220
Sample package file	221

Part 6. Appendixes 225

Part 7. Appendixes 227

Appendix A. Environment Variables.	229
Setting environment variables	235
Appendix B. Importing makefile information into TeamConnection	237
Creating a rules file	239
Appendix C. TeamConnection Merge	241

Appendix D. Enabling an OS/2 Workframe project for TeamConnection	243
Creating a TeamConnection-enabled Workframe project	243
Setting up your project options	243
Using your TeamConnection Workframe project	244
Project actions	244
Part actions	245
Using your project: a simple scenario	246
Appendix E. Enabling a Workframe/NT project for TeamConnection	247
Setting up your project options:	247
Using your TeamConnection WorkFrame project	248
Project actions	248
Part actions	249
Appendix F. Enabling and Using the VisualAge TeamConnection Enterprise Server Bridge.	251
Overview of the VisualAge TeamConnection Enterprise Server Bridge	251
Scope of this documentation	252
Description of the bridge	252
Preparing to use the VisualAge TeamConnection Enterprise Server Bridge	254
Setting up the bridge environment	254
Installing and activating the bridge	255
Using the VisualAge TeamConnection Enterprise Server Bridge	257
Setting default properties	257
Exporting VisualAge TeamConnection for Smalltalk components to TeamConnection	260
Importing VisualAge TeamConnection for Smalltalk components from TeamConnection	263
Using the bridge: a simple scenario for VisualAge Generator developers	264
Scenario assumptions	264
Exporting VisualAge TeamConnection for Smalltalk components to TeamConnection	264
Object mapping in TeamConnection	265
Build generation	266
Making a change to a member	267
Appendix G. Source Code Control User's Guide	269
Differences between other source code control providers and TeamConnection	269
Projects vs Families.	269
Installing the TeamConnection source code control DLL	270
Connecting TeamConnection to an IDE	271
Removing the TeamConnection Source Code Control DLL	271
Using TeamConnection as your source code control provider	271
Before you start	271
Opening a project	272
Integrated features	272
Full features of TeamConnection	273
Appendix H. Supported expandable keywords	277

Appendix I. Authority and notification for TeamConnection actions	279
Appendix J. Sample REXX execs, build scripts, and parsers	295
Sample REXX execs	295
Sample build scripts.	298
Sample parsers	299
Sample package files	300
Service and Support	301
VisualAge TeamConnection Services!	301
VisualAge TeamConnection Support!	301
IBM Lotus Passport Advantage Program	301
DB2 Service Maintenance and Technical Library	301
For North American Customers	302
Support for Customers Outside North America	303
Bibliography	305
IBM VisualAge TeamConnection Enterprise Server library	305
TeamConnection technical reports	306
DB2	306
Related publications	307
Glossary	309
Index	319
Readers' Comments — We'd Like to Hear from You	329

Figures

1. A sample TeamConnection client/server network	5
2. Sample of a component hierarchy	7
3. Parts, releases, and components.	8
4. Tasks window	19
5. Components window.	26
6. Accept Defects window	53
7. Create Work Areas window	54
8. Check Out Parts window	55
9. Part Filter window.	56
10. Edit Task List window	57
11. Check In Parts window	59
12. Build Parts window	61
13. Extract Parts window.	62
14. Check Out Parts	63
15. Check In Parts window	64
16. Freeze Work Areas window	65
17. Refresh Work Areas window	67
18. Build Parts window	68
19. Integrate Work Areas window	69
20. Verify Defects window	70
21. Refresh Work Areas window	71
22. Integrate Work Areas window	73
23. Reconcile Collision Record window	75
24. Modify Defect Owner window	79
25. Accept Defects window	80
26. Accept Approval Records window	82
27. Check Out Parts window	83
28. Check In Parts window	84
29. Freeze Work Areas window	85
30. Build Parts window	86
31. Complete Fix Records window	87
32. Add Driver Members window	88
33. Fix Work Areas window	89
34. Activate Fix Records window	90
35. Refresh Work Areas window	91
36. Refresh Drivers window	92
37. Build Parts window	93
38. Restrict Drivers window	94
39. Commit Drivers window	94
40. Complete Drivers window	95
41. Accept Test Records window	96
42. Notes Integrated Database Creation and Staging	112
43. The physical structure of TeamConnection	129
44. Sample build object model for msgcat.exe.	133
45. The build tree for the hello application	134
46. Two versions of a build tree	135

47. TeamConnection components on separate machines	137
48. Create Builder window	148
49. Modify Part Properties window	156
50. Modify Part Properties window	157
51. Create Builder window	160
52. A JCL fragment for an MVS compile	169
53. A JCL fragment converted to a build script	171
54. Create Parser window	176
55. Modify Part Properties window	178
56. Modify Part Properties window	180
57. Sample build tree	181
58. Sample build object model for msgcat.exe.	182
59. Create Parts window	185
60. Create Parts window	186
61. Modify Part Properties window	187
62. Connect Parts window	188
63. The build tree display	189
64. Build Parts window	190
65. The build tree for robot.dll	195
66. The build tree for robot.app	196
67. Part of the build tree for robot.app	202
68. Adding the gather step to the build tree	204

Notices

References in this publication to IBM products, programs, or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Subject to IBM's valid intellectual property or other legally protectable rights, any functionally equivalent product, program, or service may be used instead of the IBM product, program, or service. The evaluation and verification of operation in conjunction with other products, except those expressly designated by IBM, are the responsibility of the user.

IBM may have patents or pending patent applications covering subject matter in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to the IBM Director of Licensing, IBM Corporation, 500 Columbus Avenue, Thornwood, NY, USA 10594.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact the Site Counsel, IBM Corporation, P.O. Box 12195, 3039 Cornwallis Road, Research Triangle Park, NC 27709-2195, USA. Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement.

This document is not intended for production use and is furnished as is without any warranty of any kind, and all warranties are hereby disclaimed including the warranties of merchantability and fitness for a particular purpose.

IBM may change this publication, the product described herein, or both. These changes will be incorporated in new editions of the publication.

This publication contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

Trademarks

The following terms are trademarks of International Business Machines Corporation in the United States and/or other countries:

AIX®	NetView®
C/370™	OpenEdition®
C Set ++®	Operating System/2®
DB2®	OS/2®
DB2 Universal Database®	SOM®
IBM®	SOMobjects™
MVS™	TeamConnection™
MVS/ESA™	VisualAge®
MVS/XA™	XGA

Lotus and Lotus Notes are registered trademarks and Domino is a trademark of Lotus Development Corporation.

Tivoli, Tivoli Management Environment, and TME 10 are trademarks of Tivoli Systems Inc. in the United States and/or other countries.

The following terms are trademarks of other companies:

HP-UX 9.*, 10.0 and 10.01 for HP 9000 Series 700 and 800 computers are X/Open Company UNIX 93 branded products. HP-UX 10.10 and 10.20 for HP 9000 Series 700 and 800 computers are X/Open Company UNIX 95 branded products.

UNIX is a registered trademark in the United States and other countries licensed exclusively through X/Open Company Limited.

Intel and Pentium are registered trademarks of Intel Corporation.

Microsoft, Windows, Windows NT and the Windows logo are registered trademarks of Microsoft Corporation.

Java, HotJava, Network File System, NFS, Solaris and the Sun logo are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States and other countries.

Netscape Navigator is a U.S. trademark of Netscape Communications Corporation.

Adobe, the Adobe logo, Acrobat, the Acrobat logo, Acrobat Reader, and PostScript are trademarks of Adobe Systems Incorporated.

Other company, product, and service names may be trademarks or service marks of others.

About this book

This book is part of the documentation library supporting the IBM TeamConnection licensed programs. It is a guide for client users.

For additional information when performing TeamConnection tasks, refer to the *Commands Reference* when entering commands or online help when using the graphical user interface (GUI).

Getting Started with the TeamConnection Clients contains basic information for the client user.

This book is available in PDF format. Because production time for printed manuals is longer than production time for PDF files, the PDF files may contain more up-to-date information. The PDF files are located in directory path nls\doc\enu (Intel) or softpubs/en_US (UNIX). To view these files, you need a PDF reader such as Acrobat.

How this book is organized

“Part 1. Introducing TeamConnection” on page 1, gives all users an overview of the concepts of TeamConnection and introduces the terminology that is used throughout this book.

“Part 2. Developing a product using TeamConnection” on page 15, describes the different interfaces and basic TeamConnection tasks. It uses scenarios to explain how to do many tasks.

This part is for everyone using TeamConnection to do daily work. The information is meant for both the person who uses the command line interface and the person who uses the GUI, as instructions for both are provided.

“Part 4. Using TeamConnection to build applications” on page 127, tells how to use the TeamConnection build function. For information in installing and administering the build function, refer to the *Administrator's Guide*

“Part 5. Using TeamConnection to package products” on page 199, tells how TeamConnection helps you automate the packaging and distribution of your application.

“Part 6. Appendixes” on page 225, contains various pieces of information that you can refer to as you plan for and use TeamConnection.

Information on customer service, a bibliography, and a glossary are included at the back of this book.

Conventions

This book uses the following highlighting conventions:

- *Italics* are used to indicate the first occurrence of a word or phrase that is defined in the glossary. They are also used for information that you must replace.
- **Bold** is used to indicate items on the GUI.
- Monospace font is used to indicate exactly how you type the information.
- File names follow Intel conventions: **mydir\myfile.txt**. AIX, HP-UX, and Solaris users should render this file name **mydir/myfile.txt**.

Tips or platform specific information is marked in this book as follows:



Shortcut techniques and other tips



IBM VisualAge TeamConnection Enterprise Server for OS/2



IBM VisualAge TeamConnection Enterprise Server for Windows/NT



IBM VisualAge TeamConnection Enterprise Server for Windows 95



IBM VisualAge TeamConnection Enterprise Server for AIX



IBM VisualAge TeamConnection Enterprise Server for HP-UX



IBM VisualAge TeamConnection Enterprise Server for Solaris

Tell us what you think

In the back of this book is a comment form. Please take a few moments to tell us what you think about this book. The only way for us to know if you are satisfied with our books or if we can improve their quality is through feedback from customers like you.

Part 1. Introducing TeamConnection

Chapter 1. An introduction to TeamConnection.	3
TeamConnection definitions	4
TeamConnection's client/server architecture	4
TeamConnection database	5
Interfaces	5
Families	6
Users and host lists	6
Parts	6
Components	7
Releases	8
Work areas	8
Drivers	9
Defects and features	10
Processes	10
Build	11
Packaging	12
Roles people play	12

This section presents an overview of the TeamConnection product. The information in this section should be read and understood by everyone who is going to work with TeamConnection.

Additional conceptual information is provided in Parts 3, 4, 5, and 6.

Chapter 1. An introduction to TeamConnection

TeamConnection provides an environment and tools to make software development run smoothly, whether your development team is small or large. Using TeamConnection, you can communicate with and share data among team members to keep up with the many tasks in the development life cycle, from planning through maintenance.

What does TeamConnection do for you? It takes care of the following:

- *Configuration management*: the process of identifying, organizing, managing, and controlling software modules as they change over time. This includes controlling access to your software modules and providing notification to team members as software modules change.
- *Release management*: the logical organization of objects that are related to an application. The release provides a logical view of objects that must be built, tested, and distributed together. Releases are versioned, built, and packaged.
- *Version control*: the tracking of relationships among the versions of the various parts that make up an application. Version control enables you to build your product using stable levels of code, even if the code is constantly changing. It provides control over which changes are available to everyone and, optionally, allows more than one developer at a time to update a part.
- *Change control*: the controlling of changes to parts that are stored in TeamConnection. TeamConnection keeps track of any part changes you make and the reasons you make them. Your development team can build releases with accuracy and efficiency, even as the parts evolve. The product ensures that the change process is followed and that the changes are authorized. After changes are made, it allows you to integrate the changes and build the application. TeamConnection tracks all changes to the parts across multiple products and environments.

The *change control process* is configurable. Your team can decide how strict the change control should be, from loose to very tight. You can also adjust the level of control as you move through a development cycle.

- *Build support*: the function that enables you to define the structure of your application and then to create it within TeamConnection from your input parts. Independent steps in a build can run in parallel on different servers, thus reducing your build time. You can build applications for platforms in addition to the one TeamConnection runs on—currently, you can use TeamConnection to build applications on AIX, HP-UX, OS/2, Windows NT, Windows 95, Solaris, MVS, and MVS OpenEdition.
- *Packaging support*: the preparation of your application for electronic distribution to other users.

This chapter defines the basic terms and concepts you need to make the most of TeamConnection. Read this chapter first; then decide which information you need next:

Topic and description	Page
Developing products using TeamConnection:	16
<ul style="list-style-type: none"> • Getting familiar with the interfaces • The basics of using TeamConnection • More about <i>defects</i> and <i>features</i> • Following TeamConnection processes 	
Using TeamConnection to build applications:	128
<ul style="list-style-type: none"> • Build concepts • Installing build agents and processors • Working with build scripts and builders • Working with <i>parsers</i> • Building an application 	
Packaging applications:	“Chapter 16. Using TeamConnection to package a product” on page 201
<ul style="list-style-type: none"> • Using the packaging function • Using the Gather utility • Using the NVBridge utility 	

TeamConnection definitions

The following definitions are in logical order rather than alphabetical. provides additional information about these terms.

TeamConnection’s client/server architecture

Figure 1 on page 5 is an example of a network of TeamConnection clients and servers.

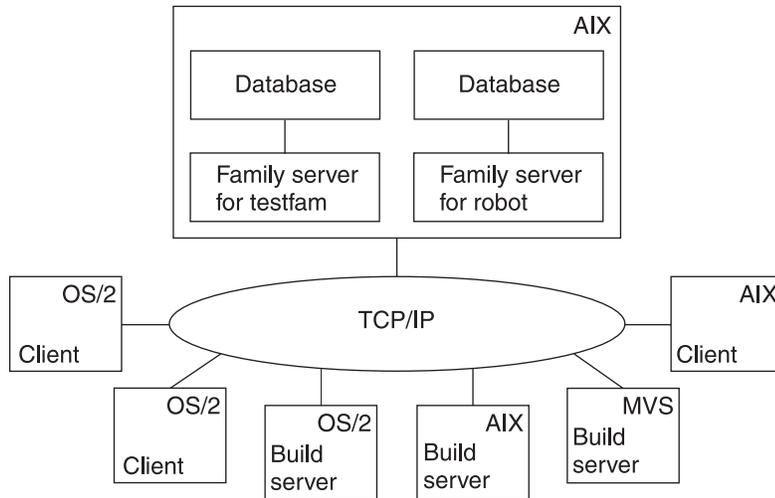


Figure 1. A sample TeamConnection client/server network

TeamConnection *family servers* control all data within the TeamConnection environment. Data stored in a family server's database includes:

- Text objects, such as source code and product documentation
- Binary objects, such as compiled code
- Modeled objects that are stored in the information model by tools such as VisualAge Generator
- Other TeamConnection objects that are metadata about the other objects

A TeamConnection *client* gives team members access to the development information and parts stored on the database server.

TeamConnection database

TeamConnection is built on IBM's DB2 Universal Database. Please refer to the DB2 documentation referenced in this document's "Bibliography" on page 305 for detailed information on DB2 database configuration, administration, and utilities.

Interfaces

TeamConnection provides the following interfaces that you can use to access data:

- A graphical user interface based on industry standards.
- A command line interface that lets you type TeamConnection commands from a prompt or from within TeamConnection
- A web client that you access through your web browser.

You can use any interface to do your TeamConnection work, or you can switch among them. This book usually gives instructions for using both interfaces.

For more information, see “Chapter 2. Getting familiar with the TeamConnection client interfaces” on page 17.

Families

A *family* represents a complete and self-contained collection of TeamConnection users and development data. Data within a family is completely isolated from data in all other families. One family cannot share data with another.

Refer to the *Administrator's Guide* for more information about families.

Users and host lists

Users are given access to the TeamConnection development data in a specific family through their *user IDs*. Each family has at least one *superuser*, who has privileged access to the family. The superuser gives other users the *authority* to perform some set of *actions* on particular data. Depending on the authority granted to a user, that user might in turn be able to grant some equal or lesser level of authority to other users. However, the ability to grant authority for some actions is reserved to the superuser. There are no actions which the superuser cannot perform.

For host-based authentication, each user ID is associated with a *host list*, which is a list of client machine addresses from which the user can access TeamConnection when using that ID.

A single user can access TeamConnection from multiple systems or logins. Likewise, a single system login can act on behalf of multiple users. The set of authorized logins for a TeamConnection user ID makes up the user's host list.

It is also possible to authenticate users through the use of passwords, either in place of host lists, or as an alternative form of authentication.

Refer to the *Administrator's Guide* for more information.

Parts

TeamConnection *parts* are objects that users and tools store in TeamConnection. They include text objects, binary objects, and modeled objects. These parts can be stored by the user or the tool, or they can be generated from other parts, such as when a linker generates an executable file. Parts can also be groupings of other TeamConnection objects for building and distribution, or simply for convenient reference. Common part actions include the following:

Create To store a part from your workstation on the server; from that time on,

TeamConnection keeps track of all changes made to the part. Or, to create a part to use as a place holder to store the output of a build.

Check out

To get a copy of a part so that you can make changes to it.

Check in

To put the changed part back into TeamConnection.

Extract To get a copy of the part *without* making changes to the *current version* in TeamConnection.

Edit To change a part from within TeamConnection using a specified editor.

Build To construct an output part from parts that you have defined to TeamConnection as input to the output part.

These are simplified definitions of part actions; there is more about the actions you can perform against parts in “Chapter 3. The basics of using TeamConnection” on page 25.

The current version of each part is stored in the TeamConnection database, along with previous versions of each part. You can return to previous versions if you need to.

Components

Within each family, development data is organized into groups called *components*. The component hierarchy of each family includes a single top component, called *root*, and *descendants* of that root. Each *child component* has at least one parent component; a child can have multiple parents.

The following figure depicts a component hierarchy.

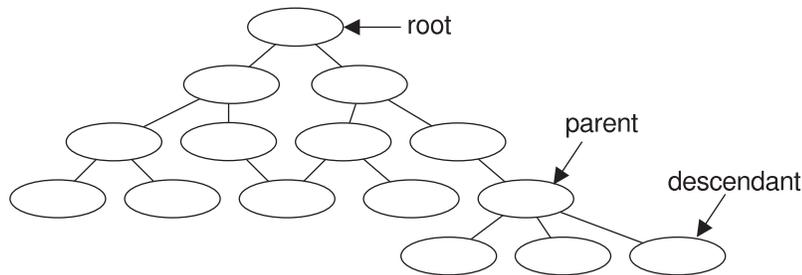


Figure 2. Sample of a component hierarchy

TeamConnection uses components to organize development data, control access to the data, and notify users when certain actions occur. Descendant components inherit access and notification information from ancestor components. Information about the components is stored in the database, including:

- The component's position in its family hierarchy.

- The user who owns the component. The component *owner* is responsible for managing data related to it, including defects or features.
- The users who have access to the component and the level of access each user has. This information makes up the component's *access list*.
- The users who are to be notified about changes to the component. This set of users is called the *notification list*.
- The *process* by which the component handles defects and features.

Releases

An application is likely to contain parts from more than one component. Because you probably want to use some of the same parts in more than one application, or in more than one version of an application, TeamConnection also groups parts into *releases*. A release is a logical organization of all parts that are related to an application; that is, all parts that must be built, tested, and distributed together. Each time a release is changed, a new version of the release is created. Each version of the release points to the correct version of each part in the release.

Each part in TeamConnection is managed by at least one component and contained in at least one release. One release can contain parts from many components; a component can span several releases. Figure 3 shows the relationships between parts, the releases that contain them, and the components that manage them.

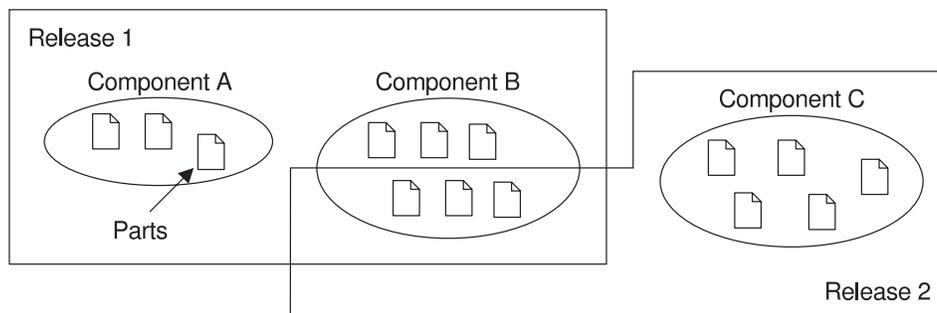


Figure 3. Parts, releases, and components

Each time a new development cycle begins, you can define a separate release. Each subsequent release of an application can share many of the same parts as its predecessor. Thus maintenance of an older release can progress at the same time as development of a newer one. Each release follows a process by which defects and features are handled.

Work areas

A release contains the latest "official" version of each of its parts. As users check parts out of the releases, update them, and then check them back in, TeamConnection keeps

track of all of these changes, even when more than one user updates the same part at the same time. To make this possible, TeamConnection uses something called a *work area*.

A work area is a logical temporary work space that enables you to isolate your work on the parts in a release from the official versions of the parts. You can check parts out to a work area, update them, and build them without affecting the official version of the parts in the release. After you are certain that your changes work, you *integrate* the work area with the release (or *commit* the driver that the work area is a member of, if you are using the driver subprocess). The integration makes the parts from your work area the new official parts in the release.

You can do the following with work areas:

- Check out parts from a release
- Update any or all of the checked-out parts
- Get the latest copies of the parts in the release, including any changes integrated by other users
- Get the latest copies of the parts in another work area
- *Freeze* the work area, making a snapshot of the parts as they exist at a particular instant in case you need to return to it later
- Build the parts in the work area
- Move all parts back into the release by integrating the work area

For more information, see “Using work areas” on page 28.

Drivers

A driver is a collector for work areas. You create drivers associated with specific releases so that you can exercise greater control over which work areas are integrated into the release and commit the changes from multiple work areas simultaneously.

When a work area is added to a driver, it is called a *driver member*. A single work area can be a member of more than one driver. By making a work area part of a driver, you associate the parts changed in relation to that work area with the specified driver. These parts must be members of the release associated with the driver.

Drivers enable you to place the following controls over work area integrations:

- Define and monitor prerequisite and corequisite work areas to ensure that mutually dependent changes are integrated in proper order.
- Monitor and resolve conflicting changes to the same part (if you use concurrent development).
- Restrict access to driver members so that they can be changed only by users with proper authority.

Defects and features

A defect is a record of a problem to be fixed. A feature is a record of a request for a functional addition or enhancement. Both may be associated with a work area, and both follow the processes defined for the component and release that are associated with the work area. TeamConnection tracks both objects through their life cycles as developers change and commit parts.

You can use defects and features to record problems and design changes for things other than the products you are developing under TeamConnection control. For example, you can use defects to record information about personnel problems, hardware problems, or process problems. You can use features to record proposals for process improvements and hardware design changes.

For more information, see “Working with defects and features” on page 38.

Processes

An application changes over time as developers add features or correct defects. TeamConnection controls these changes according to the *processes* you choose for your application’s components and releases. A process enforces a specific level of control to part changes and ensures that actions occur in a specified order.

Two separate types of processes are defined: component processes, which can be different for each component within a family, and release processes, which apply to all activities associated with a given release. Component or release processes are built from a number of lower-level processes, or *subprocesses*, that are included with the TeamConnection product.

A defect or feature written against a component moves through successive *states* during its life cycle. The TeamConnection actions that you can perform against it depend on its current state. The component processes define these actions. You can require users to do some, all, or none of the following for tracking defects and features:

dsrFeature

Design, size, and review changes to be made for features

verifyFeature

Verify that the features have been implemented correctly

dsrDefect

Design, size, and review fixes to be made for defects

verifyDefect

Verify that the fixes work

At the release level you can require some, all, or none of the following subprocesses:

track This subprocess is TeamConnection’s way of relating all part changes to a specific defect or feature and a specific release. Each work area gathers all

the parts modified for the specified defect or feature in one release and records the status of the defect or feature. The work area moves through successive states during its life cycle. The TeamConnection actions that you can perform against a work area depend on its current state.

You must use the *track subprocess* if you want to use any of the other release subprocesses.

approval

This subprocess ensures that a designated *approver* agrees with the decision to incorporate changes into a particular release and electronically signs a record. As soon as approval is given, the changes can be made.

fix

This subprocess ensures that as users check in parts associated with a work area, an action is taken to indicate that they have completed their portion. When everyone is done, the owner of the *fix record* (usually the component owner) can change the fix record to complete. The parts are then ready for integration.

driver

A *driver* is a collection of all the work areas that are to be integrated with each other and with the unchanged parts in the release at a particular time. The driver subprocess allows you to include these changes incrementally so that their impact can be evaluated and verified before additional changes are incorporated. Each work area that is included in a driver is called a *driver member*.

test

The test subprocess guarantees that testing occurs prior to verifying that the fix is correct within the release.

TeamConnection is shipped with several predefined processes. If these do not apply to your organization, you can configure your own processes by defining different combinations of subprocesses.

See “Chapter 4. The states of TeamConnection objects” on page 41 for an explanation of TeamConnection states.

Build

The TeamConnection build function automates the process of building individual parts or entire applications, both in the work group LAN environment and on an enterprise server. This function enables you to reliably and repeatedly build the same output from the same inputs. You can also build different outputs from the same inputs for different environments.

You start a build against an output part that has an associated *builder*. A builder is an object that describes how to translate input parts to get the desired output, such as a linker or compiler. An input part might have an associated parser, which determines the dependencies for the input parts in a build.

The build function does the following:

- Tracks build times of inputs and outputs so that it builds only those parts that are out of date themselves or that have out of date dependents. You can also force a build regardless of the build times.
- Enables you to spread the build over multiple machines running at the same time or into multiple processes running on a single machine, such as on MVS.

For more information, see “Part 4. Using TeamConnection to build applications” on page 127.

Packaging

Packaging is any of the steps necessary to distribute software and data onto the machines where they are to be used. TeamConnection includes two tools that you can use to automate the electronic distribution of TeamConnection-managed software and data:

Gather An automated data mover for server or file transfer-based distribution.

Tivoli Software Distribution

A bridge utility that automates the installation and distribution of software or data using Tivoli as the distribution vehicle.

For more information, see “Part 5. Using TeamConnection to package products” on page 199

Roles people play

Because TeamConnection is extremely flexible, no two projects are likely to use it in the same way, and the jobs that people perform likewise vary. Still, TeamConnection tasks can be grouped into the following general categories:

System administrator

Has *superuser* access to the family server and database administration access to the database management system. This administrator is responsible for the following:

- Installing and maintaining the TeamConnection server
- Maintaining and backing up the database used by TeamConnection

Note: On UNIX systems, the system administrator must also have root access to the host machine.

Family administrator

Has *superuser* access to the family server and database administration access to the database management system. This administrator is responsible for the following:

- Planning and configuring TeamConnection for one or more families
- Managing user access to one or more families

- Maintaining one or more families
- Creating and updating configurable fields
- Configuring release and component processes for a family
- Creating and updating user exits
- Monitoring the user activity of a family

Build administrator

This administrator is responsible for the following:

- Setting up and maintaining build servers
- Planning for builds
- Creating builders and parsers
- Starting and stopping build servers
- Defining *pools*
- Monitoring build performance
- Creating driver members
- Committing and completing drivers
- Extracting releases
- Packaging and distributing applications

End user

End users, such as project leaders, programmers, and technical writers, use one or more TeamConnection families to control and maintain application development data.

Part 2. Developing a product using TeamConnection

Chapter 2. Getting familiar with the TeamConnection client interfaces	17
Using the GUI.	17
Starting the GUI	18
Stopping the GUI	19
Performing tasks with the GUI	19
Using the Settings notebook	20
Online help information.	21
Using the command line interface	22
Using the TeamConnection web client.	23
Chapter 3. The basics of using TeamConnection	25
Laying the groundwork	25
Authority to perform tasks.	26
Finding objects within TeamConnection	27
Finding parts	27
Using work areas	28
Naming your work areas	29
Creating parts.	29
Naming your parts	30
Preparing to build your parts	30
Working with parts	31
Working in serial or concurrent development mode	31
Working with common parts	32
Getting parts from TeamConnection	33
Checking parts in to TeamConnection	34
Finding different versions of TeamConnection objects	35
Versioning releases	35
Versioning work areas	36
Versioning drivers	37
Versioning parts	38
Working with defects and features	38
Testing and verifying part changes	39
Chapter 4. The states of TeamConnection objects.	41
Defects and features	41
The states of work areas	45
The states of drivers	47
Verification and test records	49
Chapter 5. Working with no component or release processes	51
Working in serial development	51
Accepting a defect	52
Creating a work area	53
Checking out a part	54
Searching for a part.	55
Checking in a part	58

Verifying and testing part updates	60
Extracting a part	62
Checking out the part one more time	63
Checking the part back in	64
Freezing the work area.	65
Refreshing the work area	66
Building the application	67
Integrating the work area	68
Closing a defect	69
Working in concurrent development	70
Refreshing the work area from the driver.	71
Integrating the work area	72
Reconciling differences.	73
Chapter 6. Working with component and release processes	77
Moving through design, size, and review	78
Changing defect ownership	79
Accepting a defect	80
Approving the fix	81
Checking out a part	82
Checking in the changes	83
Freezing the work area.	85
Building the application	86
Accepting fix records	87
Integrating changed parts into a release	88
Adding a driver member	88
Reconciling the differences	89
Returning the work area to the fix state	89
Reactivating the fix record.	90
Refreshing the work area	90
Refreshing the driver	91
Building the driver	92
Restricting the driver	93
Integrating the parts.	94
Completing the driver	95
Testing the built application	95
Using a configured process	96
Retrieving a past version of a part	97

This section is for anyone who uses the TeamConnection client to do daily work. The information is meant for both the person who uses the command line interface and the person who uses the GUI; instructions for both are provided.

All the tasks in this part are done from a client machine.

Before reading this section, you should be familiar with the TeamConnection terminology and concepts presented in “Chapter 1. An introduction to TeamConnection” on page 3.

Chapter 2. Getting familiar with the TeamConnection client interfaces

TeamConnection provides several interfaces that you can use to access data:

- A graphical user interface based on industry standards.
- A command line interface that lets you type TeamConnection commands from a prompt or from within TeamConnection.
- A web client, that you access through your web browser.

You can use any of the interfaces to do all of your TeamConnection work, or you can switch back and forth between the them. You might find that some tasks are easier to do from the GUI or through the web, while others are easier to do from the command line.

The examples throughout “Part 2. Developing a product using TeamConnection” on page 15 give instructions for both GUI and command line interface usage.

This chapter helps you to begin using the TeamConnection client interfaces. It describes the following:

- Using the GUI
 - Starting and stopping the GUI
 - Getting around in the GUI
 - Using the Settings notebook
 - Using the online help that is provided with TeamConnection
- Using the command line interface
- Using the web client

Before you can use TeamConnection, someone in your organization with superuser or admin authority, such as your family administrator, must create for you a unique user ID and a host list entry for the workstation where you installed the client.

Using the GUI

TeamConnection provides a GUI that you can use to do all of your TeamConnection work. To use the GUI efficiently, set your default values in your Settings notebook to suit your working environment, and then become familiar with the Tasks window and how you can save time by adding your most common tasks to it.

Note: If available, use the **Select Preference File** window to specify your configuration file for the TeamConnection GUI. This is the file that TeamConnection uses to store and retrieve the information you have selected for the various GUI windows. Preferences such as which view you are using for a window (for

example Details view on the Defects window), the column order for a window, and the window size are saved in this file.

If you are using AIX, HP_UX, or Solaris and are about to use the GUI for the first time, you need to do the following tasks:

1. See *Configure the environment variables in the .profile* in the Getting Started with the TeamConnection Clients manual.
2. See *Ensure that the TeamConnection client command is accessible* in the Getting Started with the TeamConnection Clients manual.
3. Copy the sample initial tasks list for the main GUI window by typing the following command (this needs to be done only once):

```
cp $TC_HOME/nls/cfg/$LANG/teamcv3x.ini $HOME/.  
chmod u+w $HOME/teamcv3x.ini
```

Where \$TC_HOME is the location where the TeamConnection code was installed.

Starting the GUI

You can start the TeamConnection client GUI in one of the following ways:

- Select the **TeamConnection Client** icon from the TeamConnection Group folder on the desktop.
- Type teamcgui from a prompt.

The Tasks window appears.

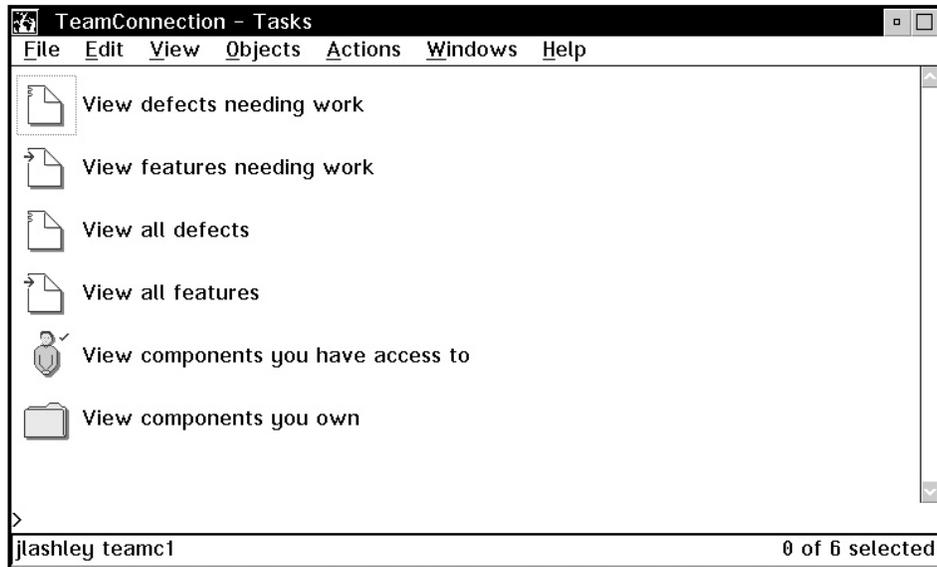


Figure 4. Tasks window

Initially, a set of default tasks appears in your Tasks window. As you become more familiar with TeamConnection and see what tasks you do most often, you can change, delete from, and add new tasks to this list. To learn how to do this, select **How do I** from the Help pull-down menu, and then select **Update tasks on the Tasks** window.

From the Tasks window, you can either select actions from the menu bar or select a task.

Stopping the GUI

To stop the TeamConnection client GUI, do one of the following:

- Select **Close** from the System menu in the Tasks window.
- Select **Exit** from the File pull-down menu of a TeamConnectionwindow.

Performing tasks with the GUI

There are several ways you can perform TeamConnection tasks with the GUI. You can:

- Select an action from the Actions pull-down menu and then select the object you want to work with. For example, if you want to view a specific defect, select **Defects** → **View** from the Actions pull-down menu; then type the name of the defect in the View Defects window.

This method is useful when you know the exact names of the objects you want to work with.

- Select the type of object you want to work with, such as **Defects**, from the Objects pull-down menu. A Filter window appears in which you can specify search criteria. You then get a list of objects that match the search criteria. Online help provides information about using the Filter window.

This method is useful when you do not know the exact name of the object you want to work with, or you want to view a list of objects.

After you have a list of objects, and if you are going to use this list at other times, you can keep the window open. Leave the window in the background as you do your other work, or minimize it. This way, you can quickly retrieve the list when you want to perform another action.

- Select a task from the Tasks window.

This method provides a fast path within the GUI. When you select a task, TeamConnection performs the underlying query or command and then displays the requested information.

- Select an object from an object window (such as the Parts, Defects, or Features window) and then select an action to perform on the selected part from the **Selected** menu. You can also display a pop-up menu listing valid actions for a specific object by placing the mouse pointer over the object and pressing mouse button 2.

Using the Settings notebook

The TeamConnection GUI provides a Settings notebook in which you can set default values for your working environment. To open the Settings notebook, select **Settings** from the Windows pull-down menu. You can set the following values; for more information about them, refer to the online help. The notebook has five pages.

Note: The environment variables you specify are relevant for the GUI only, not the command line.

<p>On the Environment page:</p>	<ul style="list-style-type: none"> • Family • Release • Component • Work area • Become user • User ID • Top • Relative directory • Working directory
<p>On the Setup page:</p>	<ul style="list-style-type: none"> • NLS path • Log file • Case • Print command • Compare command • Edit command

On the GUI page:	<ul style="list-style-type: none"> • Verbose commands • Auto refresh • Multiple object windows • Show query line • Sort pre-defined list values • Use small icons in icon views (not available in Windows or UNIX clients) • Use small icons in tree views (not available in Windows or UNIX clients) • Font for object windows (not available in Windows or UNIX clients) • Font for output windows (not available in Windows or UNIX clients) • Required field label color • Modified field label color
On the Extract page:	<ul style="list-style-type: none"> • Destination directory • Read-only • Expand keywords
On the Pool page:	<ul style="list-style-type: none"> • Pool

Online help information

Online help information is available from anywhere in the TeamConnection GUI. Use the online help when you need more information about a topic or task.

TeamConnection offers two types of help:

- General help

This is help for a specific window. General help provides an overview of the task and describes the objects on the window, such as menu-bar items, icons, fields, and push buttons. Do one of the following to access general help:

 - Select **Help** from a menu bar.
 - Select the **Help** push button.
 - Press F1.
- How do I

This is where you find step-by-step instructions for doing a specific task. How you do a task depends on the component or release process that is being followed, and this help information takes that into consideration. To access this help, select **How do I** from the Help pull-down menu. Double-click on one of the task items.

At the bottom of each Help window is a **Diagram** push button. Select this push button to view a graphical process diagram. Step your way through the diagram to better understand the processes that TeamConnection components and releases can follow. The processes that your components and releases follow depend on how the

processes are configured for your organization. The defined processes determine the actions that must occur before a defect or feature can move toward completion.

Using the command line interface

To use the command line interface effectively, you must be familiar with the actions that you can perform using TeamConnection commands. A complete description of each command, including examples for each, is available in the *Commands Reference*

To view the syntax of a TeamConnection command online, type the following at a prompt:

```
teamc commandName
```

Where *commandName* is the name of the TeamConnection command.

The *Quick Commands Reference* is a booklet that lists the syntax of each TeamConnection command.

You can also become familiar with the commands by looking at the contents of the log file where TeamConnection stores the commands that are issued as you use the GUI. This file is specified in the **Log file** field on the Setup page of the Settings notebook. The default name is teamc.log; it is stored in the directory where the client is installed (for AIX, HP-UX, and Solaris it is stored in the \$HOME directory of the user), unless you specify a different location in the Settings notebook.

You can type TeamConnection commands from a prompt within any directory; the TeamConnectionGUI does not need to be started. Or if you start the GUI, you can type a command on the command line in the Tasks window (this command line is located at the bottom of the window, just above the footer that indicates the user name and family name).

Before you start to use the command line interface, you might want to set the most used environment variables, such as TC_FAMILY or TC_COMPONENT. You are not required to set these environment variables, but if you do not, you will need to specify them in the command when required.

You set environment variables differently for different platforms:

- AIX, HP-UX, and Solaris users set environment variables in the .profile (sh, ksh environment), .dtprofile (cde environment), or .cshrc (csh environment).
- OS/2 users set environment variables in the config.sys file or from a command line prompt.
- Windows 95 and Windows NT users set environment variables in the Windows Control Panel.
- Some environment variables are set in your config.sys file during installation.

You can override the value you set for an environment variable by using the corresponding flag in the command. For example, you have the TC_FAMILY environment variable set to robot, but you need a file from another family named octo, so you issue the following command:

```
teamc part -extract hello.c -family octo -release 9501
```

“Appendix A. Environment Variables” on page 229 provides a complete list of the TeamConnection environment variables.

Using the TeamConnection web client

The TeamConnection web client provides family server connectivity and a great deal of the functionality provided by a standard TeamConnection client without the overhead required by a standard client installation. Using a web browser, anyone in the organization can access server data (provided the server is configured appropriately) by addressing a machine and a port number. Most familiar TeamConnection functions are available through the web client. If you want to disable the web client interface, you must set the environment variable TC_WWWDISABLED before starting the family.

To begin using the TeamConnection web client you must point your web browser to the correct URL. The syntax of the URL is: *http://host name of the server:port number of your family*. For example, if your server host name is *testfam* and your port number is *7890*, the URL would look like: **http://testfam:7890**

When connecting to the web client, you will be prompted for a user name and a password. The browser will cache this user name and password until it is closed. The user name is required, but the password may not be, depending on the authentication level of the family and the host on which the browser is running.

If your organization authentication level is set for PASSWORD_ONLY, you must login using a password to gain access. If your organization authentication level is HOST_ONLY, you may be denied access to the web client because either the host being used is not in the user's host list, or the browser is not bypassing the proxy. You may be able to gain access by doing one of the following:

- If the host being used is not in the user's host list, you must add the host to the user's host list to gain access.
- If the browser is not bypassing the proxy, do one of the following:
 - Temporarily disable the proxy when using the web client.
 - Reconfigure the browser to bypass the proxy when connecting to the family server.
 - Use an autoproxy so the browser will automatically bypass the proxy when connecting to the family server.
- Change the authentication level to HOST_OR_PASSWORD and login using a password.

Using the web client is much like using the TeamConnection GUI. The following are some differences you might find:

- Releases, Drivers, and Work areas cannot be extracted with the web client.
- A filter for Corequisites does not exist in the TeamConnection GUI.
- For the FeatureModify action, the following are available with the TeamConnection web client (but not the TeamConnection GUI):
 - newName
 - orginLogin.
- For the DriverMemberView action, the following are available with the TeamConnection web client (but not the TeamConnection GUI):
 - state
 - defectName
 - defectAbstract
 - committedVersion.
- For the ChangeView action, the following are available with the TeamConnection web client (but not the TeamConnection GUI):
 - ChangeView
 - workAreaState.
- For the PartBuild action, the following are available with the TeamConnection web client (but not the TeamConnection GUI):
 - cancel
 - partType.
- The PartChildInfoView action does not exist in the TeamConnection GUI.
- For the PartDelete action, force is available for use with the TeamConnection web client (but not the TeamConnection GUI).
- For the PartDisconnect action, parentType is available for use with the TeamConnection web client (but not the TeamConnection GUI.)
- For the PartModify action, fileType is available for use with the TeamConnection web client (but not the TeamConnection GUI).
- For the PartUnlock action, Source Directory is available in the TeamConnection GUI but not in the TeamConnection web client.
- For the PartViewContents, Expand Keywords is available in the TeamConnection GUI. In the TeamConnection web client, an option for turning keyword expansion on and off is provided.
- For the UserView Filter action, the following are available with the TeamConnection web client (but not the TeamConnection GUI):
 - pswStatus
 - pswModifyTime
 - pswCreateTime.

Chapter 3. The basics of using TeamConnection

All users of TeamConnection perform a number of basic tasks, such as checking parts out of TeamConnection and then back in, and testing and verifying part changes. Before you start doing these tasks, you need to understand the basic concepts behind them; that is what this chapter explains.

This chapter assumes that you have read “Chapter 1. An introduction to TeamConnection” on page 3 and are familiar with the different objects, such as components and releases. The other chapters in this part of the book define in more detail how you perform the TeamConnection tasks.

Laying the groundwork

Someone has already created your family’s component structure, and those components manage your parts and control access to the data. Your TeamConnection family also contains releases. A release identifies a version of all the parts that comprise an application at a given point in time. When you create a release, you specify the component that will manage it. One component manages a release, but many components can manage the individual parts associated with that release.

A single part can be associated with more than one release, but it is managed by one component. When you create a part, you specify the release that you want to associate with the part and the component that you want to manage it. At any time, you can link the created part to other releases so that the part can be shared, or you can change its managing component.

Before you start working with parts, you need to be familiar with your family’s component structure. This will help you when trying to locate parts within TeamConnection and when writing defects and features. You can do the following to display your family’s component structure from the GUI:

1. Select **Components** → **Components** from the Objects pull-down menu on the Tasks window. The Component Filter window appears.
2. Type the name of the component that is at the top of your component hierarchy in the **Component** field, and select **OK**. Initially this component is called root. The Components window appears, listing the component.
3. Verify that the component is displayed in tree view (a plus sign (+) appears before the component name). If not, select **Tree** from the View pull-down menu.
4. Select **Expand fully** from the Selected pull-down menu.

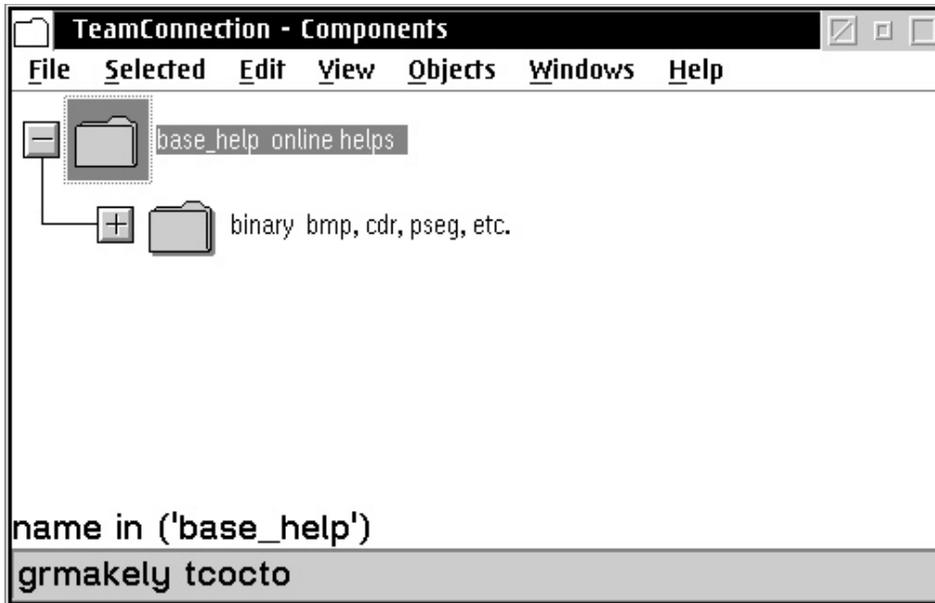


Figure 5. Components window

From a command prompt, you can issue the following command to view the component structure.

```
teamc report -view -raw bCompView -where "name='root'"
```

Authority to perform tasks

As a TeamConnection user, you are automatically given the authority to perform some basic tasks. You can:

- Open defects and features
- Add notes to existing defects and features
- Modify the information for your user ID
- Display information about any user ID
- *Search* for information within TeamConnection to create reports

You receive authority to perform additional actions when you become the owner of a TeamConnection object, such as a component or a part, or when authority is explicitly given to you by the component owners.

If you attempt an action that you do not have authority to do, TeamConnection tells you so. When this happens, you can ask the component owner, the family administrator, or a user with superuser authority to grant you the necessary authority.

Note: You can issue queries to generate reports of data from tables and views using the -view action flag. If you do not specify selection criteria, such as the fields and the search conditions you want to use, the report query selects all entries for the table or view indicated that the user has authority to access. This command does not show any objects in components that you are not authorized to access

“Appendix I. Authority and notification for TeamConnection actions” on page 279 lists the types of authority you need in order to perform various TeamConnection actions.

Finding objects within TeamConnection

All TeamConnection objects are stored on a server in a database. To find one or more of these objects within a family, do one of the following:

- Use the report command with the -view action flag from a command line or a command line within TeamConnection.

Command usage is explained in the *Commands Reference*

- Use a Filter window in the GUI.

Online helps explain how to use the Filter windows.

For now, you need to understand that the database is case-sensitive. You need to refer to and search for objects in the correct case. For example, if a component is stored in the database as hand, you would not find it if you typed Hand or HAND. This is why it is important that your organization sets a naming convention, and that everyone follows that decision when creating objects. If you do not know what naming convention has been established for your organization, talk to your family administrator.

Note: It is recommended that you use all lowercase whenever possible.

Finding parts

There are three Filter windows that you can use to find parts within TeamConnection:

Note: Use the forward slash (/) when specifying path names, such as directory/file.txt. The Intel convention of using the backward slash is not recognized in the Filter window.

Parts Use when you want to limit your search to a particular context of a work area or driver in a release, or a particular version of a release. This is generally the view users will use most often.

If you specify only a release, TeamConnection lists the committed parts for that release. However, if you want a list of all parts in a specified work area and release, TeamConnection displays all the parts visible to the work area. This includes parts that are committed to the release as well as changed parts that are visible only to the work area.

BuildView

Use when you want to search for information related to building your application, such as viewing a build tree, or when you want to do build actions.

PartFull

Use when you want to search for parts across releases, components, or work areas. For example, you want a list of all the optics.c parts. Unlike the Parts Filter, you can specify one or more release or work area names.

You can also use this filter to display only parts that have been changed in a work area. For example, you check out robot.c to work area 310:1, and that is the only part that you have changed. If you use the PartFull Filter to query for all the parts in work area 310:1, only one record is returned.

You cannot use this filter to search for build information.

Refer to the online help, in particular **How do I**, for more information on how to use the Filter windows. Select **How do I** from the Help pull-down menu to access the information.

Using work areas

A work area is a logical temporary work space that enables you to isolate your work on the parts in a release from the official versions of the parts. You can check parts out to a work area, update them, and build them without affecting the official version of the parts in the release. You must create a work area before you can create, check out, or check in parts. If your component's process includes a design, size, review subprocess for defects or features and the release follows a tracking subprocess, a work area is automatically created when sizing records exist and the associated defect or feature is accepted. TeamConnection associates these work areas with the appropriate defect or feature.

The parts in a work area do not become available in the release until the work area is integrated. Also, if your release follows a driver subprocess, parts that have been changed do not become available in the release until the associated driver is committed. However, users who have the authority to access the work area can view and work with the parts in it.

You can save intermediate versions of the parts in your work area by *freezing* your work area. Every time you freeze a work area, TeamConnection saves a revision level of the work area. When you freeze work area 123:1, for example, a version called 123:2 is created. This version contains information about each part in the work area and its current version at the time the work area was frozen. It may contain version 1 of part optics.c, for example. If you freeze the work area again later, a new version called 123:3 is created with information about the versions of the parts in the work area when it was frozen. This version may contain version 2 of part optics.c. Each of these work area versions is saved in the database and you can retrieve the versions of the parts they contain before you integrate the work area into the release. Therefore, you should freeze a work area whenever there is a possibility that you will want to return to that

version of the work area. For example, you might be adding a major feature to the code, and you want to be able to return to something that works in case the application no longer builds. When you integrate a work area or commit a driver, the work area is frozen automatically.

Naming your work areas

When TeamConnection automatically creates a work area, the work area is given the same name as the defect or feature it was created for plus the initial version number, :1. When you create a work area, you can also give it the same name as the defect or feature, or you could give it any other name. Where possible, we recommend that you name it after a defect or feature, or relate the name to the change that is being made.

Here are some things you should know before you name a work area:

- Work area names must be unique within the context of a release.
- After you create a work area, you cannot delete it. You can, however, cancel the work area in the following situations:
 - No part changes were made.
 - You undo the changes you made.
- With the proper authority, other users in your organization will be able to access your work area and make changes to the parts. This means that you need to make it easy for them to locate the work area. Following your local naming conventions will help.
- After the work area is integrated with the release, you cannot reuse the work area. If the defect is still in the working state, you can create another work area with a different name after the initial work area is integrated with the release.

Creating parts

A TeamConnection part is controlled by the TeamConnection server. A TeamConnection part is uniquely identified by the *path name* of the part, the part *type*, and the name of the release in which it is contained. You must specify both the release name and the path name whenever you perform a TeamConnection action on a part. Multiple releases can share the same part.

When you create a part, you do one of the following:

- Take an existing text or binary file that is on your workstation and place it into TeamConnection.
- Create an empty part that has no content. Empty parts are used as place holders until an application is built. For example, you can create a place holder for an executable part that will be created by a build. See “Creating the build tree for the application” on page 184 for an example of creating a place holder.

After you put a part under TeamConnection control, the official copy of the part resides in the database. The copy on your workstation is changed to read-only mode. You can then change the part by checking it out to your workstation or editing it within the GUI.

Use the online help facility if you need assistance when creating parts.

Naming your parts

If your organization has a naming convention, be sure to follow it when naming your parts. When the naming convention is not followed, everybody in your organization can have trouble locating parts. Part names created on the server are case-sensitive; they must be retrieved using the same case in which they were created.

When you name TeamConnection parts, you can specify only the base name, such as `hand.c`, or you can specify the directory path in addition to the base name, such as `robot\hand\hand.c`. Specifying the path name as part of the name lets you have several identical base name parts included in the same release, such as `robot\hand\msg.h` and `robot\optics\msg.h`.

You can also have identical part names within the context of a release as long as their part types are different, such as `TCPart` and `vgdata`.

Note: It is recommended that you use lowercase letters to name your parts.

A parts name may contain spaces provided it is enclosed in double quotation marks during processing. For example:

```
teamc part -create "This is a long file name.txt"
```

The name with spaces will be shown as-is by the GUI (without the double quotes). If the name has spaces and is not enclosed in double quotation marks, then you may get an error message repeated many times, one for each "token" separated by spaces in the long name.

Note: The base name may contain a maximum of 63 characters, not including the double quotations. The path name, which includes the base name, may contain a maximum of 195 characters, not including the double quotations.

Preparing to build your parts

If you are going to use the TeamConnection build function, you must provide certain information about each part that participates in a build. You can provide this information when you create the parts or wait until later. You can also change the information at any time.

To associate a part with a build, you must specify the following information:

- The parent part that you want to associate the part with.
- The type of relationship the part has to the parent, such as:

Input The part will be used as input to building its parent. An example of an input part is a C language source file, `x.c`, which is compiled to create its parent, `x.obj`.

Output The part will be a generated output from the same build that creates its parent part. In other words, both the parent part and this child part are outputs when the parent part is built.

Dependent

The part will be needed for the build operation of its parent to complete, but it will not be passed directly to the builder. An example of this is an include file.

If you do not provide this information when you create the part, you can provide it later using the *connect* function.

You can also specify the builder or parser that a part will use, as well as any build parameters.

“Part 4. Using TeamConnection to build applications” on page 127 explains the build function in more detail.

Working with parts

After the parts are created in TeamConnection, you will be working with these parts — getting them to your workstation so you can change them and then getting them back in to TeamConnection. This section gives a brief overview of these tasks. “Chapter 5. Working with no component or release processes” on page 51 and “Chapter 6. Working with component and release processes” on page 77 go into more detail about these and other TeamConnection tasks.

Working in serial or concurrent development mode

A release is set up for either serial development or concurrent development mode. Once the development mode is set you can change from serial mode to concurrent mode, but not from concurrent mode to serial mode. In serial development, a part is locked when a user checks it out, and no one else can update the part as long as it is checked out. In concurrent development, more than one user can simultaneously have the same part checked out.

In concurrent development, more than one user can check out and change the same part. Prior to integrating their changes, each user should refresh their work area from the driver in which they plan to put it. The first user will be able to integrate their work area, complete fix records for tracking releases, with their release. When the next user refreshes their work area, TeamConnection recognizes that the parts differ and notifies them. It is up to this user to resolve the differences, using the TeamConnection merge program, VisualAge TeamConnection VisualMerge Facility, or some other merge program. If the user fails to refresh their work area from the driver, TeamConnection will not notify them that the parts differ until they try to add the work area to the driver. They will then have to put the fix records back in the fix state, reactivate the work area,

refresh the work area, reconcile the differences, refresh their work area again, and reintegrate their work area before they can add their changes to the driver.

Before getting parts from TeamConnection, you might want to find out if the development mode for the release is concurrent or serial. To determine the mode, view the information about the specific release. To do this, select **View** from the Selected pull-down menu on the Releases window.

Working with common parts

A *common* part is a part with identical content that is shared by two or more releases or two or more work areas. For example, when an identical part is needed in two separate releases, you can link the part from one release to the other (if you have the proper authority). Both releases would then have a link to the current version of that part.

When a common part is checked out of a release, TeamConnection locks the current version of the part in all releases if one of them uses serial development. When putting the part back into the release, one of the following actions reflects the change in all releases in which the part is common:

- You integrate the work area when the driver subprocess is not followed, or
- You commit the driver when the driver subprocess is followed.

You can break the common link if you make changes to a common part and you do not want these changes reflected in other releases or work areas that link to the part. You can break the common link when you check out, check in, rename, delete, re-create, connect, or *disconnect* parts. When a part is common to more than two releases, you can maintain the common link with some of the releases while breaking the link with other releases. When a link is broken, the parts still share the same name, but the information contained in the parts is different.

Parts can also be linked between two or more work areas in the same or different releases, making the parts common to those work areas. For example, a user working in one work area can link to the latest version of a part in another work area of the same release (the part has yet to be integrated with the release). The part is then common to the two work areas within the same release. If you want to maintain the common link to all work areas, you must specify the names of the common work areas when you check in, rename, delete, or re-create the parts. As with common parts in releases, you can break the common link.

You can also link all the parts within a release to another release. This function is especially helpful when development begins on a new release of a product, and you want the parts in the new release to initially be the same as the parts in the current release. As development of the two releases continues, the common link between the parts can be broken to separate development of the new release from maintenance of the current release.

For more information about how to link parts, refer to the *Commands Reference* and online help.

Getting parts from TeamConnection

Checking out a part implies that you intend to modify it; extracting a part merely gives you a copy of the part. Normally, when you extract a part, you do not plan to change the current version in TeamConnection.

You must have the necessary authority to a component before you can check out or extract parts from that component. You need *PartExtract* authority to extract a part from TeamConnection; you need *PartCheckOut* authority to check a part out. See “Appendix I. Authority and notification for TeamConnection actions” on page 279 for a listing of all the TeamConnection actions and their authority requirements.

Parts are checked out to work areas. The work area is where you store updated parts and do builds without affecting the version of the parts in the release. When a part is checked out of the release to the specified work area, TeamConnection locks the part in the release if you are working in serial development. If you are working in concurrent development, the part is never locked. TeamConnection also puts a copy of the part on your workstation. It is here where you update the part. If a read-only copy of the part exists on your workstation, the first character of its file extension changes as follows: It becomes a “\$” (dollar sign) for OS/2 and Windows platforms. It becomes a “.” (period) for AIX, HP-UX, and Solaris platforms. This copy is a backup copy. If a backup copy already exists, it is deleted. When you are finished updating the part, you check it back in to the work area. A work area is optional when extracting a part.

The environment variables TC_BACKUP and TC_MODPERM may be set to change the the backup and read-only options. TC_BACKUP controls whether or not backup files are created. If this environment variable is set to off, backup files are not created. TC_MODPERM controls whether or not the read-only attribute is set. The default for this environment variable sets the read-only attribute on.

When you extract a part, TeamConnection copies the part to your workstation, and the part is not locked. In other words, other users can still check out the same version of the part and make changes to it, even in serial development mode. By default, TeamConnection sets the extracted part to read-only access. This is done to keep you from inadvertently changing the part on your workstation when the part in TeamConnection is not locked. You can, however, change this in the Settings window or when you are extracting the part. When you do this, be aware that someone else can change the official part in TeamConnection, making your workstation copy back level.

Where TeamConnection places a checked-out or extracted part on your workstation depends on the following:

- Your workstation’s *current working directory*
- Whether you use the -relative flag on the command line or the **Destination directory** field on the GUI
- Whether the TC_TOP environment variable is set

For more information about how these interact, refer to the part command examples in the *Commands Reference*.

When you want to make changes to a part, you can do one of the following:

- Check out one or more parts and edit the parts on your workstation. When you finish making changes to the parts, you check them back in.
- Edit a part from within the TeamConnection GUI using a specified editor. When you exit the editor, the Check In Parts window appears and you can check the part back in to TeamConnection.

In either case, if you are working in concurrent development and someone else changed a part while you had it checked out, you are asked to resolve the differences when you try to integrate your work area.

Checking parts in to TeamConnection

After you have verified the accuracy of your part changes, you are ready to check them in to TeamConnection. Any parts that you have checked out, you have the authority to check back in.

As mentioned earlier, you check parts out to a work area so you can work on them. Therefore, when you check in a part, you must specify the work area where that part is checked out. In other words, you check the part back in to the same work area. When the part is checked in, the copy on your workstation is flagged read-only. The environment variable TC_MODPERM controls whether or not the read-only attribute is set. The default for this environment variable sets the read-only attribute on.

At this time, the changed part is visible in only the named work area; it is not visible at the release or to any other work area. This lets you test your changes by building the version of the code that is in your work area.

When you are satisfied with your changes, you can integrate the parts into the release by integrating your work area. This action makes the work area visible to all the users in the release.

If you are working in concurrent development mode, TeamConnection generates a *collision record* when a changed part conflicts with a previously committed part. For example, both you and Keith have hand.c checked out. Keith makes changes to the part and then integrates the work area that contains that part. (Depending on the process being followed, Keith might have to commit the work area rather than integrate it.) Later, after making changes to hand.c, you attempt to integrate the work area that contains the part. Because the part was already integrated by Keith, you are notified of a collision and asked to refresh your work area. After the refresh, you can view the collision record and decide how you want to resolve the conflicts. “Reconciling differences” on page 73 explains in more detail how this works.

Finding different versions of TeamConnection objects

TeamConnection version control maintains different versions of the following objects:

- Releases
- Work areas (and driver members)
- Drivers
- Parts

When you want to find and retrieve previous versions of these objects, it is helpful to know how TeamConnection creates and deletes previous versions of each object.

Some basics of TeamConnection versioning will help you understand how TeamConnection identifies unique versions of objects:

- When you first create an object, the initial version name is the object name suffixed with *:1*. When you create a new work area called *myWorkArea*, for example, its version is *myWorkArea:1*. Subsequent versions are identified in numerical order: *myWorkArea:2*, *myWorkArea:3*, *myWorkArea:4*, and so on. Versions of releases and drivers are identified similarly: *myRelease:1*, *myRelease:2*, *myRelease:3*; *myDriver:1*, *myDriver:2*, *myDriver:3*; and so on.
- Unique versions of parts are identified by association with a specific version of a release, work area, or driver. Your TeamConnection family may have three different versions of a part called *myPart*, for example: one associated with release *myRelease:2*, one associated with work area *myWorkArea:1*, and one associated with work area *myWorkArea:2*.

Versioning releases

TeamConnection creates new versions of releases whenever you do the following:

- Create a release
This is the initial version of a release and contains no parts. When you create *myRelease*, for example, its version name is *myRelease:1* and it contains no parts.
- Commit a work area to the release
Committing a work area to a new release creates a new version of the release and adds the parts in the work area to the release. When you commit work area *myWorkArea:1*, for example, to *myRelease:1*, TeamConnection creates a version of *myRelease* called *myRelease:2*. It also associates the parts in *myWorkArea:1* with *myRelease:2*.
- Commit a driver to a release
Because drivers are simply collections of work areas, committing a driver to a release has the same effect as committing a work area: TeamConnection creates a new version of the release. When you commit *myDriver:2* to *myRelease:2*, for example, TeamConnection creates a version of *myRelease* called *myRelease:3*.

TeamConnection deletes versions of releases whenever you prune the release. Refer to the *Administrator's Guide* for an explanation of pruning.

Versioning work areas

TeamConnection creates new versions of work areas whenever you do the following:

- Create a work area

This is the initial version of a work area. When you create myWorkArea, for example, its version name is myWorkArea:1.

- Refresh a work area

Refreshing a work area updates it with any new versions of parts that have been integrated with the release. When a workarea is refreshed, two versions of the workarea are created. One of the contents before the refresh and one with the contents after the refresh.

- Freeze a work area

Freezing a work area is like taking a snapshot of the work area. It preserves the parts as they are at a given point in time. If you create work area myWorkArea:1, add three new parts to it — called part1, part2, and part3 — and then freeze it, your family contains a work area called myWorkArea:2, with part1, part2, and part3. The version name of each of these parts is myWorkArea:1. If you then alter part2 and freeze the work area again, your family contains the following:

- myWorkArea:1, with nothing in it
- myWorkArea:2 contains part1, part2, and part3 at version myWorkArea:1
- myWorkArea:3 contains part1 and part3 at version myWorkArea:1, and part2 at version myWorkArea:2

- Commit a work area

Committing a work area adds the parts in the latest version of the work area to the release. It also does the following:

- Creates a new version of the release
- Creates new versions of the parts in the release
- Deletes any intermediate versions of the work area

Using the previous example, if you commit myWorkArea:3 to myRelease:1, the following happens:

- TeamConnection creates a new version of myRelease called myRelease:2.
- TeamConnection creates new versions of the parts in myRelease:2.
- TeamConnection deletes myWorkArea:1, myWorkArea:2, and myWorkArea:3.

TeamConnection deletes versions of work areas whenever you commit them to a release. Once a work area has been committed, you can no longer use it for making part changes and you cannot create a new work area with the same name.

Deleting work area versions is controlled by the autopruning option of the release associated with the work area. By default, TeamConnection always deletes work area versions on commit, but you can change this option. Refer to the *Administrator's Guide* for an explanation of autopruning.

Versioning drivers

TeamConnection creates new versions of drivers whenever you do the following:

- Create a driver
When you create a new driver, TeamConnection makes two versions of it: myDriver:1, for example and myDriver:2.
- Add a work area (driver member) to a driver
If you add myWorkArea:1 to myDriver:2, for example, TeamConnection creates a new version of myDriver called myDriver:3.
- Freeze a driver
Freezing a driver is like taking a snapshot of the driver. It preserves the parts as they are when the driver is frozen. If you freeze myDriver:3, for example, TeamConnection creates a new version called myDriver:4.
- Refresh a driver
Refreshing a driver updates the driver with all changes that have been made in all of its driver members. Refreshing a driver actually creates two new versions of the driver, as follows:
 1. Freezes the driver (so that TeamConnection can have a point to roll back to if an error occurs during the refresh operation).
 2. Updates the driver with any changes from the driver members
 3. Freezes the driver again, thus preserving a copy of the updated driver.

If the current version of myDriver is myDriver:2, for example, and the parts in its driver members have been changed, then TeamConnection does the following when it refreshes the driver:

1. Freezes myDriver, creating myDriver:3.
2. Updates myDriver with changes from its driver members.
3. Freezes myDriver again, creating myDriver:4.

The result of refreshing myDriver (version myDriver:2) is two new versions: myDriver:3, containing a snapshot of the driver before the refresh, and myDriver:4, containing a snapshot of the driver after the refresh.

TeamConnection deletes versions of drivers whenever you remove driver members or commit a driver to a release.

- If you have a driver version myDriver:4 with driver members myWorkArea, yourWorkArea, and ourWorkArea, and you remove myWorkArea, then TeamConnection deletes driver versions myDriver:2, myDriver:3, and myDriver:4 and creates a new driver version called myDriver:5 containing members yourWorkArea and ourWorkArea. As a result, the family contains two versions of the driver, myDriver:1 and myDriver:5.
- When you commit a driver to a release, all intermediate versions of the driver (resulting from driver member add, driver freeze, driver refresh, or driver member remove operations) are deleted.

Versioning parts

TeamConnection versions parts in association with other TeamConnection objects, such as work areas. If, for example, you create part1 in myWorkArea:1, the current version of part1 is myWorkArea:1. If part1 is in release myRelease:2 and work area myWorkArea:2, then you can view the version of the part for either the release or the work area. The version label for part1 in myRelease:2 is myRelease:2 and in myWorkArea:2 is myWorkArea:2.

TeamConnection deletes part versions whenever it deletes versions of the object that the part is associated with. In addition to versioning in association with other TeamConnection objects, TeamConnection maintains versions of build output parts (parts that are created as the result of a build, such as an .exe file or a .hlp file). When you create a release, you can set the maximum number of versions of build output parts to maintain. If you set this maximum to 10, for example, then TeamConnection saves only 10 versions of build output parts and discards the oldest version each time a new version is created.

Working with defects and features

Defects are used to report problem information; features are used to record information about proposed design changes. After a defect or feature is opened, TeamConnection tracks the progress of the defect or feature through various states. To what degree defects and features are tracked depends on the processes followed by the release and component to which they are assigned. The following describes actions that your defined processes might require:

Analyzing defects and features

The owner is responsible for analyzing a defect or feature after it is opened. The owner can then return it if it is not valid or feasible, reassign it to another user or component, or accept it for resolution.

Designing the resolution

After a defect or feature has been accepted, the actual resolution needs to be designed so that an informed evaluation can be made. This resolution needs to be designed by users who are familiar with the product or area affected by the defect or feature.

Identifying the required resources

Sizing records are created by the owner to identify the components and releases that might be affected by the defect or feature. Each owner of a component that is referenced in a sizing record needs to evaluate the impact of the defect or feature on the parts managed by the component. If the defect or feature requires changes to parts, the sizing record is accepted and sizing information is added.

When sizing records exist and the associated defect or feature is accepted, TeamConnection automatically creates a work area.

Reviewing the design and resource estimates

After the resolution has been designed and the resources have been identified, the proposal needs to be reviewed. If the review indicates that work should continue on the defect or feature, it is accepted.

Resolving defects and implementing features

Resolving one defect or implementing one feature in one release can involve one or more users changing many parts. To change a part, a user must check out the part, make the changes required to resolve the problem or implement the design change, and check the part back in. If the release follows a tracking process, all defects or features must be associated with a work area. Parts that are checked out refer to the work areas that are monitoring the defect or feature.

Resolving a defect or implementing a feature also involves integrating the changed parts with changes made for other defects and features in that release. All changed parts are eventually integrated with the unchanged parts within the release.

Verifying the resolution of the defect or feature

The originator uses a verification record to acknowledge that the defect or feature was satisfactorily resolved or not. Accepting a verification record formally closes the defect or feature. Rejecting a verification returns the defect or feature to working state.

“Chapter 4. The states of TeamConnection objects” on page 41 explains in more detail the various states that different TeamConnection objects can go through depending on the process that is being followed. A diagram in this chapter shows the flow of these states. You might want to study this information before you start to work with defects and features.

Testing and verifying part changes

You can use TeamConnection's build function to build your program. Before you check in updated parts, you will probably want to verify the accuracy of your changes.

The scenarios in “Chapter 5. Working with no component or release processes” on page 51 and “Chapter 6. Working with component and release processes” on page 77 include information about testing and verifying part changes. “Part 4. Using TeamConnection to build applications” on page 127 provides detailed information about the build function.

Chapter 4. The states of TeamConnection objects

The actions that you can perform on certain TeamConnection objects are controlled by two factors:

- The process followed by the component and by the release
- The current state of the object

Certain TeamConnection objects follow certain states through their life cycle. An instance where an object might not follow all the possible states is when it moves through the states defined in the subprocesses of the component and the release. The following table briefly lists the component and release subprocesses. For more information on component and release subprocesses, refer to the *Administrator's Guide*.

Component subprocesses

- **dsrDefect** - Design, size, and review fixes to be made for defects
- **verifyDefect** - Verify that defect fixes work
- **dsrFeature** - Design, size, and review changes to be made for features
- **verifyFeature** - Verify that feature changes work

Release subprocesses

- **track** - Relate all part changes to a specific defect or feature and a specific release
- **approval** - Require all changes to be approved before incorporating them into a release
- **fix** - Use fix records to ensure that all required changes are made
- **driver** - Use drivers to integrate changes into a release
- **test** - Require all changes to be tested before they are integrated into the release

This chapter explains the possible states of certain TeamConnection objects and how objects are moved from one state to the next. It also explains how component and release subprocesses affect the flow of states. For a diagram showing the flow of states, refer to the poster *Staying on track with VisualAge TeamConnection processes*.

Defects and features

Use defects to report problem information; use features to record information about proposed design changes. After you open defect or feature, TeamConnection tracks the progress of the defect or feature through various states. Defects and features are tracked according to the processes followed by the release and component that they are assigned to. The possible states for defects and features are:

Open state

When you open a defect or feature, it is in the open state and you are considered the originator.

You assign the defect or feature to a component. The owner of this component becomes the feature or defect owner and is responsible for managing its resolution. The component you open a defect or feature against should be one that manages the parts affected by the enhancement or problem. Use the component descriptions and the structure of your family's hierarchy to find the most appropriate component. If you open a defect or feature in an inappropriate component, the component owner can reassign it.

While the defect or feature owner is responsible for implementation, the originator is responsible for verifying that the defect or feature is resolved correctly.

Returned state

A defect or feature owner can return a defect or feature to its originator. You can return a feature or defect from the open, design, size, or review state if you decide that the defect or feature is not feasible or not valid. You can return a defect or feature back to the working state only if it has no associated work areas. If there are associated work areas, you must cancel or undo them before you can return the defect or feature. When you return a defect or feature, add your reason for returning it so that the originator and any other users can evaluate why you believe it is not feasible or not valid.

Canceled state

A feature or defect in the open or returned state can be canceled only by its originator or by a superuser. A canceled defect or feature remains inactive unless it is reopened by the originator.

Design state

If the component to which a defect or feature is assigned includes the `dsrDefect` or `dsrFeature` subprocess, you move defects or features in the open or returned state to the design state.

In this state, the proposed change is designed, and a description of the design change is entered. The owner must describe the design change before the defect or feature can move to the next state.

If the release includes the fix subprocess, fix records are automatically created when a defect or feature is designed.

Size state

Defects or features move to this state after the owner enters design information.

In this state, users can create a *sizing record* for each release that contains parts affected by the enhancement or problem. A sizing record identifies the work that is required for and the resources affected by the defect or feature. The owner of the component that is referenced in the sizing record is the owner of the sizing record. The owner is responsible for entering information about the amount of work that is required to implement the feature or resolve the problem.

The sizing record owner can reject the sizing record if it does not affect the specified component. After all sizing records are either accepted or rejected, the defect or feature moves to the review state or returns to the design state if more design information is needed.

Review state

Defects or features move to this state after they have been sized. In this state, the design text and sizing records are reviewed to determine the feasibility of the proposal. The owner can do one of the following:

- Accept the defect or feature if all design and sizing records are acceptable. This moves the defect or feature to the working state.
- Return the defect or feature to the originator if all design and sizing records are not acceptable. If necessary, the originator can reopen a defect or feature.
- Move the defect or feature back to the design state if design modifications are needed.

Working state

Defects or features move to this state when the owner accepts the defect or feature when it is in the:

- Review state, if the component includes the `dsrDefect` or `dsrFeature` subprocess
- Open state, if the component does not include the `dsrDefect` or `dsrFeature` subprocess

When you accept a defect or feature, you accept the responsibility of resolving it. A defect or feature might require changes in more than one release.

What happens after a defect or feature is accepted varies according to the subprocesses in effect:

Component subprocesses

- **`dsrDefect` or `dsrFeature`** - TeamConnection creates a work area in the approve state for each release identified in the accepted sizing records for the defect or feature.
- **`verifyDefect` or `verifyFeature`** - TeamConnection creates verification records in the notReady state.

Release subprocesses

- **`fix`** - TeamConnection creates fix records in the notReady state based on the sizing records.
- **`approval`** - TeamConnection creates approval records for each user on the release's approver list.

If the component does not include the `dsrDefect` or `dsrFeature` subprocess, then you must manually create a work area before you can check out or create parts to address the defect or feature.

Verify state

Defects and features go through the verify state only if their component

includes the verifyDefect or verifyFeature subprocess. Defects and features are automatically moved to this state when all work areas (there can be multiple work areas for the defect or feature) for the release are integrated. If a release is specified on the defect or feature when it was created, then the defect/feature will only go to the verify state when all the work areas for the release specified move to complete state.

Note: If you do not intend to fix that defect for the release specified, then you must move the defect to verify state manually by doing a defect -verify. This should be done by the originator of the defect.

When a defect or feature is accepted, TeamConnection creates a *verification record*. This record lets the originator:

- Accept the fix if the resolution was satisfactory
- Reject the fix if not satisfied with the resolution
- Abstain if unable to assess the resolution

Once all verification records have been accepted or abstained, the defect or feature moves to the closed state. If a verification record is rejected, the defect or feature returns to the working state. The defect or feature cannot be closed until the verification records are accepted.

A defect or feature can have more than one verification record. For example, if defect 123 is returned because it is a duplicate of defect 122, a second verification record is created for defect 122. The originator of defect 123 is the owner of the second verification record for defect 122. If the originator is the same for both defects, only one verification record is created.

Note: For a discussion of verification records and test records, see “Verification and test records” on page 49.

Closed state

The closed state is the final state of a defect or feature.

If the defect is associated with multiple work areas, the defect will remain in the working state until all of the work areas are integrated.

If the component includes the verifyDefect or verifyFeature subprocess, the defect or feature automatically moves to the closed state after all verification records are in the accept or abstain state and all work areas are in the complete state. If a verification record is rejected, the defect or feature moves back to the working state. Otherwise, the defect or feature moves directly from the working state to the closed state when the first work area moves to the complete state.

You cannot re-open a defect or feature that is in the closed state. If the defect or feature was not resolved correctly, you must open a new defect or feature to address the necessary changes.

The states of work areas

A work area is a storage area where you can work on the parts in a release without affecting the "official" versions of those parts. A work area can be associated with a specific defect or feature, but it does not have to be. These attributes can affect the state of a workarea:

workareafixhold	With the workareafixhold attribute and the fix subprocess, a workarea will remain in the fix state rather than moving to the integrate state when the final fix -complete command has been issued. A workarea -integrate command must be issued to move the workarea into the integrate state.
workareacommithold	If an environment list exists for the release associated with the work areas, the automatic transition to the test state can be disabled by including the -workareacommithold attribute in the release process. A workarea -test command must be issued to move the workarea into the test state.
workareatesthold	With the workareatesthold attribute and the test subprocess, a workarea will remain in the test state rather than moving to the complete state when the final test is marked. A workarea -complete command must be issued to move the workarea into the complete state.

Approve state

When a work area is created, it goes to this state if the release includes the approval subprocess. TeamConnection creates an approval record for each user on the release's *approver list*. Each approver indicates their evaluation of the changes in their approval record:

- Accept that work should continue
- Abstain if unable to assess if work should continue
- Reject if work should not continue

When all approval records are marked as abstain or accept, the work area goes automatically to the fix state. If any approval record is marked as reject, the state of the work area remains at approve. You can change rejected approval records to accept or abstain.

Fix state

If the release does not include the approval subprocess, work areas for the release begin in the fix state.

While the work area is in this state, parts are checked out to the work area, changes are made to these parts, and builds are done to verify the accuracy of the changes.

If the release includes the fix subprocess, any fix records created for a defect or feature move to the active state when a part change is associated with the work area for the defect or feature. A fix record monitors the part changes within a single component. Fix records provide a mechanism for reviewing all part changes that apply to components before integrating those changes with changes made for other defects and features.

How fix records are handled varies according to the subprocesses in effect:

Component subprocesses

- **dsrDefect** or **dsrFeature** - TeamConnection creates fix records for features or defects when existing sizing records are accepted.

Release subprocesses

- **fix** - If a fix record does not already exist for the component, TeamConnection creates one when a part managed by that component is checked in to the database.

If neither of these subprocesses are in place and a defect or feature owner needs to create a work area manually, he or she can create fix records at the same time. Existing fix records go to the active state when a part is checked in to the work area.

Fix records provide a way of ensuring that all necessary part changes within the specified component have been made and are reviewed or inspected. The fix record owner is responsible for this review. When the fix record owner is satisfied that the part changes made within that component are complete and ready for integration with other parts in the release, the owner marks the fix record as complete. When all existing fix records for a work area are complete, the work area automatically moves to the integrate state.

Integrate state

Work areas can be moved to the integrate state as follows. If the release includes the fix and driver subprocesses, the work area automatically moves to the integrate state when all fix records are complete. If all fix records are not complete, you can force a work area to the integrate state, provided that no part changes are associated with the work area. If the release does not include the fix and driver subprocesses, you must move the work area to the integrate state manually.

While a work area is in integrate state, you must add it to an existing driver as a driver member if the release includes the driver subprocess. All work areas in the integrate state do not have to be added to the same driver. The work area stays in the integrate state until the driver in which it is a member is committed.

You can move work areas from the integrate to the following states, according to the subprocesses in effect:

Release subprocesses

- **driver** - A work area moves to the commit state when the driver it is a member of is committed, or to the restrict state when the driver is restricted. You can also force a work area to the commit state, provided that no part changes are associated with the work area.
- **test** - A work area moves to the test state so that test records can be approved or rejected.

If the release does not include these subprocesses, you can manually complete a work area in the integrate state.

Restrict state

Work areas can be moved to the restrict state only when the release includes the driver subprocess. The work area moves automatically to the restrict state when the driver to which it belongs is restricted. If a work area in this state is removed from the driver, it returns to the integrate state. Otherwise, the work area remains in the restrict state until the driver to which it belongs is committed.

Commit state

Work areas can be moved to the commit state only when the release includes the driver subprocess. The work area moves automatically to the commit state when the driver to which it belongs is committed. At this point, all parts that were changed in this release to resolve the feature or defect are committed. The work area remains in the commit state until the driver to which it belongs is completed.

Test state

Work areas can be moved to the test state only when the release includes the test subprocess. When the associated driver moves to the complete state or when a work area is committed without a driver, the work area moves to the test state. The driver is then ready for formal testing in the specified environments. Test records for the work area are created in the ready state when the work area moves to the test state. The work area stays in the test state until all test records are accepted, rejected, or abstained.

Complete state

The complete state is the final state of a work area; the work area can no longer be used. If the test subprocess is not included in the release process, the work area moves directly to this state when the associated driver is completed or when the work area is explicitly integrated.

When a work area is completed, the feature or defect associated with that work area automatically moves to the verify or complete state. The defect does not leave the working state until the work area for that release is completed.

The states of drivers

Drivers monitor and implement the integration of part changes within a release. Those part changes are included in a driver by adding the work areas containing the changed parts to the driver as driver members.

Working state

The working state is the initial state of a driver. While the driver is in this state, it is not associated with any work areas and, therefore, contains no part changes.

If the release includes the driver subprocess, drivers can be explicitly created at any time.

Integrate state

Each driver automatically moves to the integrate state when the first work area

is added to it as a driver member. If all work areas are removed from the driver, the driver automatically returns to the working state.

Work areas can be added to drivers as driver members when the driver is in the working, integrate, or restrict state and the work area is in the fix state. Adding driver members to a driver in restrict state requires proper authority.

You can extract the driver when it is in the integrate state; however, only those parts that were changed in reference to driver members are extracted. This is referred to as extracting a *delta part tree*.

Restrict state

Before a driver is committed, you can move it to the restrict state. While a driver is in this state, work areas in the integrate state can be created for or deleted from the driver by only a superuser or an individual with the special authority of memberCreateR or memberDeleteR. This allows a build administrator to have better control over what is being built. The build administrator can delete driver members that are causing build errors or add driver members to fix build errors. You can then commit an error-free driver.

When a driver moves to the restrict state, all work areas that are included as driver members also move to the restrict state.

Commit state

Committing a driver commits all work areas included as driver members and all parts that were changed in reference to those work areas. TeamConnection commits only a successfully built driver. Committing a driver changes it to the commit state. You can, however, manually commit the driver. You can also commit an unsuccessful driver by using the force option.

When a driver moves to the commit state, all work areas that are included as driver members also move to the commit state. When a work area is in the commit state, all part changes associated with the work area become the "official" versions of the parts in the release and are visible to all users of the release.

A committed driver can be extracted as a full part tree as well as a delta part tree. A full part tree is the part structure of all the parts within the release.

Complete state

The complete state is the final state of a driver. In this state, the driver is ready for formal testing in the specified environments.

If the release includes the test subprocess, the work areas that are included as driver members move to the test state. Any existing test records for the work area move to the ready state when the work area moves to the test state. The work area stays in the test state until all test records are accepted, rejected, or abstained.

Test records are used to record the outcome of environment tests for changes implemented in a driver. This record lets the owner:

- Accept the record if the test was satisfactory
- Reject the record if not satisfied with the test results

- Abstain if unable to assess the results

Once all test records have been accepted or abstained, the states of other objects change as follows:

- **Work areas** - Go to complete state.
- **Defects and features** - Go to verify state if the component includes the verifyDefect or verifyFeature subprocess; otherwise they go to the closed state.
- **Verification records** - Go to ready state and are sent to the defect or feature originators.

If the test subprocess is not configured, then work areas move to the complete state and any defects or features move to the verify state.

If the component includes a verifyFeature or verifyDefect subprocess, verification records move to the ready state and notification is sent to the originators of any defects or features that were addressed in the completed driver.

The commit and complete states of drivers differ as follows:

- When a driver is committed, all work areas are committed, but no changes occur in the states of defects or features associated with the work areas.
- When a driver is completed, then the states of other associated objects (such as test records, work areas, verification records, defects, and features) change according to the other subprocesses in effect:
 - **test** - Work areas go to the test state and test records are created in the ready state for each environment in the release's environment list.
 - **verify** - Verification records go to the ready state.

If the release includes neither of these subprocesses, then the work area goes to the complete state and all features and defects associated with the work area are closed.

Verification and test records

If you use both the verify component subprocess (verifyDefect or verifyFeature) and the test release subprocess, then TeamConnection creates both verification records for features or defects and test records for each environment defined in the release's environment list. These records serve different purposes:

- Verification records provide a means of accepting or rejecting the product changes made in response to defects or features and are thus specific in nature.
- Test records provide a means of accepting or rejecting the results of a build and are more global in nature.

These records are handled by different people and enable you to monitor your development progress in different ways. The sequence of creating and handling verification and test records is as follows:

1. Verification records are created in the notReady state when a defect or feature is accepted. This indicates that someone on the development team has begun implementing the changes warranted by the defect or feature, but the changes are not yet ready to be verified. A work area is also created for the part changes.
2. When a driver is committed all part changes associated with the driver members are integrated into the release.
3. To create test records, the driver is completed. This action creates one test record for each environment on the release's environment list. The testers on your development team use the test records to accept or reject the results of their tests on the part changes.
4. After all test records have been accepted or abstained, the verification records are moved to the ready state. This indicates that the part changes have been tested in the context of the build and each individual defect or feature is ready to be accepted or rejected by the person who opened it.
5. The defect or feature originator accepts or abstains the verification record to close the defect or feature. The originator can also reject the verification record to move the defect or feature back to working state.

Chapter 5. Working with no component or release processes

To illustrate how to work with objects in a release that does not follow a tracking process or component processes, this chapter follows an example of a programming team that is developing the control systems for a robot. They are working in a family called robot.

Instructions for performing the task are given for both the graphical user interface (GUI) and the command line interface (Command).

This chapter illustrates two scenarios: working in serial development and working in concurrent development. Working in *serial development* means that after you check out a part, TeamConnection locks the part so that it cannot be updated by anyone else. Compare this to *concurrent development*, in which more than one person can simultaneously update the same part.

The following table directs you to the scenario you need:

For this scenario,	Go to this page.
Working in serial development	51
Working in concurrent development	70

Working in serial development

Alex is one of the programmers working on the robot application within a release called robot_control. The release does not follow a tracking process, and the release supports serial development. Even though the release does not follow a tracking process, defects are opened when problems are found.

This example assumes that the parts that Alex will work with have already been created in the release, and the build tree has been established. The build tree shows the hierarchy of objects that take part in the build of an application. It identifies parts as inputs, outputs, and dependencies of a build. For more information about build trees, see "Working with a build tree" on page 133 or "Creating the build tree for the application" on page 184.

This example also assumes that the family named robot has been defined in the TC_FAMILY environment variable. Because Alex accesses information in several releases, he has not defined the release named robot_control. Therefore, he must explicitly identify the release when performing TeamConnection actions, but not the family.

A fellow team member, Carol, has discovered that the robot's aperture is not working correctly. To address this problem, she opens a defect. To fix the problem, Alex needs to make some modifications to the parts in this release. This fix will require the tasks

noted in the following table:

For information about this task,	Go to this page.
Accepting the defect	52
Creating a work area	53
Checking out a part	54
Searching for a part	55
Checking in a part	58
Verifying and testing part updates	60
Freezing the work area	65
Refreshing the work area	66
Building the application	67
Integrating the work area	68
Closing the defect	69

Accepting a defect

Alex is notified via electronic mail that defect 310 has been opened against the robot component. After some research, he agrees that there is a problem with the aperture of the robot's on-board camera, so he accepts the defect. Alex does one of the following:

GUI

From the GUI, he:

1. Selects **Defects** → **Accept** from the Actions pull-down menu on the Tasks window. The Accept Defects window appears.

Note: The Accept Defects window in this example may be different than one you may see based on your environment. Configurable fields may or may not be shown depending on any configurable fields set by you or your administrator,

2. Types 310 in the **Defects** field and selects **program_defect** from the **Answer** list.
3. Selects **OK**.

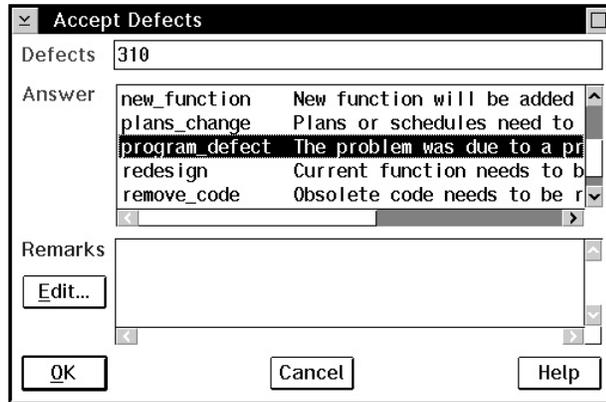


Figure 6. Accept Defects window

Command

From a command line, he issues the following command:

```
teamc defect -accept 310 -answer program_defect
```

Result

The defect goes to the working state.

Creating a work area

Because the component is not following a design, size, and review process, Alex needs to manually create a work area in which to modify and build his parts. (If the component follows a design, size, and review process, a work area is automatically created when the defect moves to the working state, provided that sizing records have been accepted for the defect.)

Before Alex checks out any parts, he creates a work area that will contain the latest view of the parts in the release by doing one of the following:

GUI

From the GUI, he:

1. Selects **Work areas** → **Create** from the Actions pull-down menu on the Tasks window.
2. Types 310 in the **Work areas** field and robot_control in the **Releases** field and selects **OK**.

Note: 310 is the name of the defect that was opened for the problem, so this is how Alex wants to identify the work area.

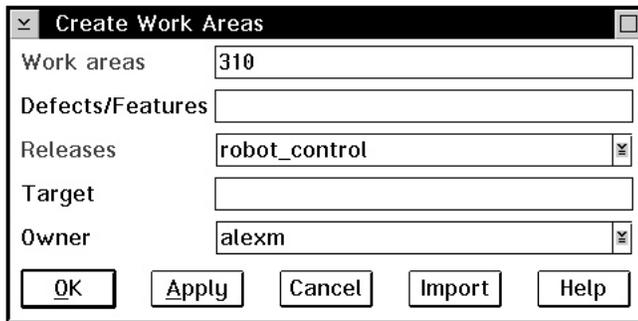


Figure 7. Create Work Areas window

Command

From a command line, he issues the following command:

```
teamc workarea -create -name 310 -release robot_control
```

Result

TeamConnection creates a work area named 310 associated with release robot_control. The following parts are currently available in the latest view of release robot_control:

brain.c	leg.c
brain.obj	leg.obj
brain.exe	foot.c
arm.c	foot.obj
arm.obj	optics.c
hand.c	optics.obj
hand.obj	

These parts are also visible in the work area 310 because the work area is associated with the release upon creation, and it contains the latest view of the entire release.

Checking out a part

Alex wants to update a subroutine within optics.c, which controls the aperture of the robot's on-board camera. He checks the part out to start the modifications.

Because Alex knows the exact name of the part, he does one of the following:

GUI

From the GUI, he:

1. Selects **Parts** → **Check out** from the Actions pull-down menu on the Tasks window.
2. Types the following:

- optics.c in the **Path names** field
- robot_control in the **Release** field
- 310 in the **Work area** field

3. Selects **OK**.

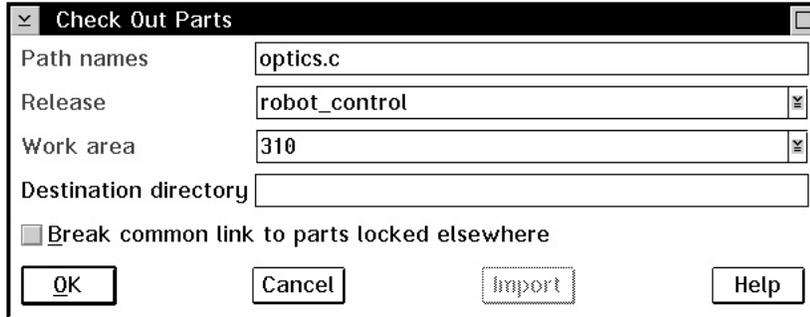


Figure 8. Check Out Parts window

Command

From a command line, he issues the following command:

```
teamc part -checkout optics.c -release robot_control -workarea 310
```

Result

A copy of the part optics.c is checked out from TeamConnection and placed in the directory specified on the Environment page of the Settings notebook of Alex's TeamConnection client. The part, optics.c, is locked. No other user can update the part until Alex integrates his work area with the release.

Searching for a part

Because Alex knows exactly what part he wants to check out, he specifies the name of the part. If he does not know the name, Alex can use the Parts Filter window or the report command to search for the name. He can do one of the following:

GUI

From the GUI, he:

1. Selects **Parts** → **Parts** from the Objects pull-down menu on the Tasks window.
He does not select the **PartFull** choice because he wants to limit his search to a particular release and work area. He uses **PartFull** when he wants to search for parts across releases, components, or work areas.
2. Types the following in the Parts Filter window:

- robot_control in the **Release** field
 - 310 in the **Work area** field
 - % in the **Base names** field and selects **like**
3. Selects **Save to Task List**.

The Part Filter dialog box is shown with the following configuration:

- Release:** robot_control
- Work area:** 310
- Version:** (empty)
- Path names:** in [] No Sort
- Base names:** like % No Sort
- Work areas:** in [] No Sort
- Components:** in [] No Sort
- Current versions:** in [] No Sort

History Table:

Query	Release	Work Area	Version
Show all Parts			

Buttons: OK, Apply, Clear, Save to Task List..., Generate Query, Cancel, Help

Figure 9. Part Filter window

Alex does this because he realizes that he is going to use this query many times, so he wants to add the query to the Tasks window.

4. Adds the necessary information to the Edit Task List window, and selects **Add/Change**.
5. Closes the Edit Task List window. The Tasks window appears.

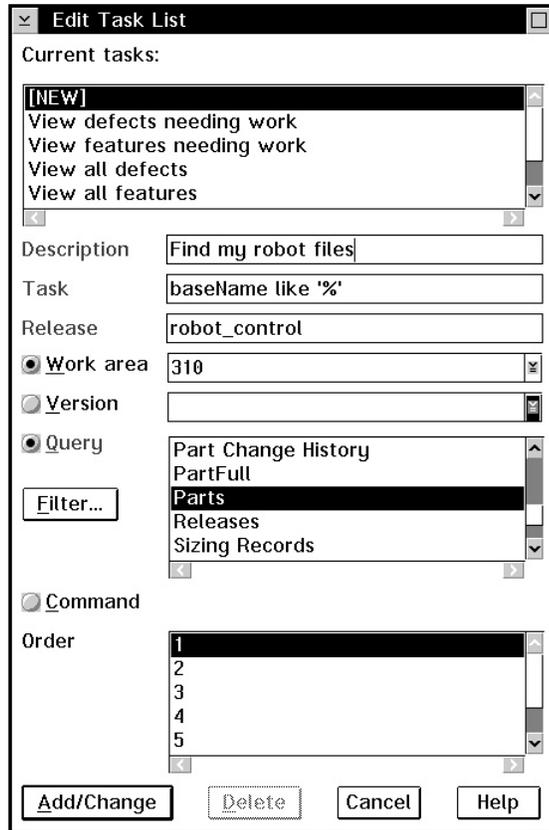


Figure 10. Edit Task List window

6. Double-clicks on the task entry he just created. The Parts window appears. Hereafter, to display the list of parts in his work area, he merely double-clicks on the task entry.
7. Places the mouse pointer over the part name optics.c and presses mouse button 2 to display the pop-up menu.
8. Selects **Check out**. The Check Out Parts window appears with the required fields pre-filled. If Alex provided directory information on the Environment page of the Settings notebook, the **Destination directory** field is pre-filled also.
9. Selects **OK** to check out the part.

Command

From a command line, he issues the following command:

```
teamc report -view partView -where "baseName like '%.c'" -release robot_control
-workArea 310
```

This command returns a list of all the parts that match the query. After Alex determines which part he wants to check out, he issues the following command:

```
teamc part -checkout optics.c -release robot_control -workarea 310
```

Result

A copy of the part optics.c is checked out from TeamConnection and placed in the appropriate directory. The part optics.c is locked. No other user can update the part until Alex integrates the work area with the release.

Checking in a part

Alex edits the part, making the modifications he thinks necessary. Now, he wants to test the modifications. First, he checks the changed part back into his work area.

GUI

From the GUI, he:

1. Selects **Parts** → **Check in** from the Actions pull-down menu on the Tasks window.
2. Types the following in the Check In Parts window, and then selects **OK**:
 - optics.c in the **Path names** field
 - robot_control in the **Release** field
 - 310 in the **Work area** field

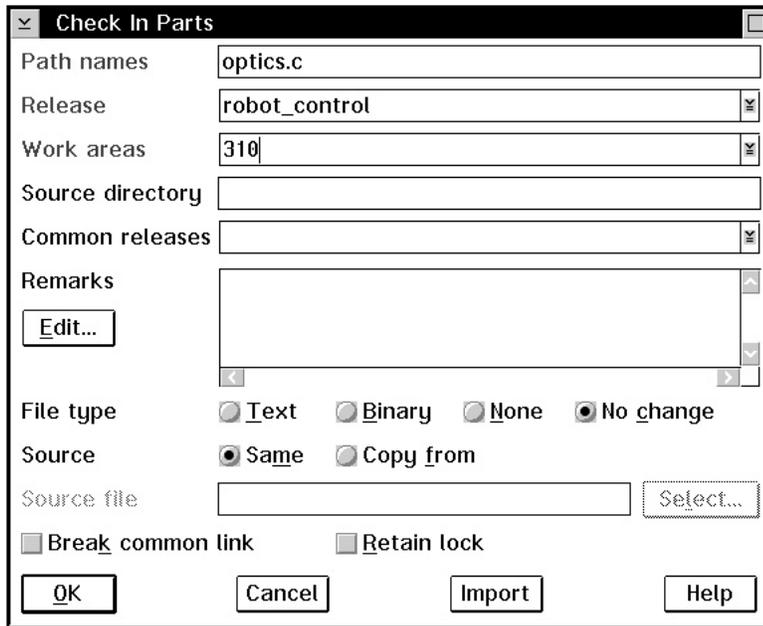


Figure 11. Check In Parts window

Note: Alex follows these steps because he knows the exact name of the part that he is checking in. If he does not know the name, or if he is checking in many parts, he can instead do one of the following to display a list of parts:

- Select the entry on his Tasks window that displays the list of parts.
- Re-open the Parts window if it was previously minimized.
- Add an entry to his Tasks window that lists all of his checked-out parts.

He then selects the parts that he wants to check in.

Command

From a command line, he issues the following command:

```
teamc part -checkin optics.c -release robot_control -workarea 310
```

Result

At this point, it is important to note that the part is checked in to work area 310 and is visible in work area 310 only. The change to optics.c is not visible at the release level or to any other work area. Only the 310 work area contains the change, which is why Alex must specify the work area on the check-out command. Because changes to parts are isolated within work areas, the check-out command must specify which work area to use so that the correct copy of the part is retrieved.

Thus, work area 310 contains the following parts:

brain.c	leg.c
brain.obj	leg.obj
brain.exe	foot.c
arm.c	foot.obj
arm.obj	optics.c (modification 1)
hand.c	optics.obj
hand.obj	

Work area 310 continues to contain the unchanged parts from the requested release view, but now the work area is overlaid with changes local to the work area — `optics.c` in this case. Alex has his own copy of the application that he can modify without impacting other developers. Alex has checked in `optics.c`; however, the modified part remains locked until the work area is integrated with the release.

Verifying and testing part updates

Alex now requests a build of `brain.exe`, the high-level program for the robot control application.

GUI

From the GUI, he:

1. Selects **Parts** → **Build** from the Actions pull-down menu. The Build Parts window appears.
2. Types the following, and then selects **OK** to start the build:
 - `brain.exe` in the **Path name** field
 - `robot_control` in the **Release** field
 - `310` in the **Work area** field
 - `normal` in the **Pool** field

The **Pool** field tells TeamConnection which set of build agents will handle this build. Alex got the name of the pool from his build administrator.

Alex could have selected `brain.exe` from a list of parts on the Parts window, and then selected **Build** from the Selected pull-down menu. This action would have placed some information in the fields, such as the path name and release name.

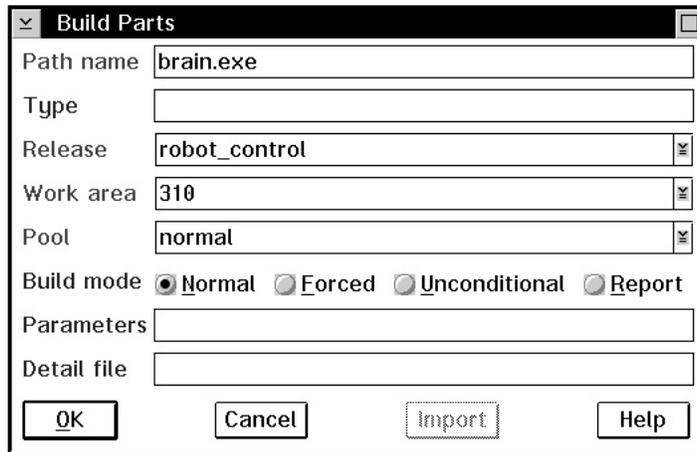


Figure 12. Build Parts window

Command

From a command line, he issues the following command:

```
teamc part -build brain.exe -release robot_control -workarea 310 -pool normal
```

Result

TeamConnection determines the parts that are needed for the build from the set of all the part versions that are currently visible from work area 310. The following part versions are selected for build:

```
brain.c          leg.c
brain.obj        leg.obj
brain.exe        foot.c
arm.c            foot.obj
arm.obj          optics.c (modification 1)
hand.c           optics.obj
hand.obj
```

After the build is complete, TeamConnection stores the resulting outputs of the build in the work area 310. After the build, the work area contains these parts:

```
brain.c          leg.c
brain.obj        leg.obj
brain.exe (contains modification 1) foot.c
arm.c            foot.obj
arm.obj          optics.c (modification 1)
hand.c           optics.obj (modification 1)
hand.obj
```

Note: For a detailed build example, see “Chapter 15. Building an application: an example” on page 181.

Extracting a part

Next, Alex tests his modifications in the robot prototype in his office. He extracts the executable part from the work area 310.

GUI

From the GUI, he:

1. Selects **Parts** → **Extract** from the Actions pull-down menu on the Tasks window.
2. Types the following in the Extract Parts window, and then selects **OK**:
 - brain.exe in the **Path names** field
 - robot_control in the **Release** field
 - 310 in the **Work area** field

Alex does this because he wants to extract the .exe part that is in his work area. If he leaves the **Work area** field blank, he gets the latest committed version of the .exe part from the release.

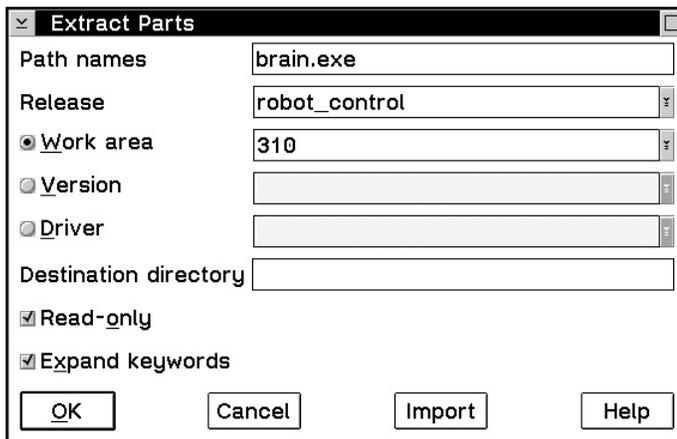


Figure 13. Extract Parts window

Command

From a command line, he issues the following command:

```
teamc part -extract brain.exe -release robot_control -workarea 310
```

Result

This action places a copy of the part brain.exe in the current directory.

Checking out the part one more time

Alex then downloads brain.exe to his robot, runs his test, and determines that the modification did not work: the robot slams into the wall. However, Alex thinks he knows what the problem is, so he needs optics.c for further modifications. First, he checks out the part.

GUI

From the GUI, he:

1. Does one of the following to display the Check Out Parts window:
 - Selects **Parts** → **Check out** from the Actions pull-down menu on the Tasks window.
 - Selects the entry on his Tasks window that displays the list of parts, and then selects the part.
 - Re-opens the Parts window if it was minimized, and then selects the part.
2. When the Check Out Parts window appears, he types the necessary information and selects **OK**.

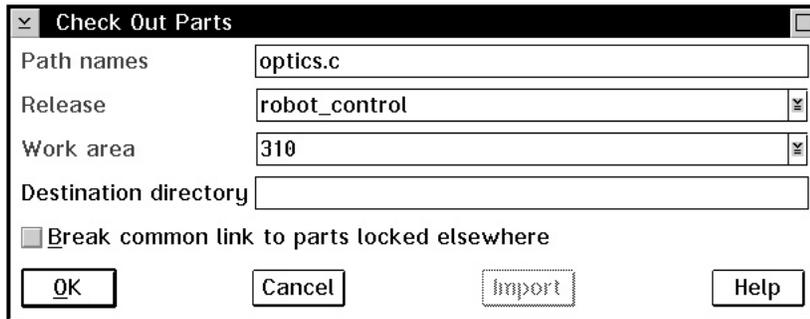


Figure 14. Check Out Parts

Command

From a command line, he issues the following command:

```
teamc part -checkout optics.c -release robot_control -workarea 310
```

Result

A copy of the previously modified optics.c from work area 310 is checked out and placed in the current directory.

Checking the part back in

Alex makes his modification and checks the part in.

GUI

From the GUI, he:

1. Does one of the following to display the Check In Parts window:
 - Selects **Parts** → **Check in** from the Actions pull-down menu on the Tasks window.
 - Selects the entry on his Tasks window that displays all the parts he has checked out, and then selects the part.
 - Re-opens the Parts window if it was minimized, and then selects the part.
2. When the Check In Parts window appears, he types the necessary information and selects **OK**.

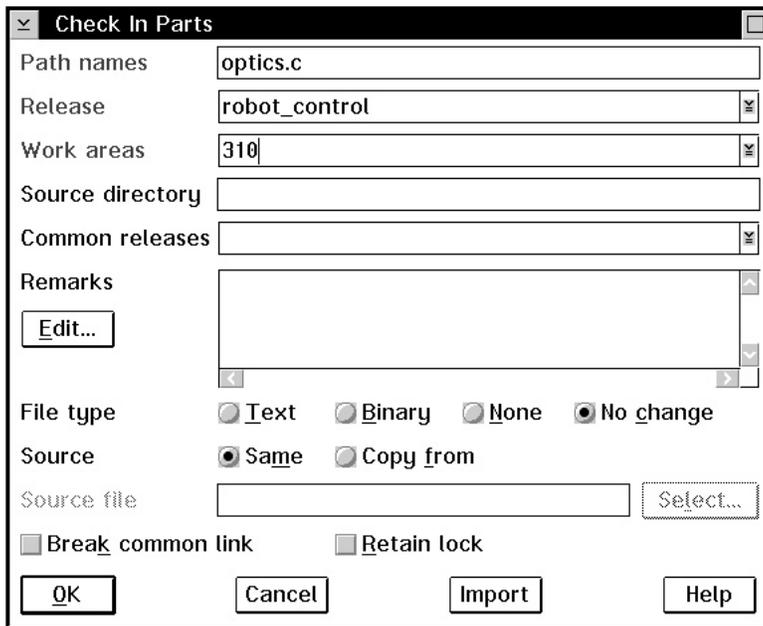


Figure 15. Check In Parts window

Command

From a command line, he issues the following command:

```
teamc part -checkin optics.c -release robot_control -workarea 310
```

Result

Now the work area contains the following parts:

brain.c	leg.c
brain.obj	leg.obj
brain.exe (contains modification 1)	foot.c
arm.c	foot.obj
arm.obj	optics.c (modification 2)
hand.c	optics.obj (modification 1)
hand.obj	

Because Alex did not specify that he wanted to save a copy of the work area by freezing it, optics.c (modification 1) was overwritten.

Freezing the work area

Alex builds the application again, extracts the executable part, and runs his test. This time, everything works, and the robot successfully finds its way to the snack machine down the hall without hitting anything. Alex is very pleased, but he notices an unrelated problem in the robot's autofocus system. Before Alex begins repairing the autofocus subroutine, he wants to save a copy of the application as it exists now in his work area. So, Alex does one of the following to freeze the work area:

GUI

From the GUI, he:

1. Displays the Freeze Work Areas window in one of the following ways:
 - Selects **Work areas** → **Freeze** on the Actions pull-down menu from the Tasks window.
 - Selects 310 from the list of work areas on the Work Areas window, then selects **Freeze** from the Selected pull-down menu.
2. Types 310 in the **Work areas** field and robot_control in the **Releases** field if the data is not already in the fields.
3. Selects **OK**.

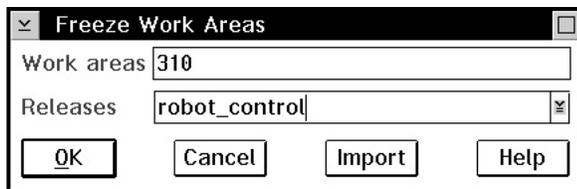


Figure 16. Freeze Work Areas window

Command

From a command line, he issues the following command:

```
teamc workarea -freeze 310 -release robot_control
```

Result

The freeze command saves the work area 310. Thus, TeamConnection takes a snapshot of the work area, with all its parts and their visible versions, and saves it. Alex can come back to this stage of development in the work area if he wants. Note, however, that a freeze action does not make the changes visible to the other people working in the release, nor does it unlock the parts.

Refreshing the work area

Alex finally finishes his work on the robot's optical systems after making three additional attempts at modifying `optics.c` and rebuilding the application. Alex modified and rebuilt the application a total of five times in the work area. Now, he wants to share his work with the rest of the team. His work area currently contains the following parts:

<code>brain.c</code>	<code>leg.c</code>
<code>brain.obj</code>	<code>leg.obj</code>
<code>brain.exe</code> (contains modification 5)	<code>foot.c</code>
<code>arm.c</code>	<code>foot.obj</code>
<code>arm.obj</code>	<code>optics.c</code> (modification 5)
<code>hand.c</code>	<code>optics.obj</code> (modification 5)
<code>hand.obj</code>	

While Alex worked in his work area, other members of the team were working on their own modifications. Some of these modifications have been integrated with the release, so the copy of the release that Alex has is probably stale. If he were to integrate his changes at this time with the release, he might cause the application to break.

Alex first refreshes his work area with parts from the release by doing one of the following:

GUI

From the GUI, he:

1. Selects **Work areas** → **Refresh** from the Actions pull-down menu on the Tasks window.
2. Types 310 in the **Work areas** field and `robot_control` in the **Releases** field.
Alex wants to refresh from the release, so he does not specify a source.
3. Selects **OK**.

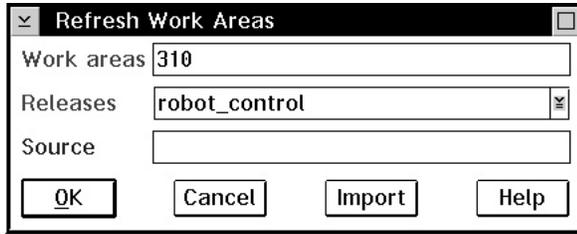


Figure 17. Refresh Work Areas window

Command

From a command line, he issues the following command:

```
teamc workarea -refresh 310 -release robot_control
```

Result

This action updates work area 310 with any changes from the release, and it also freezes work area 310, if it is not already frozen. Now Alex's work area contains the following versions of parts:

brain.c (Jenny's modification)	leg.c
brain.obj (from Alex's build after refresh)	leg.obj
brain.exe (contains modification 5)	foot.c
arm.c	foot.obj
arm.obj	optics.c (modification 5)
hand.c (Joy's and Ken's modification)	optics.obj (modification 5)
hand.obj (from Alex's build after refresh)	

None of the objects that Alex modified and none of the objects built as a result of Alex's modifications is overwritten by the refresh.

Building the application

Alex again builds the application brain.exe within his work area to determine whether his changes integrate with Jenny's, Joy's, and Ken's modifications.

GUI

Alex has a Parts window open with a list of all the parts that exist in work area 310. He highlights the part brain.exe, and then does the following:

1. Selects **Build** from the Selected pull-down menu.
2. Types normal in the **Pool** field. The other required fields have the correct information.
3. Selects **OK** to start the build.

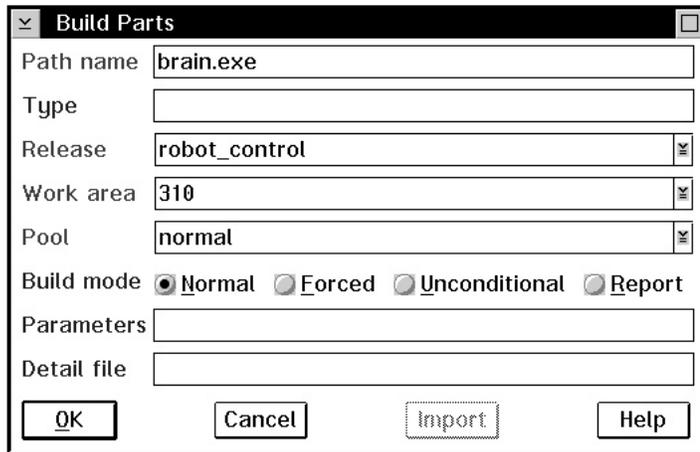


Figure 18. Build Parts window

Command

From a command line, he issues the following command:

```
teamc part -build brain.exe -release robot_control -workarea 310 -pool normal
```

Result

Fortunately, nothing breaks, so Alex is ready to integrate his changes with the release.

Integrating the work area

To integrate his changes with the release, Alex must integrate the work area he has been using with the release. This will make the work area visible to all the users in the release. He does one of the following:

GUI

From the GUI, he:

1. Selects **Work areas** → **Integrate** from the Actions pull-down menu on the Tasks window. The Integrate Work Areas window appears.
2. Types 310 in the **Work areas** field and robot_control in the **Releases** field.
3. Selects **OK**.

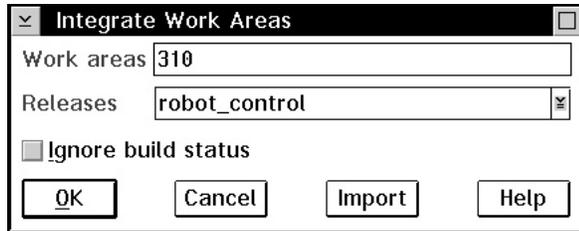


Figure 19. Integrate Work Areas window

Command

From a command line, he issues the following command:
teamc workarea -integrate 310 -release robot_control

Result

TeamConnection first determines that Alex's changes were built against the latest version of the release. Then TeamConnection makes Alex's changes visible at the release level so that the other team members can see and use them. The following part versions are now visible from the release:

```
brain.c (Jenny's modification)
brain.obj (from Jenny's build)
brain.exe (from Alex's build)
arm.c
arm.obj
hand.c (Joy's modification, Ken's modification)
hand.obj (from Ken's build)
leg.c
leg.obj
foot.c
foot.obj
optics.c (Alex's modification 5)
optics.obj (from Alex's build)
```

TeamConnection also makes a copy of the release before integrating Alex's changes. If something doesn't work, the users or the administrator can go back to the release prior to Alex's integration. The part, optics.c, is now unlocked in the release. The work area is now in the complete state and can no longer be used.

Closing a defect

Now that Alex is finished making changes to fix the problem reported in defect 310, he is ready to close the defect. He does one of the following:

GUI

From the GUI, he:

1. Selects **Defects** → **Verify** from the Actions pull-down menu on the Tasks window.
The Verify Defects window appears.
2. Types 310 in the **Defects** field.
3. Selects **OK**.

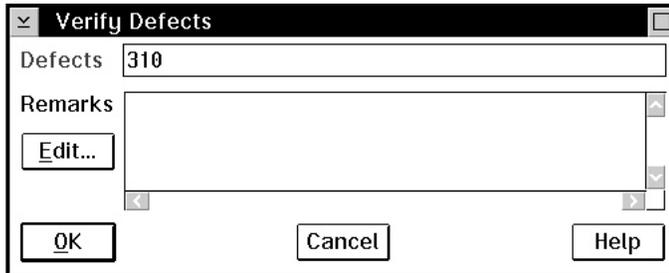


Figure 20. Verify Defects window

Command

From a command line, he issues the following command:

```
teamc defect -verify 310 -release robot_control
```

Result

Because the component does not include the verifyDefect subprocess in its process, the defect moves directly to the closed state.

Working in concurrent development

The previous section discussed working in a serial development environment. While Alex had optics.c in his work area, no one else on the team could check out the part. TeamConnection allows you to hold the part until you are sure that it integrates with the rest of the application. Therefore, the lock is not released until the work area as a whole is integrated with the release.

The scenario changes slightly for concurrent development. In this case, several users can work on the same part at the same time. These users must reconcile their changes as they integrate their work areas with the release.

The following tasks are required:

For information about this task,	Go to this page.
Refreshing the work area from the driver	71
Integrating the work area	72
Resolving differences	73

Refreshing the work area from the driver

If Alex and Jenny are working on `optics.c` at the same time, they must resolve their part differences at some point, because both want to make their changes visible to the release. If Alex and Jenny were not required to do this before committing their work areas, the last developer to commit would always overlay the other's changes. For this scenario, assume that Jenny finishes her changes first. The first thing she does is refresh her work area.

GUI

From the GUI, she:

1. Selects **Work areas** → **Refresh** from the Actions pull-down menu on the Tasks window.
2. Types 415 in the **Work areas** field.
3. Types `robot_control` in the **Releases** field. Jenny wants to refresh from the driver.
4. Types driver name in the **Source** field.
5. Selects **OK**.

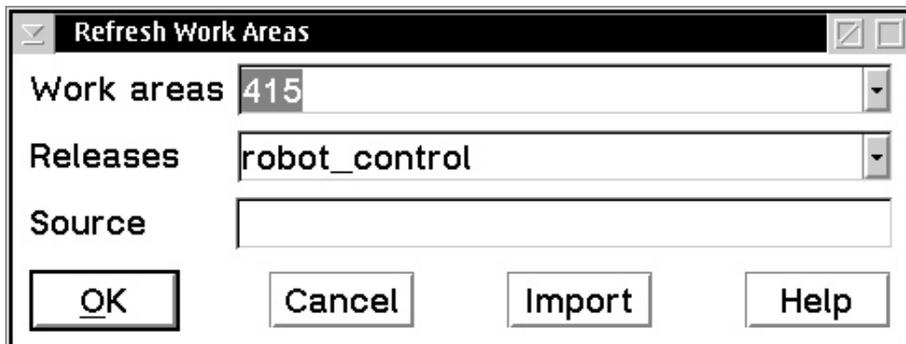


Figure 21. Refresh Work Areas window

Command

From a command line, she issues the following command:

```
teamc workarea -refresh 415 -release robot_control
```

Result

This command refreshes her work area with the latest view of the release. Her work area now contains the following part versions:

```
brain.c (Jenny's modification 3)
brain.obj (Jenny's modification 3)
brain.exe (has Jenny's brain.c modification 3 and optics.c modification 4)
arm.c
arm.obj
hand.c (Joy's modification, Ken's modification)
hand.obj (Joy's modification, Ken's modification)
leg.c
leg.obj
foot.c
foot.obj
optics.c (Jenny's modification 4)
optics.obj (Jenny's modification 4)
```

Integrating the work area

The refresh shows Jenny only the parts integrated with the release. She does not see Alex's work because he has not integrated his work area yet. Jenny rebuilds the application, tests it, and decides she is ready to integrate her changes. She does one of the following:

GUI

From the GUI, she:

1. Selects **Records** → **Fix records** → **Complete** from the Actions pull-down menu on the Tasks window. The Complete Fix Records window appears.
2. Types 415 in the **Work areas** field, robot_control in the **Releases** field, and robot_component in the **Component** field.
3. Selects **OK**.

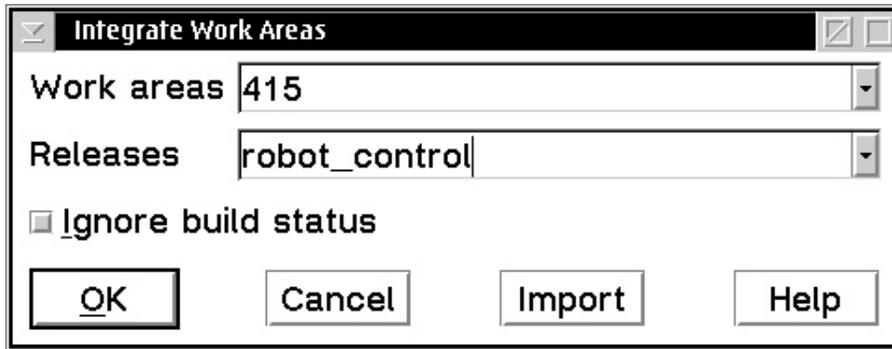


Figure 22. Integrate Work Areas window

Command

From a command line, she issues the following command:

```
teamc workarea -integrate 415 -release robot_control
```

Result

Because Jenny is up-to-date with the latest view of the driver, her changes are integrated after TeamConnection preserves a copy of the previous version of the release.

Reconciling differences

Later, Alex is ready to integrate his modifications. Alex issues a refresh command from the driver, as Jenny did (see page 71 for instructions).

This time, Alex receives a message that collision records were generated, because both he and Jenny have updated the same parts. At this time he does not know which parts collided. TeamConnection refreshes work area 310 with the exception of the part optics.c, which had the collision. Alex's work area shows the following parts:

```
brain.c (Jenny's modification 3)
brain.obj (Jenny's modification 3)
brain.exe (Contains Alex's modification 5)
arm.c
arm.obj
hand.c (Joy's modification, Ken's modification)
hand.obj (Joy's modification, Ken's modification)
leg.c
leg.obj
foot.c
foot.obj
optics.c (Alex's modification 5)
optics.obj (Alex's modification 5)
```

Alex can use either the GUI or the command line to reconcile the differences. Four steps are required from the command line:

1. Check out the latest uncommitted version.
2. Extract the latest committed version.
3. Run the merge program against the two parts.
4. Check in the resultant part.

However, on the GUI the reconcile action automatically does the preceding steps for you, which can save you a considerable amount of work if several parts require reconciliation.

GUI

From the GUI, he:

1. Selects **Parts** → **Collision Records** from the Objects pull-down menu. The Collision Record Filter window appears.
2. Types 310 in the **Work areas** field and selects **OK**. The Collision Records window appears with optics.c listed as the part having the collision.
3. Highlights the optics.c entry and selects **Reconcile** from the pop-up menu. The Reconcile Collision Record window appears with the required information pre-filled.
Alex does not have to reconcile every part for which a collision record is created. He can choose either his copy or the copy at the release rather than combining the two. For example, if Alex wants to use his copy of optics.c without merging with the copy at the release level, he selects the reject action (of course, he would not do that without first talking with Jenny). If he wants to use the copy of optics.c at the release level without merging any of his changes into the copy at the release level, he selects the accept action.
4. Because Alex wants to combine the two sets of changes, he selects **Merge** to start the TeamConnection merge program, or any merge program of his choice. Alex merges the changes and then saves and exits from the merge program. TeamConnection checks the resultant part back in as part of this merge step.
The online help provides information on how to use the merge program.
5. Selects **OK** from the Reconcile Collision Record window.

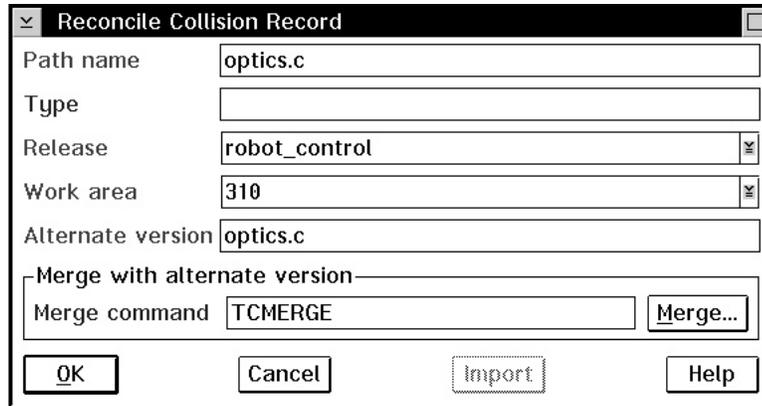


Figure 23. Reconcile Collision Record window

Command

From a command line, he does the following steps:

- Issues a report command to determine which parts are in conflict:

```
teamc report -view collisionView -workarea 310
```

This report tells him that `optics.c` is the part that collided and gives the *alternate version ID* of the part that caused the collision. Alex makes note of the alternate version ID, `robot_control:2`, because he needs to specify that in a later step.

- Extracts a copy of `optics.c` from the release:

```
teamc part -extract optics.c -release robot_control -relative d:\temp
```

By not specifying a work area on the part `-extract` command, Alex ensures that he receives the last committed copy of the part at the release. Also, Alex specifies a *relative* path for the part extract. By specifying the relative directory, he prevents TeamConnection from placing the part in his default directory, where he normally works on checked-out parts. For more information about the `-relative` flag, refer to the *Commands Reference*

- Checks out his copy of `optics.c` from his work area:

```
teamc part -checkout optics.c -release robot_control -workarea 310
```

Because he did not specify a relative path, this part is checked out to his working directory `d:\robot`.

- Uses the merge program to reconcile the two copies of `optics.c`:

```
tcmerge d:\temp\optics.c d:\robot\optics.c -out  
d:\robot\optics.c -prime d:\temp\optics.c
```

If Alex decides not to merge the two parts, but instead wants to use his copy of optics.c, he uses the collision -reject command. Or, if he wants to use the copy of optics.c at the release level, he uses the collision -accept command.

- Checks the resultant copy of optics.c into his work area and builds it against the rest of the system.
- After he is satisfied with the reconciled changes, he lets TeamConnection know that the previously discovered conflict is reconciled. Alex does this by completing the collision record that TeamConnection created when Alex attempted to integrate his copy of optics.c. He does the following:

```
teamc collision -reconcile -path optics.c -release robot_control  
-workarea 310 -altversion robot_control:2
```

Result

Alex is now ready to make his changes visible to the release. He can use either the GUI or the command line to integrate the work area.

He refreshes from the driver again. The integrate is permitted because a completed collision record exists for the conflict between the two versions of optics.c. However, if Ken or Joy had integrated a new version of optics.c while Alex was busy resolving the last collision, Alex's driver add would fail. He would have to repeat the collision resolution process.

Chapter 6. Working with component and release processes

The previous chapter described how to work with parts when the release does not follow a tracking process. This chapter describes how to work with parts when a tracking process *is* followed and how to use component processes for features and defects.

When tracking is part of the process, users must associate any changes to their parts with the defects or features active for the release. This association is made through a work area. The work area is the object that ties a defect or feature with a specific release. When checking out a part, the user must specify the work area with which the modification is associated. For any release and defect or feature pair, there can be multiple work area objects.

Aside from their association with a defect or feature, the work areas for a full-tracking process environment are identical to those defined for working in a no-tracking process environment. Work areas maintain a separate view for the user working on the modifications associated with a defect or feature without affecting the release. This view can be integrated with the release at some point. A work area is implicitly created when a defect or feature is accepted if the managing component follows a design, size, and review process for defects and features and if a sizing record is created. The work area that TeamConnection creates is based on the sizing record and has the same name as the defect or feature. If sizing records were not created, you must explicitly create the work area.

As an example of how this all works, suppose that the robot project from the previous chapter is entering system test. The administrator decides to turn on a full-tracking process for the release, such as `track_full`. This process includes the track, approval, fix, driver, and test subprocesses. The release follows concurrent development, and the component follows a design, size, and review process for both defects and features.

On a weekly basis the project leader, Carol, creates a driver. A *driver* monitors and implements the integration of part changes within a release. These part changes are included in a driver by adding the work areas referenced by the changed parts to the driver as *driver members*.

One of the testers for the robot project discovers that the autofocus mechanism in the robot's eye fails when the robot is placed in front of striped wallpaper. The tester must open a defect against the component optics, which is owned by Carol. Carol verifies that the problem does exist, accepts the defect, and assigns it to Alex. This fix will require the tasks noted in the following table:

For information about this task,	Go to this page.
Changing the defect owner	79
Accepting the defect	80
Approving the fix	81

For information about this task,	Go to this page.
Checking out a part	82
Checking in the changes	83
Freezing the work area	85
Building the application	86
Accepting fix records	87
Adding a driver member	88
Returning the work area to the fix state	89
Reactivating the fix record	90
Refreshing the work area	90
Refreshing the driver	91
Building the driver	92
Restricting the driver	93
Integrating the parts	94
Completing the driver	95
Testing the built application	95

Moving through design, size, and review

Because the defect was created against a component that follows the design, size, and review process for defects, Carol must move the defect through this process before the defect can be accepted and parts can be checked out. As the names imply, the process requires that the following be done:

- Design what needs to be done in order to resolve the problem. She must enter design text before the defect can move to the size state.
- Size the amount of work that is required to resolve the problem. At this time, Carol creates a sizing record and specifies robot_control as the release that contains the parts that require changing. If parts in other releases require changing because of the defect, a sizing record is created for each release. A sizing record assures that a work area is created when the defect is accepted. It identifies the work that is required for and the resources affected by the defect or feature. The owner of the component that is referenced in the sizing record is the owner of the sizing record. The owner is responsible for entering information about the amount of work that is required to implement the feature or resolve the problem.
- Review all design text and sizing records and determine if work should continue on the defect.

Changing defect ownership

Because Carol is the component owner, she is currently defined as the owner of defect 456. But the problem is in Alex's code, so she wants him to own the defect. To reassign ownership, she does one of the following:

GUI

From the GUI, she:

1. Selects **Defects** → **Modify** → **Owner** from the Actions pull-down menu on the Tasks window. The Modify Defect Owner window appears.
2. Types 456 in the **Defects** field and types Alex's user ID, alexm, in the **New owner** field.
3. Selects **OK**.

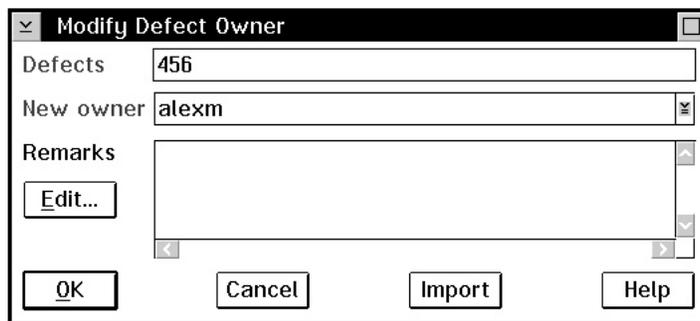


Figure 24. Modify Defect Owner window

Command

From a command line, she issues the following command:

```
teamc defect -assign 456 -owner alexm
```

Results

Alex is now the owner of defect 456. He is responsible for fixing the problem and moving the defect through its various states.

Accepting a defect

When you accept a defect or feature, you accept the responsibility of resolving it. A defect or feature might require changes in more than one release. If the component includes the design, size, and review process, these releases were identified during the size state, and TeamConnection created a work area for each identified release. If the component does not include the design, size, and review process, you will need to create a work area manually.

When the first work area moves to the complete state, the defect or feature automatically moves to the verify state or closed state.

Alex, now the owner of the defect, accepts the defect by doing one of the following:

GUI

From the GUI, he:

1. Selects **Defects** → **Accept** from the Actions pull-down menu on the Tasks window. The Accept Defects window appears.

Note: The Accept Defects window in this example may be different than one you may see based on your environment. Configurable fields may or may not be shown depending on any configurable fields set by you or your administrator,

2. Types 456 in the **Defects** field and selects **program_defect** from the **Answer** list.
3. Selects **OK**.

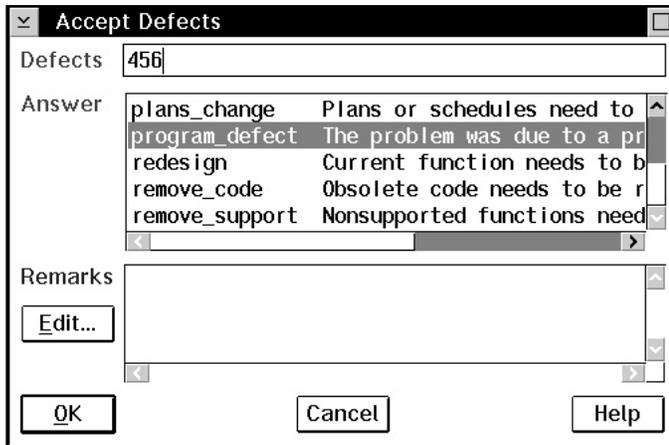


Figure 25. Accept Defects window

Command

From a command line, he issues the following command:

```
teamc defect -accept 456 -answer program_defect
```

Results

Defect 456 moves to working state, and TeamConnection creates a work area called 456. The work area is associated with the release specified on the sizing record, which in this example is robot_control. When the work area is created, a fix record is also created based on the sizing record. Because the approval subprocess is included in the release's process, the work area is created in the approve state and the fix record is created in the notReady state.

Just as with a work area that is explicitly created, the defect work area contains a view of the current versions visible to the release. In this case, the contents of the work area are:

brain.c	leg.c
brain.obj	leg.obj
brain.exe	foot.c
arm.c	foot.obj
arm.obj	optics.c
hand.c	optics.obj
hand.obj	

Approving the fix

Because the full-tracking process includes the approval subprocess, each person identified on the approval list must approve the proposed changes before Alex can begin work on the defect.

Linda and Sam are both listed as approvers. They have been notified by TeamConnection that they have approval records. After reviewing the defect, they do one of the following to indicate their approval:

GUI

From the GUI, they:

1. Select **Records** → **Approval records** → **Accept** from the Actions pull-down menu.
2. Type 456 in the **Work areas** field and robot_control in the **Release** field.
3. Select **OK**.

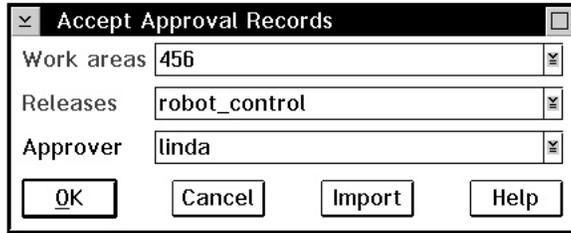


Figure 26. Accept Approval Records window

Command

From a command line, they both issue the following command for the approval record that they have:

```
teamc approval -accept -workarea 456 -release robot_control
```

Results

After both Linda and Sam accept the approval records, TeamConnection moves the work area to the fix state.

Checking out a part

Now that the approval records have been accepted, Alex can check out the necessary parts. He decides that modifications are again required to the part `optics.c`. So, that is the part he checks out.

Alex must specify the work area on the check-out command so that the part is obtained from the defect's work area. He does one of the following:

GUI

From the GUI, he:

1. Selects **Parts** → **Check out** from the Actions pull-down menu on the Tasks window.
2. Types the following:
 - `optics.c` in the **Path names** field
 - `robot_control` in the **Release** field
 - `456` in the **Work area** field
 - `d:\robot\src` in the **Destination directory** field
3. Selects **OK**.

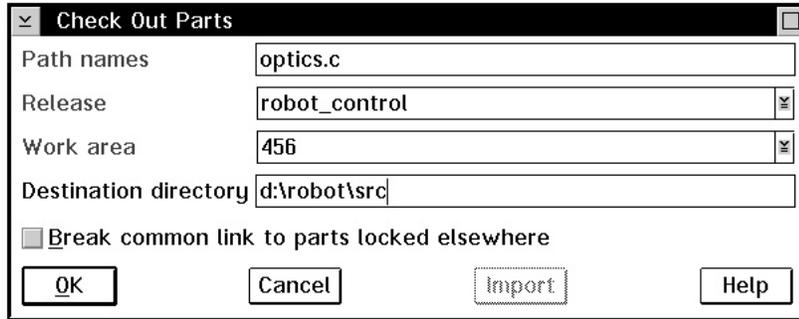


Figure 27. Check Out Parts window

Command

From a command line, he issues the following command:

```
teamc part -checkout optics.c -release robot_control -workarea 456
  -relative d:\robot\src
```

Results

A copy of the part `optics.c` is checked out from TeamConnection and placed in the directory `d:\robot\src`. If the directory name is not specified in the command, TeamConnection uses the directory specified in the `TC_RELATIVE` environment variable. Because the release is following concurrent development mode, other users can also check out and change this part while Alex has it checked out.

Checking in the changes

Alex makes his modifications and wants to test his corrections. First, he must check the part into the work area. He does one of the following:

GUI

From the GUI, he:

1. Selects **Parts** → **Check in** from the Actions pull-down menu on the Tasks window.
2. Types the following in the Check In Parts window, and then selects **OK**:
 - `optics.c` in the **Path names** field
 - `robot_control` in the **Release** field
 - `456` in the **Work areas** field
 - `d:\robot\src` in the **Source directory** field

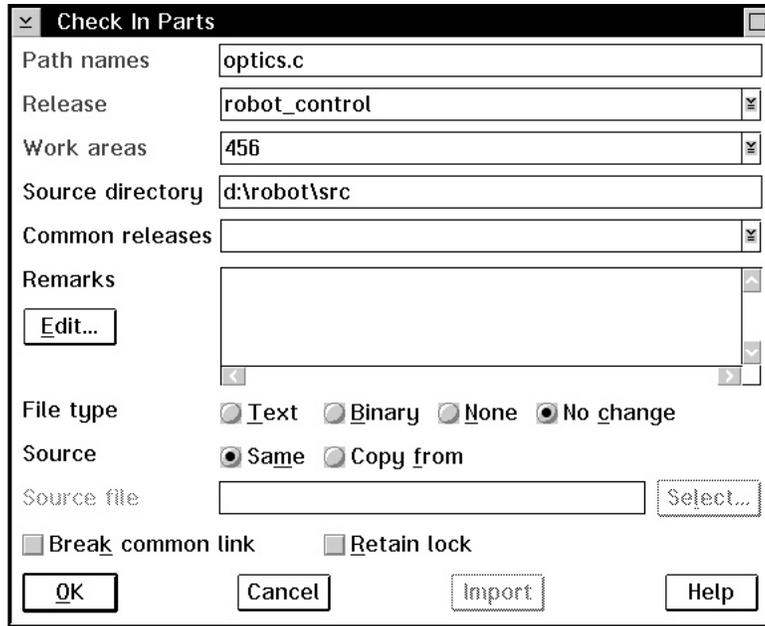


Figure 28. Check In Parts window

Note: Alex follows these steps because he knows the exact name of the part that he is checking in. If he does not know the name, or if he is checking in many parts, he can instead do one of the following to display a list of parts:

- Select the entry on his Tasks window that displays the list of parts.
- Re-open the Parts window if it was previously minimized.
- Add an entry to his Tasks window that lists all of his checked-out parts.

He then selects the parts that he wants to check in.

Command

From a command line, he issues the following command:

```
teamc part -checkin optics.c -release robot_control -workarea 456
```

Results

Now the work area contains the following part versions:

```
brain.c           leg.c
brain.obj         leg.obj
brain.exe        foot.c
```

```
arm.c          foot.obj
arm.obj       optics.c (Alex's modification 1)
hand.c        optics.obj
hand.obj
```

Freezing the work area

Alex now wants to save, or freeze, the working system. He does one of the following:

GUI

From the GUI, he:

1. Displays the Freeze Work Areas window in one of the following ways:
 - Selects **Work areas** → **Freeze** from the Actions pull-down menu on the Tasks window.
 - Selects **Work areas** from the Objects pull-down menu on the Tasks window. Types the appropriate search information on the Work Area Filter window to get a list of work areas. Selects 456 from the list of work areas on the Work Areas window, and then selects **Freeze** from the Selected pull-down menu. This method is useful when you are going to be working with several work areas or you are unsure of the work area name.
2. Types 456 in the **Work areas** field and robot_control in the **Releases** field if the information is not already there.
3. Selects **OK**.

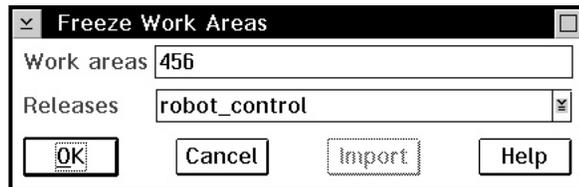


Figure 29. Freeze Work Areas window

Command

From a command line, he issues the following command:

```
teamc workarea -freeze 456 -release robot_control
```

Results

The freeze command saves the work area 456. Thus, TeamConnection takes a snapshot of the work area, with all its parts and their visible versions, and saves it. Note, however, that a freeze action does not make the changes visible to the other people working in the release. This does not occur until the work area is integrated.

Building the application

Alex now builds the application to verify that the changes he has made have fixed the problem. He does one of the following:

GUI

From the GUI:

Alex has a Parts window open with a list of all the parts that exist in work area 456. He highlights the part brain.exe and then does the following:

1. Selects **Build** from the Selected pull-down menu.
2. Types normal in the **Pool** field. The other required fields are pre-filled with the correct information.
3. Selects **OK** to start the build.

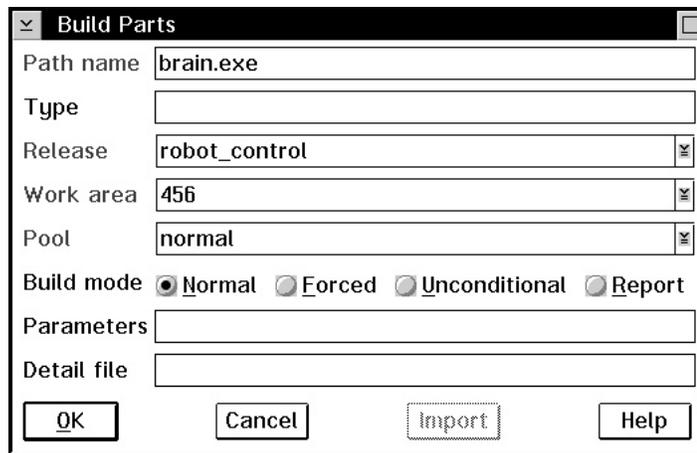


Figure 30. Build Parts window

Command

From a command line, he issues the following command:

```
teamc part -build brain.exe -release robot_control -workarea 456 -pool normal
```

Results

Alex builds the application and tests the results. The modification seems to solve the problem.

Note: For a detailed build example, see “Chapter 15. Building an application: an example” on page 181.

Accepting fix records

Alex is satisfied that the changes are complete and the part is ready to be integrated with other parts in the release. He does one of the following:

GUI

From the GUI, he:

1. Selects **Records** → **Fix records** → **Complete** from the Actions pull-down menu on the Tasks window.
2. Types the following in the Complete Fix Records window, and then selects **OK**:
 - 456 in the **Work areas** field
 - robot_control in the **Releases** field
 - optics in the **Component** field

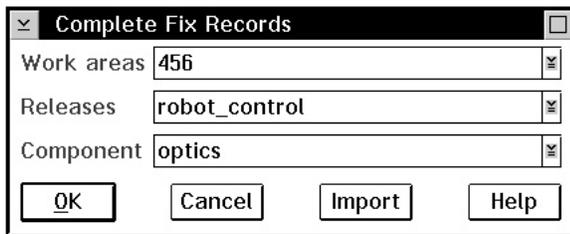


Figure 31. Complete Fix Records window

Command

From a command line, he issues the following command:

```
teamc fix -complete -workarea 456 -component optics -release robot_control
```

Results

The fix record moves to the complete state. Because only one fix record was generated for this defect, the work area moves to the integrate state at the same time. When more than one fix record exists, they all must be completed before the work area moves to the integrate state.

Integrating changed parts into a release

The changes that Alex has made are now ready to be put into the next set of changes scheduled to be integrated with the release. This set of changes is known as a *driver*.

A driver named 0105 currently exists, and several driver members have already been added to the driver. Therefore, the driver is in the integrate state.

Adding a driver member

Carol, the project lead, adds work area 456 as a *driver member* of driver 0105:

GUI

From the GUI, she:

1. Selects **Drivers** → **Add driver members** from the Actions pull-down menu on the Tasks window.
2. Types the following:
 - 0105 in the **Driver** field
 - robot_control in the **Release** field
 - 456 in the **Work areas** field
3. Selects **OK**.

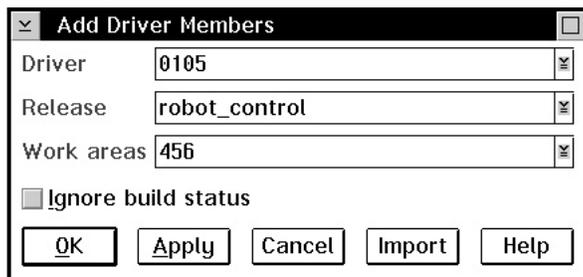


Figure 32. Add Driver Members window

Command

From a command line, she issues the following command:

```
teamc driverMember -create -driver 0105 -workarea 456 -release robot_control
```

Results

Carol previously created a driver member for driver 0105 that included changes to optics.c, so Carol is notified that collisions were detected. (Remember, the release is in concurrent development mode.)

Carol deletes the driver member for work area 456. She then asks Alex to reconcile the collisions.

Reconciling the differences

Before Alex can reconcile the differences, he needs to do the following:

1. Return the work area to the fix state
2. Reactivate the fix record
3. Refresh his work area

Returning the work area to the fix state

The first step in reconciling the differences is for Alex to return work area 456 to the fix state. He does one of the following:

GUI

From the GUI, he:

1. Selects **Work area** → **Fix** from the Actions pull-down menu on the Tasks window.
2. Types 456 in the **Work areas** field and robot_control in the **Releases** field.
3. Selects **OK**.

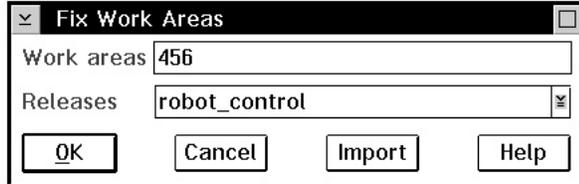


Figure 33. Fix Work Areas window

Command

From a command line, he issues the following command:

```
teamc workarea -fix 456 -release robot_control
```

Results

Work area 456 is in the fix state. After the fix record is reactivated, Alex will check out optics.c from this work area to reconcile the differences.

Reactivating the fix record

Currently, the fix record for work area 456 is in the complete state. Alex must reactivate the fix record to move it back to the active state so that he can make the necessary changes to optics.c. He does one of the following:

GUI

From the GUI, he:

1. Selects **Records** → **Fix records** → **Activate** from the Actions pull-down menu on the Tasks window.
2. Types 456 in the **Work areas** field and selects robot_control from the **Releases** field and optics from the **Component** field.
3. Selects **OK**.

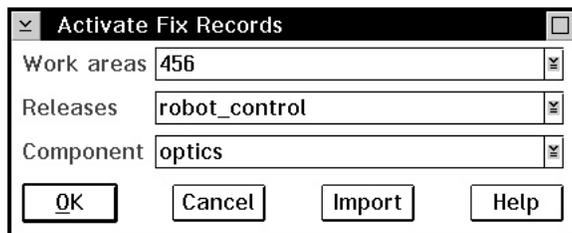


Figure 34. Activate Fix Records window

Command

From a command line, he issues the following command:

```
teamc fix -activate 456 -release robot_control -component optics
```

Results

The fix record returns to the active state.

Refreshing the work area

Alex now needs to refresh his work area with the parts that are already in driver 0105. He does one of the following:

GUI

From the GUI, he:

1. Selects **Work areas** → **Refresh** from the Actions pull-down menu on the Tasks window.

2. Types the following in the Refresh Work Areas window and selects **OK**:
 - 456 in the **Work areas** field
 - robot_control in the **Releases** field
 - 0105 in the **Source** field

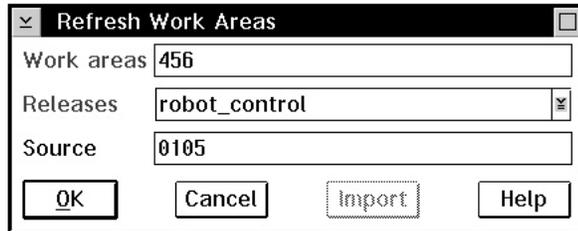


Figure 35. Refresh Work Areas window

Command

From a command line, he issues the following command:

```
teamc workarea -refresh 456 -release robot_control -source 0105
```

Results

TeamConnection notifies Alex of the collision, so his next step is to reconcile the differences. He follows the same procedure that is described on page 73.

Alex completes the fix record and then tells Carol that he has reconciled the part differences and that she can now create the driver member. She creates the driver member without any collisions this time.

Refreshing the driver

Carol is ready to integrate the changes in driver 0105 with the release. Because other team leads have integrated changes as well, she wants to build her driver with the most current release part versions. She does one of the following:

GUI

From the GUI, she:

1. Selects **Drivers** → **Refresh** from the Actions pull-down menu on the Tasks window.
2. Types 0105 in the **Drivers** field and robot_control in the **Release** field.
3. Selects **OK**.

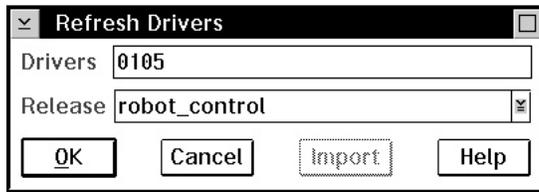


Figure 36. Refresh Drivers window

Command

From a command line, she issues the following command:

```
teamc driver -refresh 0105 -release robot_control
```

Results

This command refreshes driver 0105 with any committed updates to the release.

Building the driver

Carol builds the application using the parts current to driver 0105. She does one of the following:

GUI

From the GUI, she:

1. Selects **Build** from the Action pull-down menu on the Tasks window.
2. Types the following in the Build Parts window:
 - brain.exe in the **Path name** field.
 - robot_control in the **Release** field.
 - 0105 in the **Work area** field.
 - normal in the **Pool** field.
3. Selects **OK** to start the build.

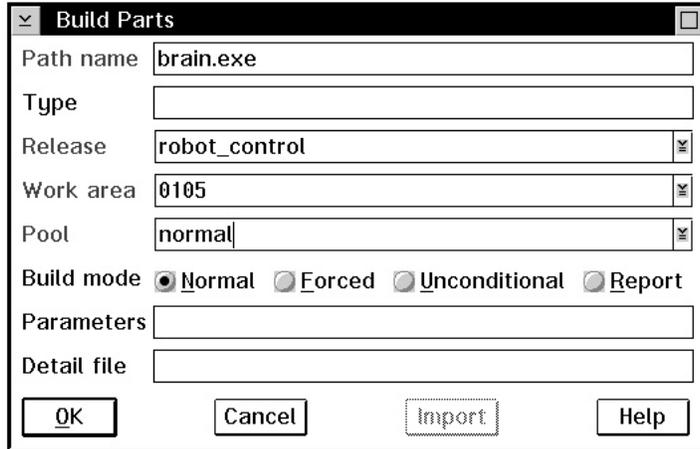


Figure 37. Build Parts window

Command

From a command line, she issues the following command:

```
teamc part -build brain.exe -release robot_control -workarea 0105 -pool normal
```

Results

Carol runs some simple regression tests to verify that the application built properly. She is satisfied with the results, and is ready for the next step — committing the driver changes to the release.

Restricting the driver

After all changes have been integrated with the release, Carol needs to make some final changes before building the driver. To enable her to make these changes while protecting the driver from access by anyone else, she needs to restrict access to it. She does one of the following:

GUI

From the GUI, she:

1. Selects **Drivers** → **Restrict** from the Actions pull-down menu on the Tasks window.
2. Types 0105 in the **Drivers** field and robot_control in the **Release** field.
3. Selects **OK**.

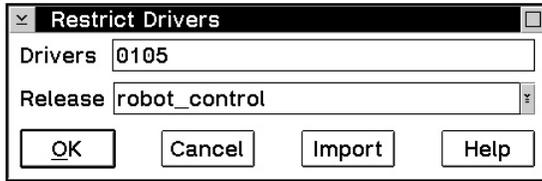


Figure 38. Restrict Drivers window

Command

From a command line, she issues the following command:

```
teamc driver -restrict 0105 -release robot_control
```

Results

This command restricts driver 0105 so that only Carol is able to make changes to it. Carol is now ready to build the application.

Integrating the parts

Carol commits the changes in the driver to the release by doing one of the following:

GUI

From the GUI, she:

1. Selects **Drivers** → **Commit** from the Actions pull-down menu on the Tasks window.
2. Types 0105 in the **Drivers** field and robot_control in the **Release** field.
3. Selects **OK**.

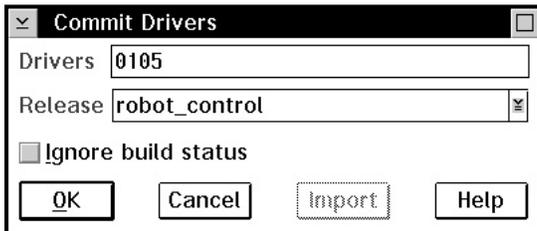


Figure 39. Commit Drivers window

Command

From a command line, she issues the following command:

```
teamc driver -commit 0105 -release robot_control
```

Results

TeamConnection moves the part versions associated with driver 0105 into the release. Other members of the team can now view the changes. Committing a driver commits all work areas designated as driver members and all parts changed in reference to those work areas.

Completing the driver

The driver is ready for formal testing in the specified release's environment list. Testing is tracked using test records for each environment in which testing is to be done. To create the test records, Carol must complete the driver.

GUI

From the GUI, she:

1. Selects **Drivers** → **Complete** from the Actions pull-down menu on the Tasks window.
2. Types 0105 in the **Drivers** field, and selects robot_control from the **Release** field.
3. Selects **OK**.

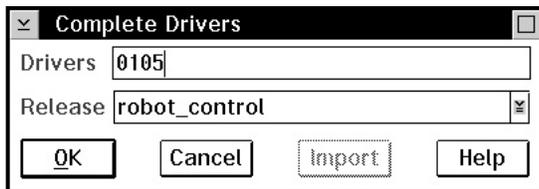


Figure 40. Complete Drivers window

Command

From a command line, she issues the following command:

```
teamc driver -complete 0105 -release robot_control
```

Results

All the work areas in the driver are changed to the test state, and test records are created.

Testing the built application

Annmarie is the tester for the MVS version of the robot application. When she receives notification that the test record is in the ready state, she tests the part changes that were made within the release by Alex and several of his team members. The tests complete successfully, so she accepts the test record by doing one of the following:

GUI

From the GUI, she:

1. Selects **Records** → **Test records** → **Accept** from the Actions pull-down menu on the Tasks window.
2. Types 456 in the **Work areas** field, and selects robot_control from the **Releases** field and MVS from the **Environments** field.
3. Selects **OK**.

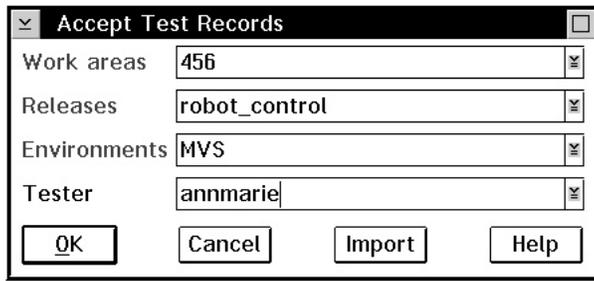


Figure 41. Accept Test Records window

Command

From a command line, she issues the following command:

```
teamc test -accept -workarea 456 -release robot_control -env mvs
```

Results

Annmarie's test record moves to the accept state. However, work area 456 will not go to the complete state until Tim, who is the tester for the OS/2 environment, marks his test record.

After all test records are moved from the ready state, the work area moves to the complete state. Because the component process includes the verifyDefect subprocess, defect 456 moves to the verify state. A verification record for the defect is created in the ready state.

Using a configured process

The scenarios in this chapter and the preceding chapter illustrate one release with no process management enabled and another release with full process management enabled. However, administrators can define a release that requires users to work with some intermediate level of process management. That is, the administrator can remove some of the subprocesses from the full-tracking scenario.

For example, the administrator might want to eliminate the driver subprocess. If the driver subprocess is eliminated, the user cannot create driver members to associate the changes in a work area with a driver. Likewise, users cannot commit drivers to integrate several work areas with the release. Instead, users integrate the changes for each work area by integrating the work area with the release.

To demonstrate how this works, assume that Carol and Alex are trying to fix the robot's dislike of striped wallpaper using a release without the driver subprocess enabled. Initially, the scenario is not affected by the absence of the driver subprocess. The defect is opened, and a work area is created. Alex, after receiving notice that he needs to solve the problem, goes through the process of checking out the faulty part, making fixes, checking the fixes into the work area, and rebuilding. He can still freeze the work area whenever he wants to save its current content.

The difference occurs when Alex is ready to integrate his changes with the release. When the driver subprocess is not enabled, Alex issues the following command:

```
teamc workarea -integrate 456 -release robot_control
```

This command moves the part versions associated with work area 456 into the release so they are visible to other developers. However, if collision records are created, TeamConnection flags the concurrent changes and stops the integration until the changes are reconciled and the corresponding collision records are completed.

Retrieving a past version of a part

Versioning is an insurance policy for the developer. By freezing the work area, the developer knows that the parts currently visible in the work area will be retained in their current form.

For this example, assume that Alex just updated the `optics.c` module to add support for a new zoom lens. Alex did a considerable amount of work on this task, and it required a dozen check-out, check-in, and build cycles before he finished. Alex's work area now contains the following:

<code>brain.c</code>	<code>leg.c</code>
<code>brain.obj</code>	<code>leg.obj</code>
<code>brain.exe (from Alex's build 12)</code>	<code>foot.c</code>
<code>arm.c</code>	<code>foot obj</code>
<code>arm.obj</code>	<code>optics.c (modification 12)</code>
<code>hand.c</code>	<code>optics.obj (from Alex's build 12)</code>
<code>hand.obj</code>	

Next Alex must update the `brain.c` part to set the appropriate conditions for activating the new zoom capability. He does not yet want to integrate his changes to `optics.c` for the zoom lens with the release because they are of little value without his changes to `brain.c`. Also, he is not certain that he is completely done with `optics.c` until he completes the modifications to `brain.c`. Rather than integrate an incomplete change, he freezes his work area by issuing the following command:

```
teamc workarea -freeze 1208 -release robot_control
```

This command takes a snapshot of the work area and its parts in their current state.

As Alex works on the brain.c module, he makes sweeping modifications to optics.c to simplify the interface between brain.c and optics.c. Unfortunately, he realizes too late that the simplification he is pursuing will not work. Rather than spend several hours removing his updates to optics.c, he wants to start fresh from a copy of optics.c that does not contain the changes for the simplification.

Alex has frozen his work area three times since beginning work on the zoom lens integration. Also, he has done additional check-ins to his work area since his last freeze. He cannot remember the particular version of his work area that contains the copy of optics.c that he wants. So, he wants to see all the versions of his work area that he has saved. He issues the following report command:

```
teamc report -view versionView -where "workAreaName='1208' and  
releaseName='robot_control'" -stanza
```

This command returns a list of the versions frozen from work area 1208. The report looks like this:

```
name                1208:1  
workAreaName        1208  
releaseName          robot_control  
predecessor          robot_control:5  
hasSuccessor         yes  
releaseVersion       no  
addDate              1995/01/11 14:30:26  
freezeDate           1995/01/11 15:00:00  
  
name                1208:2  
workAreaName        1208  
releaseName          robot_control  
predecessor          1208:1  
hasSuccessor         yes  
releaseVersion       no  
addDate              1995/01/12 09:25:13  
freezeDate           1995/01/12 17:15:58  
  
name                1208:3  
workAreaName        1208  
releaseName          robot_control  
predecessor          1208:2  
hasSuccessor         yes  
releaseVersion       no  
addDate              1995/01/14 11:13:25  
freezeDate           1995/01/15 09:01:35  
  
name                1208:4  
workAreaName        1208  
releaseName          robot_control  
predecessor          1208:3
```

hasSuccessor	no
releaseVersion	no
addDate	1995/01/16 08:10:15
freezeDate	1995/01/16 10:05:11

So what does it all mean?

- name is the name of the version in the work area.
- workAreaName is the name of the work area that owns the version.
- ReleaseName is the name of the release that owns the version.
- Predecessor is the name of the version that precedes, or is the parent of, this version.
- hasSuccessor has a value of *yes* if the version has a successor, *no* if it does not.
- releaseVersion has a value of *yes* if the version is part of the release's main version history; the value is *no* if the version belongs to a work area.
- addDate is the date and time the version was created.
- freezeDate is the date the version was frozen.

This report seems erroneous. TeamConnection returned four versions in the report even though Alex has executed the freeze command against his work area only three times. The fourth version, 1208:4, is the unfrozen version in which Alex is currently making his changes.

Another concern might be the predecessor of the first version returned in the report. Why is its predecessor robot_control:5? At some point Alex began his work by making modifications to the latest code in the release. The first version of Alex's changes is based on the release version robot_control:5.

After reviewing the report, Alex thinks that his last working copy of optics.c was saved when he created version 1208:2. However, to make sure, he wants to see the parts modified in version 1208:2. He issues the following command:

```
teamc report -view partView -version 1208:2 -release robot_control
           -where "currentVersion='1208:2'" -stanza
```

This report returns a list of parts visible to version 1208:2 that have a currentVersion (or version ID) of 1208:2. If a part has such a version ID, the part was modified in the version 1208:2.

Note: If the -where clause were not specified, the report would return all of the parts visible from version 1208:2.

The TeamConnection system returns the following report:

baseName	optics.c
releaseName	robot_control
compName	robot_dev
versionSID	1208:2
addDate	02/02/94
lastUpdate	04/15/94

```
pathName          smarts\eyes\optics.c
nuVersionSID      1208:2
nuAddDate
nuDropDate
nuPathName
userLogin         alexm
fmode             0640
```

Because optics.c is the only part modified in version 1208:2, Alex assumes it is the copy he wants. He extracts the part by issuing the following command:

```
teamc part -extract optics.c -version 1208:2 -workarea 1208 -release robot_control
```

This command extracts the desired copy of optics.c from the frozen version 1208:2. Alex can then overlay the corrupted copy of optics.c that he has checked out with the copy he just extracted, and he can start over fresh. He can also check in the overlaid optics.c to his work area.

This method works only for parts with a file type of TcPart. If your part has a type of something other than TcPart, you can do one of the following to restore the part:

- Use the undo action if restoring to the previous version.
- Use the link action to link to a previous version.

In addition to the reporting features mentioned above, Alex can also obtain a list of work areas by issuing the following command:

```
teamc report -view WorkAreaView -where "releaseName='robot_control'" -stanza
```

The report that is returned lists the work areas in the release robot_control. A user can also see the parts changed for each work area by specifying the -long parameter on this command.

Part 3. Using TeamConnection Notes Integrated Databases

Chapter 7. Introduction to TeamConnection Integrated Notes Databases	103
Getting started	103
Prerequisites and dependencies	104
Using TeamConnection with Lotus Notes	104
Sources of user information	104
Database types	105
Forms and subforms	106
Views	107
Reviews	109
Document archiving	109
Chapter 8. Creating and Maintaining Integrated Notes Databases	111
Initializing the original template and creating a database.	112
Creating customized production databases	116
Performing reconciliation	118
Database maintenance: refreshing design from a template	119
Chapter 9. Database Design Strategies and Advanced Customization	121
Rules of thumb and general advice.	121
Using the Customization setup facility	122
Notes Database Customization	123
Modify TeamConnection Access	124
Reconciliation of Notes and TeamConnection Data	124
Advanced customization	125

This section is for anyone who will be responsible for accessing the Integrated Notes Database code shipped with TeamConnection and customizing predefined database templates to best serve your project's needs.

Chapter 7. Introduction to TeamConnection Integrated Notes Databases

The VisualAge TeamConnection Integrated Notes Database feature provides a documentation facility to support software development. A software development group can use this database to communicate with TeamConnection objects from within Lotus Notes documents. The TeamConnection Integrated Notes Database links the technical documents to the TeamConnection objects involved in software enhancement and maintenance.

A single database template is provided that you use to define the databases that will assist you with all phases of your development process. From the template, databases can be produced that can assist with requirements, design and development documentation, and test case management, as well as other purposes.

Getting started

This portion of the User's Guide is intended to direct administrators and users on the fundamentals of preparing for, configuring, and using the IBM VisualAge TeamConnection Enterprise Server Integrated Notes Database feature.

The feature requires some degree of administrative intervention, including set up, customization, and tuning, before the databases provided by the feature are available to development or test team members in your organization. These issues are addressed in "Chapter 8. Creating and Maintaining Integrated Notes Databases" on page 111 and "Chapter 9. Database Design Strategies and Advanced Customization" on page 121. Chapter 8. describes post-installation tasks at a high level. If you have no intention of performing *advanced* customization, Chapter 8. provides enough information to activate and begin testing an integrated database.

It is recommended that you read "Chapter 9. Database Design Strategies and Advanced Customization" on page 121 if you plan on customizing subforms, or if you have no experience with Lotus Notes database design.

For an overview of database templates and the user interface, begin with "Using TeamConnection with Lotus Notes" on page 104. The database itself contains the most detailed user documentation, and is provided in the following forms:

- Database *About* document
- Database *Using* document
- Help (contextual)

Prerequisites and dependencies

In preparation for setting up and administering the Lotus Notes Integration feature, ensure that the machine you are using to implement this feature (generically, the *Notes server*) is equipped as follows:

- To use this function, the client machine requires Lotus Notes Release 4.5.3a or higher and a Version 3 TeamConnection client. There are no prerequisites or dependencies on the TeamConnection server.
- The TeamConnection Web Client interface is also required.

Using TeamConnection with Lotus Notes

The VisualAge TeamConnection Integrated Notes database template provides a generalized technical documentation facility to support software development. This section is intended as an overview for Notes Integrated Database users.

The TeamConnection Integrated Notes database links technical documents to the TeamConnection parts involved in software enhancement and maintenance. Your software development team can use this database to communicate with TeamConnection from within Lotus Notes documents. A reconciliation facility is provided that synchronizes the data in the TeamConnection family with the Notes databases that use it.

Your team can create software design documentation, project requirements documents, and test cases using Lotus Notes. Then utilize the integrated tracking and change control functions TeamConnection provides to identify, organize, manage, and control software parts as they change over time.

You can associate Notes documents with existing defects and features within TeamConnection and use the default web browser, defined in Notes, to view additional details of the TeamConnection defects and features or perform other TeamConnection actions.

You can also access the TeamConnection client GUI from Notes.

Sources of user information

For more information about TeamConnection Integrated Notes databases, see the *About*, *Using*, and contextual help information available from the database **Help** menu. *About* and *Using* are available from the **Navigator** window **Help** menu. It is advisable to examine these documents before entering your integrated database.

Note: The *About* document is also available from the **Help** choice on the **Navigator** window.

The *About* document provides graphical previews for the Requirements, Development, and Test database types. These previews provide a graphical depiction of available documents and the document hierarchy for each database, along with the default setup values available to an administrator.

The *Using* document provides detailed information that users will need to take advantage of Notes Integrated Database functionality. For each database type, *Using* documents the specific steps required to create and respond to the available document types. In addition, *Using* provides a listing of each Notes element (form or view), along with a brief description of the context in which it is appropriate.

As you navigate through the integrated database, you can access help in the several ways:

For a view

Select the **Help** button in the action bar to look at help for the view.

For the document that you are editing

Select the **Help** button in the action bar to look at help for a document or form.

For a field in a form

Click in the field and look at the bottom of the screen for a short description of the field. Look to the right of the field to see if there is descriptive text for help in entering the correct format. Select the **Help** push button in the action bar to get more detailed help on the field.

Database types

A single database template is provided that you use to define the databases that will assist you with all phases of your development process. Using the different databases you can create various documents, open TeamConnection features and defects, associate, modify, or cancel existing TeamConnection features and defects, and use the Web browser to view information on TeamConnection features and defects.

The Notes database is linked to the TeamConnection database to integrate information in Notes with information in TeamConnection. In this way, TeamConnection defects and features are linked to Notes documents. The following database types can assist you with the overall management your software project:

Requirements

The Requirements database assists with documenting and categorizing Notes documents that describe requirements. TeamConnection features can be opened for software design requirements as they move through the formal phases of development. You can also associate Notes documents with existing features in TeamConnection.

Software Design and Development

The Software Design and Development database assists with software design and development, including change control to develop a complete and comprehensive set of software design documentation. TeamConnection features and defects can be opened to identify and track the implementation of

the design. Documents existing in Notes can be associated with existing features and defects in TeamConnection.

Test Case Management and Tracking

The Test Case Management and Tracking database assists with test case management and tracking. In this database, test case definitions and their execution results are tracked. Defects can be opened for failing test cases to track their progress.

Generic

A Generic database is available for other purposes that you choose and also has access to TeamConnection.

User Defined

A User Defined database is also available. Advanced users can define a database that utilizes the Notes documentation facility, combined with the TeamConnection integration built into the template, with their own local definitions of forms and views.

Each of the databases also allow you to capture and track information about documents and link the documents to TeamConnection. With each of type of database, you can add a document to the database, respond to a document, respond to a response, flag a document as private, and route a document for review to your team members.

Refer to the database *Using* document for a description of the forms, views, and tasks associated with each of the database types available.

Forms and subforms

Forms provide the structure and organization of elements in the documents used by your database. The integrated database forms come preloaded with default values. There are one or more subforms defined for all the forms that allow user-defined subforms. This allows you to augment or replace the subform supplied by TeamConnection. The TeamConnection-supplied database teamc.nsf has all of the user subforms defined. See "Advanced customization" on page 125 for more details on using user-defined subforms.

Table 1 provides a listing of all forms currently available to each database type.

Table 1. Integrated database forms by database type

Form	Use for this form
Common	
Basic Document	Document basic information that would not be included in a requirements, design, or test case document.
Response	Document your response to a requirements, design, test case, or basic document.
Response to Response	Document your response to another's response to a requirements, design, test case, or basic document.

Table 1. Integrated database forms by database type (continued)

Form	Use for this form
Requirements	
Requirement	Document and open a requirements document in the database.
TeamConnection Feature	Open a TeamConnection feature against a requirement document to track the implementation of the requirement.
Development	
Design	To open a design document in the database.
Enhancement description	To document a small design issue or improvement (not associated with a design document).
Lower level design	To document information associated to a design document.
Design change description	To document design change information.
TeamConnection Feature	Open a TeamConnection feature against a design to track the implementation of the design.
TeamConnection Defect	Open a TeamConnection defect against part of a design to track a design defect.
Test	
Test case	To document a test case that should be run against the design or specific function of a software product, or TeamConnection defect.
Execution record	To document information about test case execution.
TeamConnection Defect	Open a TeamConnection defect to track a test case failure.
Generic	
Basic Document	Document basic information for future reference.
Response	Document your response to a basic document.
Response to Response	Document your response to a response of a basic document.
User Defined	To define the documents within a User Defined database, you must modify the forms and subforms that are provided. Modifying forms and subforms requires knowledge of Notes Forms design and Designer authority on the Notes database.

Views

Views provide various ways to organize and access the documents in your integrated database. By selecting a view, you can access and sort specific categories of documents in the database. You can also navigate through the database to search for a specific document title. By using buttons on the action bar, you can choose to view the children documents of a parent document or category, all active features, or all active defects.

Views are available through the **Navigator** and the **View** menu.

Table 2 on page 108 provides a listing of all views currently available to database users.

Table 2. Integrated database views

View	Documents accessed (purpose)
All documents	To list all of the documents stored in the database.
My favorite documents	Contains documents you have placed in the folder.
Administration	
Document Control	To list all of the documents by their document control number.
Documents	
All	To list all of the user documents stored in the database.
Archiving	Allows you to archive the database.
By Author	To list all of the user documents stored in the database organized by author.
Hierarchy	To list the user documents in the database's hierarchy.
Review Status	To list the review status of user documents in the database.
State Summary	To list the states of user documents in the database.
Development	
Design Changes	To list all of the design change documents in the database.
Enhancements	To list all of the enhancement documents in the database.
Hierarchy	To list all of the design documents in a hierarchy.
My Designs	To list all of the current user's documents.
State Summary	To list the state of design documents in the database.
Requirements	
Final Priority	To list the final and originator priority of documents in the database.
Hierarchy	To list the requirements and associated features in a hierarchy.
List	To list all of the requirements documents in the database.
My Requirements	To list all of the current user's requirements documents.
Originator Priority	To list the originator priority of documents in the database.
State Summary	To list the state of requirements in the database.
Test	
Execution Records	To list all of the execution records in the database.
Hierarchy	To list test case, execution records, and defects in a hierarchy.
My Execution Records	To list all of the execution records that you have created in the database.
Testers	To list the tester name(s) and a summary of their test case executions.
Test Case	To list all of the test cases in the database.
TeamConnection	
Defects	To list all of the TeamConnection defects associated with the database.
Features	To list all of the TeamConnection features associated with the database.
Reconcile Log Records	To list a summary of the reconcile log results for the database.
Agents	To list the agents configured by your administrator.

Reviews

You can make your design, development, requirements, and test case documents available to your team for review through the Notes review function. The author of a document can set up a document review cycle for a document. Documents can be routed to the reviewers chosen by the author one at a time, in sequence, or all at once. A reviewer log is automatically generated to show the state of the documents and the names of the reviewers.

This application was designed with the intention that all users except the manager should have Author access. If users have Editor access, the review cycle may not function correctly.

See the database *Using* document for a description of the tasks associated with the document review process.

Document archiving

To keep the document library current with only the latest topics, document authors can choose to move documents from the current database and store them in a different database. In addition, Notes databases can be checked into TeamConnection.

Most of the Archiving activities take place from the **Archiving** view. You must switch to this view in order to initiate archiving on a document library database.

See the database *Using* document for a description of the tasks associated with the document archiving process.

Chapter 8. Creating and Maintaining Integrated Notes Databases

This section and the one that follows (“Chapter 9. Database Design Strategies and Advanced Customization” on page 121) provide instructions for the members of your organization, referred to as the *Lotus Notes Database Administrator*, or *administrator* for brevity’s sake, who will be responsible for accessing the Integrated Notes Database code shipped with TeamConnection, customizing predefined database templates to best serve your project’s needs, and maintaining your organization’s production integrated databases.

Important notice to administrators

For any administrative activities that involve teamc.nsf and fhcnote.ntf, you can locate these files on the TeamConnection server in the NLS\CFG\LANG path (\$TC_HOME/NLS/CFG/[LANG] for UNIX platforms), where [LANG] specifies the language directory for your version of the product (ENU or en_US, for example).

As the designated administrator works through the tasks described in this section, it may become apparent that other members of the organization should assume the role of individual database owner, and make decisions that affect specific databases. This segmentation or hierarchy should become more apparent as the tasks described move from installation and setup issues to those that involve design and maintenance.

After you have installed TeamConnection and verified the items listed in “Prerequisites and dependencies” on page 104, you can proceed to performing administrative tasks necessary to activate this feature from a Notes standpoint. This is likely to be staged process. After loading initial database templates, as described in “Initializing the original template and creating a database” on page 112, you can begin the recursive process of customization and tuning. Setup for the initialization stage is required. Design strategies and advanced customization (which is optional) are described in detail in “Chapter 9. Database Design Strategies and Advanced Customization” on page 121.

The primary administrative tasks discussed in this section are summarized in Figure 42 on page 112:

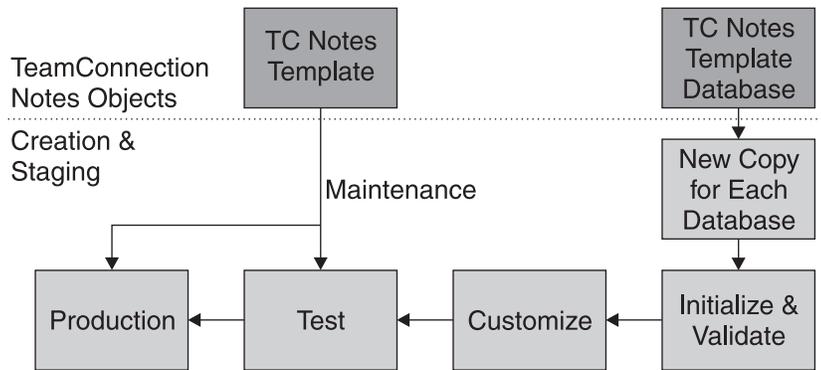


Figure 42. Notes Integrated Database Creation and Staging

The flow presented in the figure applies to all potential user databases available through the Notes Integrated Database feature.

Initializing the original template and creating a database

To help you with developing Notes applications, a database master template is included on the IBM VisualAge TeamConnection Enterprise Server installation CD-ROM. The file, which serves as a single source for a number of database uses, is placed on the TeamConnection server during installation and will stay there until you move it to a Notes directory.

Note: Regular Notes databases have the .NSF extension. Notes templates have an .NTF extension. You will use teamc.nsf database to create new integrated notes databases. Use the database *template* fhcnotes.ntf to refresh the design of existing databases. Although teamc.nsf and fhcnotes.ntf can be accessed, you should never modify these databases directly.

To activate this function in Lotus Notes, perform the following steps:

1. Obtain a copy of teamc.nsf file from the server and set it up to meet your local requirements in the following manner:
 - a. Copy teamc.nsf from the server to your local workstation into your local notes data directory (something like *x:\notes\data*). Use the method of file transfer of your choice, but be aware that teamc.nsf is sizable (more than 5 Mb).
 - b. You must create a copy of the teamc.nsf file that will act as your target database. Create a copy of the empty teamc.nsf database on your Notes workspace as follows:
 - 1) Select the TeamConnection Database (teamc.nsf).
 - 2) Select **File->Database->New Copy**. This displays the **Copy Database** window.

- 3) The **Server** field should contain the Local value. You can supply the **Title** of your choice.
 - 4) Rename the file in the **File Name** field to match the naming conventions of your organization. This copy of the template must keep the .nsf suffix. For example: development.nsf.
 - 5) Select the folder icon button, which will open the **Choose a folder** window. Choose the directory on your local file system that will store the database.
 - 6) In the **Copy** field select the **Database design only** radio button, and leave the **Access Control List** checkbox checked, so that the new database will inherit those characteristics from the empty template database.
 - 7) Select **OK**.
- c. Add yourself as an Administrator to the **Access Control List** for the new database as follows:
- 1) Using mouse button 2, select the notes database icon on your Notes workspace page.
 - 2) Select **Access Control**.
 - 3) From the **Access Control List** window select **Add**.
 - 4) Enter your full Notes name, as it would appear in your Notes Address Book.
 - 5) Assign yourself the role of **Administrator** and **Author** by selecting them in the **Roles** box of the panel.
 - 6) Assign a value of Person in the **User type** field and Manager in the **Access** field.
 - 7) Select **OK**.
- d. Open the new empty database. The **About** overview window will display showing what this Notes database is. When you read and exit (Escape) this window, the main window for this database should display.
- e. There are 3 choices on this window: **Setup**, **Proceed to Database**, and **Help**. We suggest that you select **Help** to view the *About* document before proceeding to **Setup**. *About* provides details about database and document types that may be beneficial during the setup process.
- f. Select **Setup**.
- Note:** If you are not defined with the role of Administrator in the Access Control List, you will only be able to view the setup in read—only mode.
- g. From the Setup window follow the instructions on completing the **Initialization** setup.

Note

Setup is a staging process that may require some planning. See “Chapter 9. Database Design Strategies and Advanced Customization” on page 121 for details on setup strategies and customization options.

You will need to decide what type of database you need: Requirements, Design and Development, Test Case Management, Generic, or a User Defined database.

Note: The **Preview** button enables you to view database forms and related information to help determine which type of database is most suitable for your group.

See “Using TeamConnection with Lotus Notes” on page 104 or the Notes database documents for a description of each database type.

You will need to know your TeamConnection family and server port name to complete the initialization portion of setup. You must also have authorization to connect to and use the server. You should use a test family, initially, until you have the details of your database design and customization stabilized and tested before pointing the setup to a “live” TeamConnection family.

Note: You must identify your full TeamConnection family specification (family@hostname@port) as an alias in your hosts and services files, and supply only the family name in the ***Family** field. In addition, all database users must have the same entry in their hosts and services files.

Be sure to use the **Test Connection** and **Test Web Client** buttons to verify that all necessary installation configurations are valid.

Select the **Finish** button to update the integration with TeamConnection and complete the **Initialization**.

- h. When the **Initialization** is complete, you can either begin the **Customization** setup (see Creating customized production databases) or return to the main window. If you want to begin working with the initialized database, select **Proceed to Database**. Your new, but empty database is available. You can now test that it is working to your satisfaction.

In the most basic terms, testing would include the following activities:

- Create documents.
- Update documents. This should include testing the state system, which would involve moving a document’s state to Approved.
- Reopen documents.
- Select **Feature** and **Defect** push buttons.
- Open Defects and/or Features.
- Examine views.

You should also add any users involved in the testing process to the access list for this database. The default **Access** of Author is recommended.

- i. When the database is working to your satisfaction (you can go back to the setup several times to make changes and refine it), you are ready to make it available to your user community.

Be sure to delete any documents you created during local testing and evaluation.

Note: If document numbering is turned on, you may want to reset it when the database goes into production. To reset document numbering, perform the following steps:

- 1) Select the Notes **Navigator**.
 - 2) Select **Administration->Document Control**.
 - 3) Edit the document and set the number to 0 (zero).
 - 4) Select **File->Save**.
2. If, after testing the current setup, you find that the default settings for the database are adequate for the needs of your organization, proceed to the next step. Otherwise, proceed to "Creating customized production databases" on page 116 for further customization and testing of the database.
 3. Put the database that you just setup on a Notes server using the normal Notes process.

Note: To enhance performance for database users, it might be advisable to turn replication on, using the standard Notes process.

4. Identify the users of the application and the access level each will need. List the name of each user and user group in the **Access Control List** accordingly. Database users must have Author access. The template has been set up so that any database created with it will have the default access of Author.

Note: If users will create documents *and* documents are automatically numbered (as defined in setup), then a **Role** of Author is also *necessary*.

You can also choose to assign the role of Project Leader. Certain document states can be reserved for only the Project Leader to set (such as, "Approved"). This role allows a project leader permission to place a document into one of these reserved states. The Administrator assigns these states during Setup. One or more individuals can be defined as a Project Leader or an Administrator.

If you want to limit this database to only members of a certain user group, the default user should have No Access, and members of the user group will need to be listed in the **Access Control List** as either individuals or groups with Author access. At least one person needs to have the role of Administrator.

5. Arrange for adding the database icon to the Notes workspace on the appropriate client machines, using the normal Notes process used by your organization.

Creating customized production databases

Now that you have initialized the Integrated Notes Database feature and performed a preliminary setup of the database template for your organization, you can proceed to customize the initialized database.

See “Using the Customization setup facility” on page 122 for additional details on the customization options.

Note: A **Customization Wizard** is available from the **Setup** window action bar to assist you in customizing your database.

Perform the following steps to customize your database and make it available at a production level:

1. From the main database window, select **Setup**.
2. Use the **Customization** setup facility to customize the database design. Using **Customization**, you can modify the default values that are provided.

Note: If your customization will extend to subforms, you will need some familiarity with the Lotus Notes environment.

You can adjust or refresh values are for the following:

- Notes document names, titles, and subtitles.
- Structure of the Notes document hierarchy.
- States each Notes document may progress through.
- TeamConnection family, component, and configurable field information.
- TeamConnection feature and defect attributes to store in Notes documents.
- Reconciliation of Notes and TeamConnection data. See “Performing reconciliation” on page 118 for details.

Make any design changes that are necessary for your particular application and organization’s workflow. This process is detailed in “Chapter 9. Database Design Strategies and Advanced Customization” on page 121 and within the Customization setup facility itself.

3. When the setup is complete, you will return to the main window. Select **Proceed to Database**. Your new, but empty database is available. You can now test that it is working to your satisfaction.

In the most basic terms, testing would include the following activities:

- Create documents.
- Update documents. This should include testing the state system, which would involve moving a document’s state to Approved.
- Reopen documents.
- Select **Feature** and **Defect** push buttons.

- Open Defects and/or Features.
- Examine views.

You should also add any users involved in the testing process to the access list for this database. The default **Access** and **Role** of Author is recommended.

4. When the database is working to your satisfaction (you can go back to the setup several times to make changes and refine it), you are ready to make it available for to your user community.

Be sure to delete any documents you created during local testing and evaluation.

Note: If document numbering is turned on, you may want to reset it when the database goes into production. To reset document numbering, perform the following steps:

- a. Select the Notes **Navigator**.
 - b. Select **Administration->Document Control**.
 - c. Edit the document and set the number to 0 (zero).
 - d. Select **File->Save**.
5. Put the database that you just setup on a Notes server using the normal Notes process.

Note: To enhance performance for database users, it might be advisable to turn replication on, using the standard Notes process.

6. Identify the users of the application and the access level each will need. List the name of each user and user group in the **Access Control List** accordingly. Database users must have Author access. The template has been set up so that any database created with it will have the default access of Author.

Note: If users will create documents *and* documents are automatically numbered (as defined in setup), then a **Role** of Author is also necessary.

You can also choose to assign the role of Project Leader. Certain document states can be reserved for only the Project Leader to set (such as, "Approved"). This role allows a project leader permission to place a document into one of these reserved states. The Administrator assigns these states during Setup.

If you want to limit this database to only members of a certain user group, the default user should have No Access, and members of the user group will need to be listed in the **Access Control List** as either individuals or groups with Author access. At least one person needs to have the role of Administrator.

7. Arrange for adding the database icon to the Notes workspace on the appropriate client machines, using the normal Notes process used by your organization.

Performing reconciliation

The Reconciliation facility synchronizes the data in the TeamConnection family and the Notes databases that use it. The reconciliation facility is an agent that should be run regularly. As a default, reconciliation is activated during Setup. It can be run on an established schedule that you set up, or manually when needed. We suggest that the reconciliation facility be set up to run after your regular TeamConnection build completes.

To run reconciliation on a regular schedule, the scheduled reconcile agent needs to run on either your Notes server or a Notes client using a replica of the server database. In either case, the TeamConnection client must also be available to the server or the client where reconciliation is scheduled to run. Typically, an administrator will set up this client to run automatically on his or her own workstation.

You must be assigned the role of Administrator to activate the reconciliation facility. In addition, for the Notes client that will run the agent, the local user preferences (available from **File->Tools->User Preferences**) must have **Enable local scheduled agents** checked.

To **Activate Reconcile**, verify that **Yes**, the default, is selected in the **Customization** setup. To activate the reconciliation facility on a Notes client, do the following:

1. Using standard Notes replication procedures, create a replica of the server database on the desired Notes client.
2. Verify that the Notes user is enabled with a **Role** of Administrator.
3. Make sure replication is active between the local replica and the server.
4. Select the local replica database icon using mouse button 2 and select **Go to Agents** from the pop-up menu.
5. Open the **Reconcile Notes and TeamConnection** option that has a trigger value of **Scheduled**.
6. The **Options** button will provide details about the named reconciliation. Select the **Schedule** button.
7. Supply values for the fields on the **Schedule** window that meets the needs of your team, and then select **OK**.
8. Select **File->Save**.
9. Check the agent's box. Provide a server name in the **Choose Server To Run On** window. Use Local if you are running the agent on your Notes client.
10. Select the **OK** push button.

The result is that reconciliation will be performed between your local replica Notes database and TeamConnection. Any changes to Feature or Defect documents will be replicated to your server Notes database.

To enable and run the reconciliation facility manually, select **Actions->Reconcile Notes and TeamConnection** from any database window.

Database maintenance: refreshing design from a template

In the course of using Notes Integrated databases, you may have occasion to refresh current databases with a template. As you receive code updates from IBM VisualAge TeamConnection Enterprise Server, it is likely that you will want your current databases to reflect the most current template. In the course of maintaining database consistency across your enterprise, you may want to refresh all active databases from a common template. In either event, you should use caution. The steps that follow will address this issue.

Note: After TeamConnection maintenance has been applied, obtain an updated version of the fhcnote.nsf template from the server and put it in your local notes data directory (for example, x:\notes\data).

To refresh a database from a template, do the following:

1. Make a new copy of the database that you want to refresh, using the following procedure:
 - a. Select the TeamConnection Database you want to refresh.
 - b. Select **File->Database->New Copy**. This displays the **Copy Database** window.
 - c. The **Server** field should contain the Local value. You can supply the **Title** of your choice.
 - d. Rename the file in the **File Name** field to match the naming conventions of your organization. This copy of the template must keep the .nsf suffix. For example: development.nsf.
 - e. Select the folder icon button, which will open the **Choose a folder** window. Choose the directory on your local file system that will store the database.
 - f. In the **Copy** field select the **Database design and documents** radio button, and leave the **Access Control List** checkbox checked, so that the new database will inherit those characteristics from the empty template database.
 - g. Select **OK**.
2. Select the local database copy that you have just created from your workspace.
3. From the **File** menu, select **Database->Replace Design**.
4. You will locate the template used to refresh the database, fhcnote.nsf, in the **Replace Database Design** window. If the template is not on your Local server, which primes the list of template choices, you may have select the **Template Server** button to locate the desired template.

Important notice to administrators:

Be sure that you select the right template. An inappropriate template may have a destructive impact on your database! Select **Replace**.

5. Test the refreshed database copy extensively, including the following activities (at a minimum):
 - Create documents.
 - Update documents. This should include testing the state system, which would involve moving a document's state to Approved.

- Reopen documents.
 - Select **Feature** and **Defect** push buttons.
 - Open Defects and/or Features.
 - Examine views.
6. When you are satisfied that refreshing the local database copy was a successful operation, select the production (original) database from your workspace and repeat the refresh process, beginning with step 3.

Important notice to administrators:

Be sure that you select the right template. An inappropriate template may have a destructive impact on your database!

7. You should test the refreshed database extensively before alerting users that it is in production mode.

Chapter 9. Database Design Strategies and Advanced Customization

This section includes some general strategies for setting up, testing, and tuning your Notes Integrated Database implementation. The primary administrative functions are described in "Chapter 8. Creating and Maintaining Integrated Notes Databases" on page 111, but if you have limited experience with Lotus Notes database administration, wish to familiarize yourself with terminology used by templates, or expect to significantly modify the template database used to create your organization's integrated databases, it may be valuable to read this section of the document.

Rules of thumb and general advice

There are some simple, general rules that will make your implementation of Notes Integrated Databases:

- Always work on a copy of a database, rather than a server (common) version when altering design. This guarantees that you will not overwrite your customizations when performing maintenance.
- Stage database setup and implementation. You will always begin with an **Initialization** setup and test phase. After the high-level verification of a Notes Integrated Database template's "fitness" for your organization, it is likely to require customization

It is also likely that a staging scheme may involve changes in database ownership, as increasing customization is necessary. Individual database administrators may have more to say about the details of implementation, and therefore control of databases has a tendency to become decentralized.

- In terms of customizability, document hierarchies are relatively static, as compared to forms and views.

Note: The primary exception to the customizability of forms is the set of subforms supplied by TeamConnection *as they are currently named*. It is important to copy and rename any forms that you plan to customize, so the template refreshes do not overwrite your customizations.

- Testing at every stage of refinement is crucial to successful implementation, related to stability and improved mapping to your organizations needs.
- To enhance performance for database users, it might be advisable to turn replication on, using the standard Notes process.
- Plan and monitor access control as needed.

To summarize the basic steps for moving from the originally installed TeamConnection template database through opening a specific database to your user base for production-level work:

1. Download the teamc.nsf template to your local machine and rename it to something that makes sense for your organization.
2. Complete **Initialization** setup.

3. Assess and implement any access control/authorization schemes necessary at this level of implementation.
4. Open the newly initialized database.
5. Create test documents for testing.
6. Open features and defects if your database addresses these elements.
7. Determine whether existing (TeamConnection—supplied) documents, subdocuments, and forms can be mapped into your organization's terminology, workflow, hierarchies, and state system.
8. Perform **Customization** setup, iteratively, to enact structural changes required by your implementation. Customization enables you to address structural categories, such as menus, document names and hierarchies, and workflow states.
See "Using the Customization setup facility" for a more detailed treatment of this topic.
9. Assess and implement any access control/authorization schemes necessary at this level of implementation.
10. Open the newly customized database.
11. Create test documents for testing.
12. Open features and defects if your database addresses these elements.
13. Validate as much as possible in your testing to determine the database's production readiness.
14. Delete all test documents, etc., and place the database in production status on a Notes server. If you have used a test TeamConnection family to this point, it will be necessary to specify the appropriate production family in the **Customization** setup.
15. Make the "live" database available to the appropriate user base.

Using the Customization setup facility

Use the **Customization** setup facility to customize the database design. Using **Customization**, you can modify the default values that are provided. This requires familiarity with the Lotus Notes environment.

Note: A **Customization Wizard** is available from the **Setup** window action bar to assist you in customizing your database.

You can adjust or refresh values are for the following:

- Notes document names, titles, and subtitles.
- Structure of the Notes document hierarchy.
- States each Notes document may progress through.
- TeamConnection components and configurable fields information (refresh only).
- TeamConnection feature and defect attributes to store in Notes documents.
- Reconciliation of Notes and TeamConnection data.

The **Customization** setup facility addresses the following areas of database manipulation:

- Notes Database Customization
- Modify TeamConnection Access
- Reconciliation of Notes and TeamConnection Data

In order to perform the Customization setup, you must set your role to Administrator in the **Access Control List** window.

Notes Database Customization

The following areas of customization are described in detail in the Setup help.

Modify Database Optional Information

In the **Modify Database Optional Information** section, you can define levels of categories and valid (default) attribute values for the database type you have chosen.

Common Database Options

You can set permanent values that correspond to categories, subcategories, and subsubcategories, which are used for organizing documents within views. If you do not supply default categories, users can supply them in the course of working in the database.

Administrators can set permanent values for category levels in any of the database types.

Database-specific Options

Within categories and their children, you can assign a set of valid attribute values. Each database type has a unique set of attributes that are made available in the database forms. For a Test Case Management or a Requirements database, the predefined attributes have been tailored to suit the intent of the database type. You can provide allowable values for each attribute provided.

Modify which documents your project will use

In the **Modify which documents your project will use** section, you can select the Notes documents your project will use and provide the document name, title and subtitle. The documents available for your database are determined by the database type. The setup forms come preloaded with default values, although you have the option to overtype existing field values, and turn documents and IBM-supplied subforms and document numbering on and off. See the database Previews (available through the *About* document) for a description of documents available for each database type.

Using this function, you can deactivate documents you don't need now and reactivate them at a later time if they are needed. For example, it is possible to deactivate **Design Change**, until a later stage in the project development cycle.

Modify the document hierarchy

In the **Modify the document hierarchy** section, you can identify valid

response documents for the base documents you defined in the **Modify which documents your project will use** section. This level of customization allows you to control how response documents can build upon a base document and other response documents.

Modify the states for documents

In the **Modify the states for documents** section, you can integrate the documents you defined in the **Modify which documents your project will use** section with a state system that best reflects your organization's workflow process. For each document type, the setup provides default state names that might be applicable to your integrated database. However, you can change the order of the state names, or the names themselves, if it provides a closer mapping to familiar terminology and processes.

Modify TeamConnection Access

The following areas of customization are described in detail in the Setup help.

Modify Family Information

The **Modify Family Information** section is very similar to the **Initialization** setup, except that you are not provided an opportunity to change the database type. The required fields will be primed with the current TeamConnection family configuration values, which were either defined at initialization or during a subsequent **Customization**. Whenever the family is modified, the TeamConnection component and configurable field information is automatically fetched from the current family server.

Refresh Components List

The **Refresh Components List** section enables you to retrieve a list of TeamConnection components in the TeamConnection family currently connected to your Notes Integrated Database. Select the **Get or Refresh** push button to refresh the list of components in the TeamConnection family linked to this database.

Refresh Feature Configurable Fields

The **Refresh Feature Configurable Fields** section enables you to retrieve a list TeamConnection configurable fields for features in the TeamConnection family currently connected to your Notes Integrated Database, and modify the list of attributes stored in your **Feature** documents.

Refresh Defect Configurable Fields

The **Refresh Defect Configurable Fields** section enables you to retrieve a list TeamConnection configurable fields for defects in the TeamConnection family currently connected to your Notes Integrated Database, and modify the list of attributes stored in your **Defect** documents.

Reconciliation of Notes and TeamConnection Data

The following areas of customization are described in detail in the Setup help.

Activate Reconcile

The Notes Integrated Database feature provides a reconciliation facility that synchronizes the data in the TeamConnection family and the Notes databases that use it. The reconciliation facility is an agent that should be run regularly. The reconciliation facility can be run on an established schedule that you set or manually when needed. We suggest that the reconciliation facility be set up to run after your regular TeamConnection build completes, so that you will have the latest document states reflected.

Note: See “Performing reconciliation” on page 118 for additional information about setting up and initiating the reconciliation facility.

Defects and Features

During reconciliation you can choose to have the TeamConnection defects and features that are in the cancel state deleted from the Notes database.

Log Document Options

You can choose to have the log that is created during reconciliation mailed to database users and/or stored in the database itself.

Advanced customization

The template that integrates TeamConnection and Lotus Notes was designed for maximum flexibility. The Setup document allows considerable database customization without requiring any skills in Notes database design. For users that wish to further customize their TeamConnection/Lotus Notes database to use customized Forms (Forms define the layout of Notes documents), so their documents are suited to the specific needs and requirements of their organization, it is possible to do this and still retain all the base functionality of the template. At a minimum this requires basic knowledge of Notes Forms design and Designer authority on the Notes database.

The customizable forms in the TeamConnection/Notes databases are called *subforms*. A subform is a form that is contained in a form. One or more subforms might make up a complete form that defines the layout of document. A default subform is supplied by IBM for each document that is supported. It is dynamically loaded when the main form is loaded.

For example, in a requirements database, the main form known as Document A loads a subform called fhcSm.A.Requirements. This subform has all the detail definition of a requirement. Document A has base document definition that is common to all documents.

Each document allows for several User subforms to be defined and used in addition to or in place of the IBM subform. A user subform can be defined that specifies the layout and content that more closely matches your organization than the corresponding IBM subform. A user subform is automatically loaded. The IBM subform is also loaded by default but this can be disabled using Setup. User subforms exist for most primary forms and help forms.

There are one or more subforms defined for all the Forms that allow user-defined subforms. This allows you to augment or replace the subform supplied by TeamConnection. The TeamConnection-supplied database teamc.nsf has all of the user subforms defined. All of them are empty. You should define your initial database by copying teamc.nsf so that you have all the subforms available (This process is described in “Initializing the original template and creating a database” on page 112.)

These user subforms are not modified if you refresh or replace your design database using the TeamConnection template fhcnote.nsf, so you can be assured that your subforms will remain intact if you update your database with a new version of this template (as described in “Database maintenance: refreshing design from a template” on page 119). User defined subforms follow a naming convention. All begin with the a prefix of UserSm. followed by a suffix that indicates what Form they are associated with. They are listed in the Subforms list of the Design section of the database. You cannot create subforms. You must use the empty subforms supplied.

You can see where these subforms are specified on the main Form by opening that form and locating one of the lines that display [Computed Subform]. Select the line and look in the bottom pane. The name of the user subform that is loaded at that location in the form is listed.

Note: See the **Advanced Customization** help topic for a detailed listing of forms and subforms available.

Part 4. Using TeamConnection to build applications

Chapter 10. Basic build concepts	129
The physical structure of the build function	129
The build object model	130
Parent-child relationships in a build tree	131
Working with a build tree	133
Putting the pieces together	135
Chapter 11. Installing, starting, and stopping build servers	137
Installing the build function	137
Creating a build server on MVS	137
Creating a build server on MVS/OE	140
Starting build servers using teamcbld	140
Starting an MVS build server	141
Starting the MVS/OE build server	143
Creating build startup files (for non-MVS environments)	144
Stopping the build servers	145
Stopping an MVS build server	145
Chapter 12. Working with build scripts and builders	147
Creating a builder	147
Writing a build script	150
Passing parameters to a build script	151
Writing a simple build script	153
Writing an executable file for a build script	153
Testing a build script	154
Modifying the contents of a build script	155
Putting a builder to work	155
Removing a builder from a part	156
Working with VisualAge C++ and Templates	157
Chapter 13. Working with MVS build scripts and builders	159
Creating a builder for MVS builds	159
Writing an MVS build script	163
File name conversions for MVS	164
Passing parameters to an MVS build script	164
TeamConnection syntax for MVS build scripts	165
Supported JCL syntax	166
EXEC statement	166
DD STATEMENT	167
Example of a build script for a C compile	168
Example of a build script for a COBOL compile	171
Example of a build script for a link	172
Chapter 14. Working with parsers	175
Creating a parser	175
Writing a parser command file	177

Putting a parser to work	178
Removing a parser from a part	179
Chapter 15. Building an application: an example	181
Starting the build servers	182
Creating builders and parsers	183
Creating the build tree for the application.	184
Starting the build on the client	189
Putting the build scripts to work	191
Finishing the job and reporting the results to the user	192
Monitoring the progress of a build	192
Running a build in spite of errors	193
Building all parts, regardless of build times	193
Finding out which parts will be built	194
Canceling a build	194
More sample build trees	195
Defining multiple outputs from a single build event	195
Synchronizing the build of unrelated parts	196

This section tells how to install and use the TeamConnection build function.

Though build administrators will be most interested in this section, anyone who builds an application using TeamConnection will find the first and last chapters helpful.

Chapter 10. Basic build concepts

This chapter defines terms and briefly describes the TeamConnection pieces that work together in building an application. For more details, continue to the other chapters in this section.

The TeamConnection build function has numerous features:

- It builds applications for platforms in addition to those it runs on. Currently you can build applications using TeamConnection on the following platforms: AIX, HP-UX, Solaris, MVS, MVS/OE, OS/2, Windows NT, and Windows 95.
- Its graphical representation of the structure of an application makes it easier to visualize and change.
- It lets you build an application using any number of machines working in parallel.
- Because it is fully integrated with TeamConnection's version control system, it ensures that the correct versions of parts are used in a build.
- It can work not only with parts that represent files, such as C source files, but also with parts that represent objects, such as VisualAge Generator applications.
- It can manage other steps related to software packaging and distribution.

For more information, see "Part 5. Using TeamConnection to package products" on page 199.

The physical structure of the build function

Figure 43 shows the structure of the TeamConnection build function:

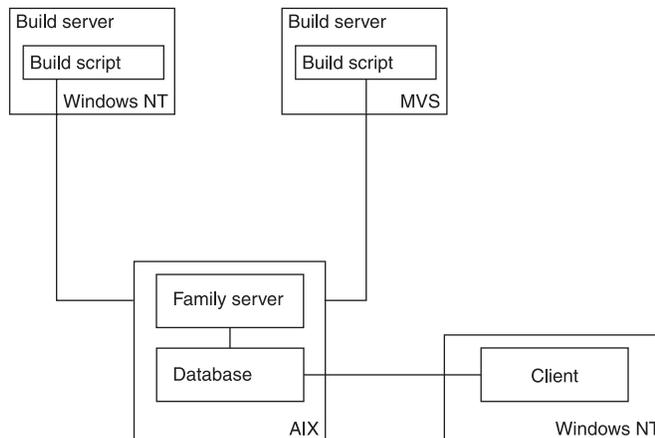


Figure 43. The physical structure of TeamConnection

Build servers are started by a TeamConnection administrator. For more information, see “Chapter 11. Installing, starting, and stopping build servers” on page 137.

The build object model

Figure 44 on page 133 shows the TeamConnection objects and events that constitute the build function, as illustrated in a sample application named msgcat.exe. This build object model is a conceptual model of the build function. When you use TeamConnection to define a build, you work with a *build tree* (a simplified graphical illustration of the build object model), which you can access through the TeamConnection GUI. “Working with a build tree” on page 133 explains build trees. This section explains the build objects and events represented in a build tree.

In TeamConnection, the build function is always described and discussed in terms of the final output of the build: the product or executable file that the build produces. For the sample application shown in this illustration, msgcat.exe is the build output and appears at the top of the build object model and as the top branch of the build tree illustrated on page 133. When you want to actually build the product, you request a build of msgcat.exe. TeamConnection uses the build tree that you define for this product to determine which objects and build events it needs to generate the final output. The objects and events that TeamConnection uses for a build include the following:

TeamConnection part

An object produced or used during a build, containing any data produced or used by the build. For example, a part called hello.c contains the source code for the application called msgcat. A part might be a text or binary file, or an object such as a VisualAge Generator generic collector.

Build event

An individual step in the build of an application, such as the compiling of hello.c into hello.obj.

A specific build request typically contains many build events. For example, if you start a build of an entire application, TeamConnection initiates build events for each compile and link operation.

Build requests are processed in a round-robin fashion for each TeamConnection family involved in a build. Build events are initiated in the order that they are received by each build machine involved in the build request, sometimes in parallel.

Build event processing is internal to TeamConnection; you cannot interact with these processes directly.

Builder

An object that can transform a build event's input parts into output parts by calling tools such as linkers or compilers. For example, one builder might know how to transform the input part hello.c into the output part hello.obj. A different builder might know how to transform hello.obj into msgcat.exe. Builders are associated with the parent, or output part, rather than the child, or input.

Build script

An object that a builder uses in transforming inputs to outputs; it is essentially a binding between TeamConnection and a transformation tool, such as a linker or compiler. In OS/2, Windows, UNIX, or MVS/OE environments, a build script is usually a command file, but it can be a string that calls the tool. In MVS, it is a file containing JCL-like instructions.

Parser A tool that can read a source file and report back a list of dependencies of that source file. It frees a developer from knowing the dependencies one part has on other parts to ensure a complete build is performed. For example, a C parser can read a C source code file and report back a list of the files included by the source file or by the included files.

TeamConnection will re-verify all parser dependencies:

1. When the user creates or checks in the part, TeamConnection will add all parser dependencies that it can find.
2. During build, TeamConnection will again check all parser dependencies and update as needed.

Parent-child relationships in a build tree

One relationship that is important to understand and distinguish is the relationship between parent and child parts in a build tree.

Though parent-child relationships usually imply that the parent part generates the child part, in a TeamConnection build it is the opposite. Because TeamConnection places the build output at the top of the tree, it refers to the build output as the parent and to the build input as the child.

A child part can be related to a parent part one of three ways: it can be an input part, an output part, or a dependent part.

Input parts

A part used as direct input to your build. An example of this is a C language source part. If you start a build and this part has changed, the changed part will be part of the new build.

Output parts

A generated output from a build, such as an OBJ or EXE part, or a part with no contents that serves as an organizer object. If you start a build and this part has changed, the changed part will be included in the new build.

Dependent parts

A part needed for the build operation to complete but that is not passed directly to the compiler. An example of this is an include part. If you start a build and this part has changed, the changed part will be included in the new build.

Though parent-child relationships usually imply that the parent part generates the child part, in a TeamConnection build it is the opposite. Because TeamConnection places the build output at the top of the tree, it refers to the build output as the parent and to the build input as the child.

To understand how build output is generated, it may be easier to start at the bottom of the build object model and work your way up. In Figure 44 on page 133, `hello.h` and `bye.h` are C source files that are embedded in `hello.c` and `bye.c`, respectively. The parser, `parser1`, is able to read `hello.c` and `bye.c` to determine files they embed. This build object model contains three build events:

- The builder `compiler1` compiles `hello.c` into `hello.obj`.
- The builder `compiler1` compiles `bye.c` into `bye.obj`.
- The builder `linker1` links `hello.obj` and `bye.obj` into `msgcat.exe`

This build object model contains the following parent-child relationships:

- `msgcat.exe` is the parent of `hello.obj` and `bye.obj`.
- `hello.obj` is the parent of `hello.c`
- `bye.obj` is the parent of `bye.c`

You establish these parent-child relationships between parts when you create the parts in TeamConnection.

Before you can build `msgcat.exe`, for example, you need to create a place-holder part for it and designate `linker1` as its builder. You then create place-holder parts for `hello.obj` and `bye.obj` and designate `compiler1` as their builder and `msgcat.exe` as their parent.

“Creating the build tree for the application” on page 184 walks you through an example of creating the build tree for this object model.

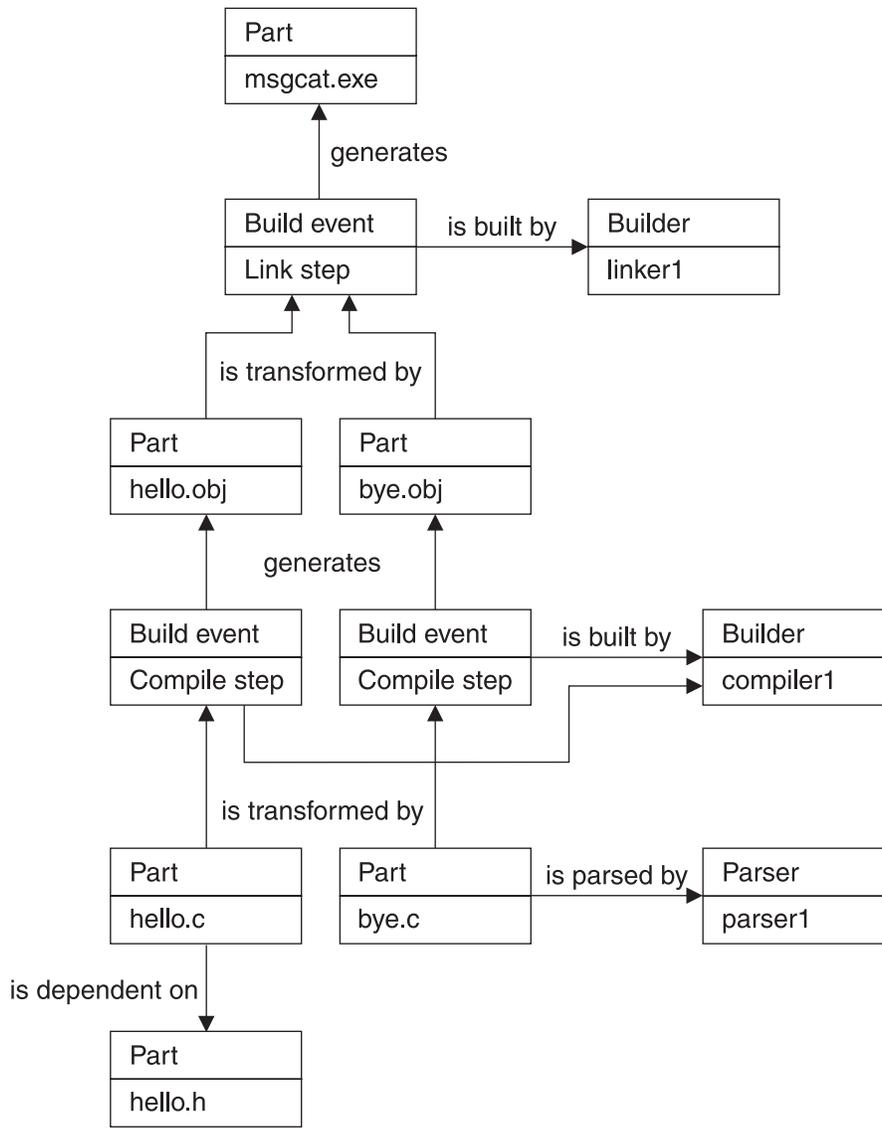


Figure 44. Sample build object model for msgcat.exe

Working with a build tree

Software developers must provide the information by which TeamConnection determines the build events that make up a build request. An application's build tree shows this information graphically.

A build tree is a simplified version of the build object model, showing the dependencies that the parts in an application have on one another. If you change the relationship of one part to another, the build tree changes accordingly. Figure 45 shows a build tree for the hello application:

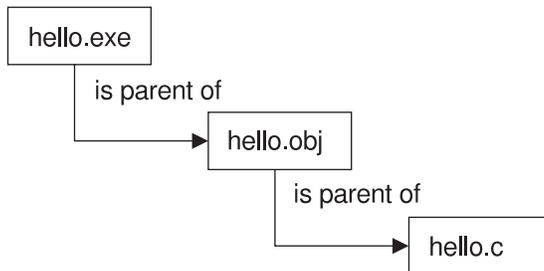
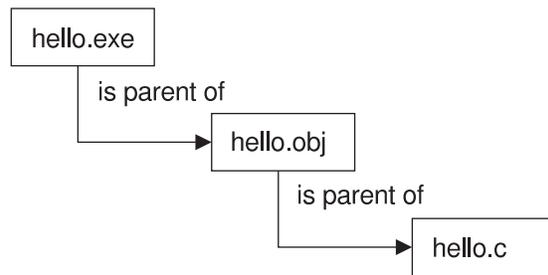


Figure 45. The build tree for the hello application

In this simple application, `hello.c` is compiled to create `hello.obj`, which in turn is linked to create `hello.exe`. The build tree shows that `hello.exe` is the parent of `hello.obj`, which in turn is the parent of `hello.c`. To build the entire application, you request to build `hello.exe`.

Just as the parts that make up an application are versioned, the relationships between these parts are versioned. Thus, more than one version of the build tree can exist. For example, Figure 46 on page 135 shows two different versions of the same build tree:

Work area 817



Work area 915

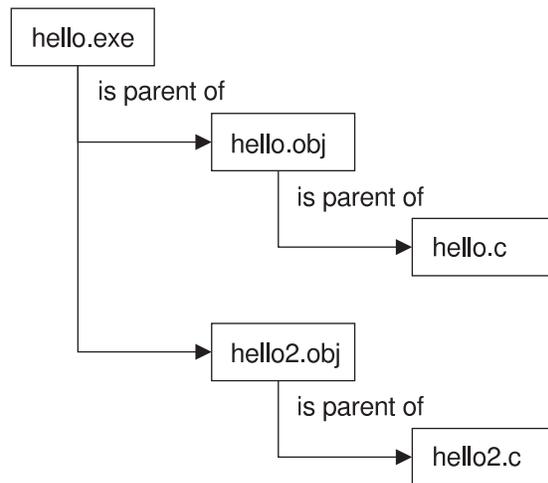


Figure 46. Two versions of a build tree

Putting the pieces together

The table that follows lists the tasks involved in preparing for building an application and in actually building it. Usually an administrator does the preparation steps, but anyone with the proper authority can do so.

For more information about this task,	Go to this page.
Creating build startup files	"Creating build startup files (for non-MVS environments)" on page 144
Starting build servers	"Chapter 11. Installing, starting, and stopping build servers" on page 137

For more information about this task,	Go to this page.
Stopping build servers	"Stopping the build servers" on page 145
Writing a build script	150
Creating a builder	147
Creating a parser	175
Defining a build tree	184
Starting a build	189
Stopping a build	194
Verifying the parts to be built	194

Chapter 11. Installing, starting, and stopping build servers

This chapter explains how to install, start, and stop a build server.

Installing the build function

Before installing build servers, you should be familiar with the build concepts found in “Chapter 10. Basic build concepts” on page 129.

For hardware requirements for the build server machines, refer to the requirements for your specific operating system in the *Installation Guide*.

When you install the various parts of TeamConnection, you can choose to group them on a single server machine, or you can distribute them in various combinations. For example, Figure 47 is a configuration that shows each component on a different machine:

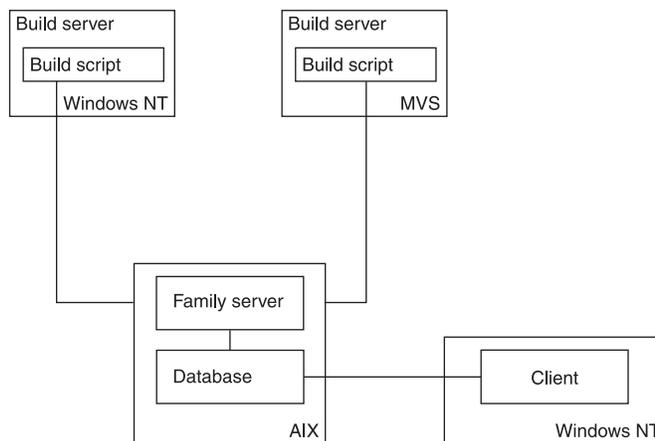


Figure 47. TeamConnection components on separate machines

The TeamConnection installation programs allow you to select specific components to install. You can use this program to install a build server, with or without other TeamConnection components. For instructions on installing TeamConnection, refer to the installation chapter for your platform in the *Installation Guide*.

Creating a build server on MVS

The code for an MVS build server is shipped with TeamConnection and installed on your TeamConnection server when you install the product.

The following are software requirements for the MVS build server:

- TCP/IP Version 3.2 for MVS
- OS/390 R3 LE

To install the build server on MVS, you create MVS data sets and then upload the TeamConnection files to these data sets. Follow these steps:

1. On your TeamConnection server, install the MVS build server component, following the instructions in the *Installation Guide*.
2. On MVS, create data sets with the following characteristics:
 - An object data set [1] to contain object files: LRECL=80, RECFM=FB, BLKSIZE=3120, DSORG=PO (Approximate size: 15 cylinders on a 3390.)
 - A load module data set [2] to contain TEAMCBLD and DLLs: LRECL=80,RECFM=U, BLKSIZE=3120, DSORG=PO (Approximate size: 15 cylinders on a 3390.)
 - A data set to contain JCL for creating load modules [3]: LRECL=80, RECFM=VB, BLKSIZE=3120, DSORG=PO
 - A JCL data set for the MVS build scripts and samples [4]: LRECL=80, RECFM=VB, BLKSIZE=3120, DSORG=PO
 - An environment variable data set for the EDCENV DDname in runpgm.jcl [5]:LRECL=80, RECFM=VB, BLKSIZE=3120, DSORG=PS (optional, needed at runtime)
3. From the mvs subdirectory where TeamConnection is installed, do the following to upload the files to MVS:
 - a. Type the following at a prompt and press Enter: ftp *hostname*. Specify your name and password, if required.
 - b. Type the following, and press Enter after each:
 - binary
 - cd *data set for object code* [1]
 - put fhccmnc.mvs fhccmnc
 - put fhcrscli.mvs fhcrscli
 - put teamcbld.mvs teamcbld
 - Issue only one of the following commands according to the platform from which you are transferring the MVS files:

From Intel platforms:

```
put nls\msg\${LANG}\fhbmsg.mvs fhbmsg
```

From UNIX platforms:

```
put /usr/teamc/mvs/${LANG}/FHBMSG.MVS fhbmsg
```

Note: For *LANG*, substitute the name of the language subdirectory, for example, enu for US English.
 - put fhbtclnk.mvs fhbtclnk
 - put getdsn.mvs getdsn
 - ascii

- cd *JCL data set for load module* [3]
 - put fhblink.jcl fhblink
 - put runpgm.jcl runpgm
 - put runpgmt.jcl runpgmt
 - put fhbmenv.var *environment variable data set* [5]
 - quit
- c. This optional step installs the sample build scripts on MVS. It must be done from a machine that has the TeamConnection client installed.

Note: The details in this step are subject to change. Please refer to updated documentation when attempting this in the future.

If you are doing these steps from a different machine than in the previous step, repeat step 3.a from the bin subdirectory where the client is installed on this machine. Otherwise, change to the bin subdirectory where the client is installed. Then, type the following statements and press Enter after each:

- ascii
 - cd '*data set for teamproc jcl*' [4]
 - put fhbmc.jcl fhbmc
 - put edcc.jcl edcc
 - put fhbmasm.jcl fhbmasm
 - put fhbplked.jcl fhbplked
 - put fhbcobm.jcl fhbcobm
 - put fhbmpi.jcl fhbmpi
 - put fhbtclnk.jcl fhbtclnk
 - put fhbm370.jcl fhbm370
 - put fhbmlink.jcl fhbmlink
 - quit
- d. From MVS, do the following:
- 1) Modify fhblink to customize the JCL to your MVS system.
 - 2) Submit this JCL to create the required load modules.
When this JCL is executed successfully, the following modules are created:
 - TEAMCBLD (main executable)
 - FHBMSG (national language-specific message catalog DLL)
 - FHCCMNC (supporting DLL)
 - FHCRSCLI (supporting DLL)

Note: All except for FHCCMNC are re-entrant.
 - 3) The object dataset [1] can be deleted.

Creating a build server on MVS/OE

The steps that follow describe how to install a build server on MVS/OE.

1. Install the MVS/OE build server component.
2. From the **oe** subdirectory where TeamConnection is installed, use ftp to upload the following files to MVS/OE, in binary.
 - rename `teamcbld.oe` to `teamcbld` (main executable)
 - rename `fhccmnc.oe` to `fhccmnc` (supporting DLL)
 - rename `fhcrscli.oe` to `fhcrscli` (supporting DLL)
 - `teamcv3.cat` (for required national language)
3. Using `chmod` in the OE shell, set the access flags for `teamcbld`, `fhccmnc`, and `fhcrscli` so that they are executable.

Starting build servers using `teamcbld`

You can start the build servers using the following line command:

```
teamcbld [-c buildDir] [-e environment]  
        [-p pool] [-f family]  
        [-u userID] [-l login_password]  
        [-k local_codepage] [-r remote_codepage]  
        [-b become] [-s] [-n]
```

Where:

- *buildDir* specifies the directory `teamcbld` should be run in.
- *environment* specifies the environment that you are building for, such as OS/2 or MVS. The value you specify here can be anything you like, but it must exactly match the environment specified for a builder in order for the builder to use this build agent. This value is case-sensitive. You can also set this value using the `TC_BUILDENVIRONMENT` environment variable.
- *pool* is the name of the build pool. You can also set this value using the `TC_BUILDPPOOL` environment variable. This value is case sensitive, and should match the parameter specified in the Part `-build` command.
- *family* is the name of the TeamConnection family. If a *family* is not specified, the value of the `TC_FAMILY` environment variable is used as the default.
- *userID* is used only when the family server authentication level requires a TeamConnection user ID and password. A password parameter is not necessary when `TC_LOGIN` has been used.
- *login_password* is used only when the family server authentication level requires a TeamConnection user password. A password parameter is not necessary when the command line `tclogin` command has been used.
- *local_codepage* local machine's code page used for translation, if not specified, no codepage conversion will be done.
- *remote_codepage* family machine's code page used for translation, if not specified, no codepage conversion will be done.

- *become* is a user become value. The default is the value specified for the TC_BECOME environment variable.
- **-s** sends log file messages to the screen. The build server generates a log file called teamcbld.log. Build server log messages can be routed to the screen using the **-s** parameter.
- **-n** writes files:
 - writes the changed environment variables to a file called **tcbldev.lst** instead of setting them in program's environment. The format of the file is variable=value.
 - writes the list of input files to a file called **tcblin.lst**. One file per line, format is pathName type.
 - writes the list of output files to a file called **tcblout.lst**. One file per line, format is pathName type.



You can also set the **-s** and **-n** build options using the TC_BUILDOPTS environment variable.

Two environment variables that directly affect build performance are:

- TC_BUILDMINWAIT - Minimum amount of time to wait (in seconds) between queries for new jobs. Default setting is 5, minimum setting is 3.
- TC_BUILDMAXWAIT - Maximum amount of time to wait (in seconds) between queries for new jobs. Default setting is 15, maximum setting is 300.

The command teamcbld will check the family for work to do every TC_BUILDMINWAIT. If it is not busy it will slow down to the TC_BUILDMAXWAIT time. The build administrator can adjust these variables to tune the frequency that the build server will check the family.

Starting an MVS build server

The MVS build server is a special client to the family server. It uses a user ID that must be defined to the family server. Depending on the authentication level in effect, a host list entry must be created for the user, and, if password authentication is required, the **-l** parameter must be used to specify the password when starting the build server.

The user ID used by the build server is the TSO user ID under which the build job is started. The user ID must always be in uppercase and must be created in uppercase on the family server. The user ID is not determined by the TC_USER environment variable.

You can start the MVS build server program TEAMCBLD, in Batch or under TSO. The RUNPGM JCL executes the MVS build server in batch. The RUNPGMT JCL executes the MVS build server under a TSO environment.

The following summarizes the actions to run the build server on MVS:

Modify the RUNPGM JCL for your installation.

Modify the environment variable dataset as needed. Two environment variables that directly affect build performance are:

- TC_BUILDMINWAIT - Minimum amount of time to wait (in seconds) between queries for new jobs. Default setting is 5, minimum setting is 3.
- TC_BUILDMAXWAIT - Maximum amount of time to wait (in seconds) between queries for new jobs. Default setting is 15, maximum setting is 300.

The teamcbld will check the family for work to do every TC_BUILDMINWAIT. If it is not busy it will slow down to the TC_BUILDMAXWAIT time. The build administrator can adjust these variables to tune the frequency that the build server will check the family.

Submit the JCL.

To start an MVS build server, do the following to modify the RUNPGM JCL:

1. Add a job card.
2. Modify the STEPLIB DD statement to point to the data set that contains the following load modules: TEAMCBLD, FHBMSG, FHCCMNC, and FHCRSCLI.
3. Modify the TEAMPROC DD statement to point to the data set that will contain all your MVS build scripts.

Note: The TEAMPROC DD defines the dataset into which the build script from TeamConenction family will be placed during the build. If the builder is of type 'none', then the build script does not exist in the family, but must already be in this dataset. If the build script refers to other build scripts (using the EXEC card), then those scripts are expected to be in this dataset also. However an optional ddname, PROCLIBS, can be defined which will be searched for the referred build scripts which are not found in TEAMPROC.

4. Modify the EDCENV DD statement to point to the data set that contains the environment variables.
5. Modify the following statement. For RUNPGMT, the statement is equivalent. Parameters can be specified using the PARM field or, in some cases, environment variables.

```
//RUNPGM EXEC PGM=TEAMCBLD,  
// PARM='ENVAR("_CEE_ENVFILE=DD:EDCENV")/[-e environment]  
//      [-p pool] [-f family] [-u unit_name]  
//      [-l login_password] [-k local_codepage]  
//      [-r remote_codepage]'
```

Where:

- *environment* specifies the environment that you are building for, such as OS/2 or MVS. The value you specify here can be anything you like, but it must exactly match the environment specified for a builder in order for the builder to use this build agent. This value is case-sensitive. You can also set this value using the TC_BUILDENVIRONMENT environment variable.

- *pool* is the name of the build pool. You can also set this value using the TC_BUILDPOOL environment variable. This value is case sensitive, and should match the parameter specified in the Part -build command.
- *family* is the name of the TeamConnection family. If a *family* is not specified, the value of the TC_FAMILY environment variable is used as the default.
- *unit_name* indicates the default unit type for dynamic data set allocations. VIO is the default.
- *login_password* is required only when the family server authentication level requires a TeamConnection user password.
- *local_codepage* indicates the code set that text data is converted to for the build. The default is IBM-1047.
- *remote_codepage* indicates the code set for data stored in TeamConnection. The default is IBM-850.

Notes:

- TEAMCBLD converts the text files between the two codepages using the **iconv** codepage conversion functions, and therefore can support only those codepages that have the iconv conversion tables installed on the MVS system. If not specified, TEAMCBLD uses internal tables which correspond to single-byte codepages IBM-1047 and IBM-850.
 - The -k (*local_codepage*) and the -r (*remote_codepage*) flags are also used to convert any messages generated in the MVS environment which are sent back to the workstation client.
6. Submit the job.

Starting the MVS/OE build server

Under MVS/OE, the build server program can be started with the following teamcbld command:

```
teamcbld -e environment -p pool -f family [-l login_password] [-s] [-v] [-n]
[-k local_codepage] [-r remote_codepage]
```

Where:

- *environment* specifies the environment that you are building for, such as OS/2 or MVS. The value you specify here can be anything you like, but it must exactly match the environment specified for a builder in order for the builder to use this build agent. This value is case-sensitive. You can also set this value using the TC_BUILDENVIRONMENT environment variable.
- *pool* is the name of the build pool. You can also set this value using the TC_BUILDPOOL environment variable. This value is case sensitive, and should match the parameter specified in the Part -build command.
- *family* is the name of the TeamConnection family. If a *family* is not specified, the value of the TC_FAMILY environment variable is used as the default.

- **-s** sends log file messages to the screen. The build server generates a log file called `teamcbld.log`. Build server log messages can be routed to the screen using the **-s** parameter. **-s** must be lowercase. An uppercase **-S** turns it off.
- **-v** increases the level of messages written to the log. This option is equivalent to the verbose option on TeamConnection `teamc` commands **-v** must be lowercase. An uppercase **-V** reduces the level.
- `login_password` is used only when the family server authentication level requires a TeamConnection user password. The user ID used by the build server is the MVS/OE user ID under which the build job is started.
- `local_codepage` indicates the code set that text data is converted to for the build. The default is IBM-1047.
- `remote_codepage` indicates the code set for data stored in TeamConnection. The default is IBM-850.

Notes:

1. The command `teamcbld` converts the text files between the two codepages using the `iconv` codepage conversion functions, and therefore can support only those codepages that have the `iconv` conversion tables installed on the MVS system. If not specified, `teamcbld` uses internal tables which correspond to single-byte codepages IBM-1047 and IBM-850.
2. The **-k** (`local_codepage`) and the **-r** (`remote_codepage`) flags are also used to convert any messages generated in the MVS environment which are sent back to the workstation client.

`TC_CATALOG`, if specified, should be the pathname of the message catalog file. For example, `/xxx/enu/teamcv3.cat`. If not specified, default (English) messages are used. The file `teamcbld` must be included in the `PATH` environment variable. The files `fhccmnc` and `fhcrscli` must be included in `LIBPATH`.

Two environment variables that directly affect build performance are:

- `TC_BUILDMINWAIT` - Minimum amount of time to wait (in seconds) between queries for new jobs. Default setting is 5, minimum setting is 3.
- `TC_BUILDMAXWAIT` - Maximum amount of time to wait (in seconds) between queries for new jobs. Default setting is 15, maximum setting is 300.

The command `teamcbld` will check the family for work to do every `TC_BUILDMINWAIT`. If it is not busy it will slow down to the `TC_BUILDMAXWAIT` time. The build administrator can adjust these variables to tune the frequency that the build server will check the family.

Creating build startup files (for non-MVS environments)

You can create startup files for concurrently starting build servers with the family server using the `teamcd` command. This is the preferred method for starting build servers. When starting the build servers in this manner, you need to create a startup file.

Information about the build servers is put in a startup file and the name of the startup file is specified in one of two ways:

- In the teamcd command, using the -b parameter.
- In the TC_BUILD_RSSBUILDS_FILE environment variable.

You can store the build startup files wherever you like, provided that you give the full file path names for them in the -b parameters, or in the TC_BUILD_RSSBUILDS_FILE environment variable.

Stopping the build servers

To stop a build server, do one of the following:

- Close the window in which the build server is running.
- Press Ctrl+C when the build server window has focus.
- Close the window in which the family server was started if the build server was started with the teamcd command.

Stopping an MVS build server

To stop an MVS build server, cancel the RUNPGM job that was used to start it.

Chapter 12. Working with build scripts and builders

A builder is an object that can transform one set of TeamConnection parts into another by invoking tools such as compilers and linkers. For example, one builder might transform a COBOL source file into an object file. Another might transform a set of object files into an executable file. Builders use build scripts to invoke the tools that actually transform TeamConnection parts.

Usually a build administrator creates build scripts and builders, but anyone with the proper authority can do so. For more information about the required authority, see Appendix I. Authority and notification for TeamConnection actions.

This chapter tells how to create and maintain build scripts and builders. It assumes that you have read “Chapter 10. Basic build concepts” on page 129. The following table directs you to the tasks to be done:

For more information about this task,	Go to this page.
Creating a builder	147
Writing a build script	150
Testing a build script	154
Updating a builder	155
Putting a builder to work	155
Removing a builder from a part	156

Creating a builder

As with most other TeamConnection operations, there are two ways you can create a builder: using the graphical user interface (GUI) or the command line interface.

To create a builder using the GUI:

1. From the Actions pull-down menu on the Tasks window, select **Builders** → **Create**.
2. On the Create Builder window, type the requested information.

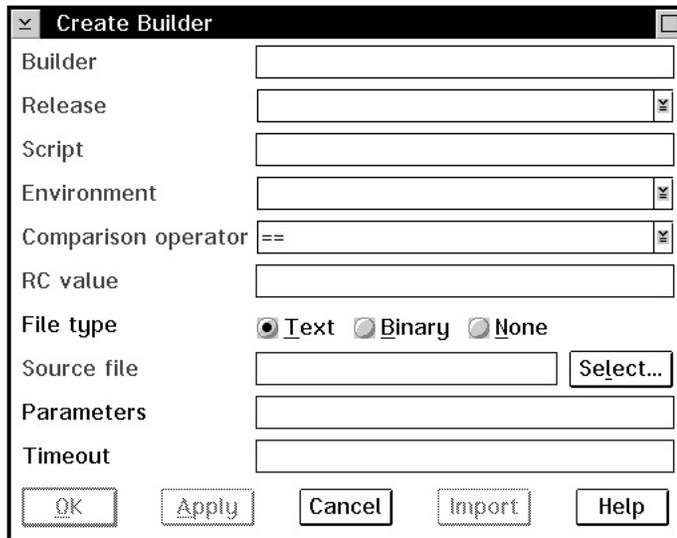


Figure 48. Create Builder window

To create a builder using the command line:

From a command line, type the `teamc builder -create` command and press Enter.

The complete command syntax is the following:

```
teamc builder -create name -condition RC_expression
                -environment name
                -from script_filespec
                -script name
                -value RC_value -release name
                -family name
                [-text | -binary | -none]
                [-parameters Parameters]
                [-timeout number] [-become user_name]
                [-verbose]
```

No matter which way you create a builder, you must specify a number of attributes for it. Together with the contents of the build script and the tools you use (the compilers, linkers, and so on), the following attributes define how a transformation takes place.

Builder

The name of the builder must be unique within a release. It can be anything you want; we recommend you establish and follow a meaningful naming convention. An example of a builder name is `c_set_2`.

Release

This is the name of the release that contains the builder. Builders are release-specific objects. They are not versioned within a release; therefore you can have only one version of a builder at any time in a release.

To use the builder from a previous release, you can link to a part that uses it in that release. This action copies the builder to the new release. Otherwise, you must create the builder again in the new release.

Script, File type, and Source file

These fields work together to define the build script that the builder invokes to accomplish the transformation. (The **File type** field on the GUI corresponds to `-text`, `-binary`, and `-none` in the command. The **Source file** field on the GUI corresponds to the `-from` attribute.)

- If the build script is simple enough to be expressed in one line, you can specify it in the **Script** attribute when you create the builder, and specify a file type of `none`. At minimum, the script must specify the name of the transformation tool. For example, to invoke the C Set/2 compiler, you might specify these values:

- **File type** - `none`
- **Script** - `icc`

See “Writing a simple build script” on page 153 for more information.

- If the build script is more complex, you must first create a separate file containing it; see “Writing an executable file for a build script” on page 153 for more information about how to write it. Specify the fully qualified path name of your file as the source file, and specify the file type as `text` or `binary`. TeamConnection can also detect the file type and store it in the proper format.

When the builder is created, this source file is stored as part of the builder in the TeamConnection database; during a build, the build server creates and runs a local version of this file. Specify the name you want for this local file in the **Script** field. For example, you might specify these values:

- **File type** - `text`
- **Script** - `c_compile.cmd`
- **Source file** - `c:\src\c_compile.cmd`

When this builder is created, the contents of `c:\src\c_compile.cmd` are stored in the builder. When this builder is invoked, TeamConnection creates a file named `c_compile.cmd`, writes the build script into this file, and then runs it.

- If the builder is being used to only collect a set of build objects (for example, the **ALL:** rule in makefiles), specify these values:

- **File type** - `none`
- **Script** - `null`

This prevents the build process from extracting input and output parts. See “Synchronizing the build of unrelated parts” on page 196 for an example.

Environment

This is the name of the environment supported by the builder, such as OS/2, Windows, AIX, HP-UX, Solaris, or MVS. The value that you specify here can be anything you like, but it must exactly match the environment value specified in the command used to start the build server. (See “Chapter 11. Installing,

starting, and stopping build servers” on page 137 for more information.) It is recommended that you follow a naming convention for this attribute, using values such as `os2` and `mvs`.

Comparison operator and RC value

Together, these two attributes make up a Boolean expression that defines the criteria used to decide whether a specific build event was successfully accomplished, when evaluated against the value returned by the build script. (The **Comparison operator** and **RC value** fields on the GUI correspond to the `-condition` and `-value` attributes in the command.)

The values allowed for **Comparison operator** are as follows:

- **EQ** or `==` - Equals
- **LT** or `<` - Less than
- **LE** or `<=` - Less than or equals
- **GT** or `>` - Greater than
- **GE** or `>=` - Greater than or equals
- **NE** or `!=` - Not equal to

RC value can be any positive integer. An example of a Boolean expression formed from these two attributes is `return_value LE 4`, meaning that the build event is considered a success if the build script returns a value less than or equal to four.

Parameters

This is a string passed to the build script as its argument. If the string includes blanks, enclose the string in double quotes. For example, for a builder used for VisualAge C++ compiling, you might specify a parameter string of `"/Ss /Ge-`. If the string includes a double quote, precede the double quote with a backslash (`\`). If the string includes a dash (`-`), precede the dash with a blank space, otherwise the string is interpreted as the start of a TeamConnection action flag.

Timeout

This attribute specifies the number of minutes that a build server will be allowed to complete a build event after it receives the build job from the TeamConnection family server. The default is 1440 minutes (24 hours). If the build event does not complete within this time, the build event fails.

Writing a build script

When you create a builder, you must specify a build script. The build script actually invokes the transformation tool and passes it parameters defined in the **Parameters** attribute of the builder.

Passing parameters to a build script

There are three places where parameters can be specified that affect the outcome of a build.

As attributes of a builder

Builder parameters are passed to the build script, after variable substitution is performed. Variables are substituted based upon the following syntax:

`$(variable_name)`

To pass parameters to your build script, specify them in the **Parameters** attribute of the builder. TeamConnection sets these variables before invoking the build script.

Note: If the command `teamcbld` included the `-n` flag, then the build script will process the `tcbldev.lst`, `tcblin.lst`, and the `tcblout.lst` files instead of the variables set by the **Parameters** attribute.

In UNIX environments, you need to include an escape character before the `$` character, for example: `\$(TC_INPUT)`.

You can use the following TeamConnection environment variables:

TC_BUILD_USER

The user ID of the TeamConnection user who submitted the build.

TC_FAMILY

The TeamConnection family.

TC_RELEASE

The release of the parts that are being built.

TC_LOCATION

The current directory where the build script runs.

TC_INPUT

A list of the TeamConnection parts that are input to the object being built.

TC_INPUTTYPE

Identifies each input type.

Note: There is a one-to-one relationship between each object in the `TC_INPUT` list and this list of types (TCPart, for example).

TC_OUTPUT

A list of the parts that are being built in this build event.

TC_OUTPUTTYPE

Identifies each output type. The default is file.

Note: There is a one-to-one relationship between each object in the TC_OUTPUT list and this list of types (TCPart, for example).

TC_WORKAREA

The name of the work area in which the build is being performed.

You can define other variables. These can be set when you start the build by specifying a value for **parameters** in the part -build command (from the command line or through the GUI). These variables are set in the parameters string passed to the build script.

These variables are also used to set environment variables before the build script is invoked.

As attributes of a part in the build tree

Parameters that are unique to a particular part are specified on the part -create and part -modify commands. Like the builder parameters, these parameters allow variable substitution.

When parameters are specified for a part, these parameters are used *in place* of the parameters specified for the builder. In other words, if both builder and part parameters are specified, the part parameters take precedence.

In addition, whenever parameters are specified for *any* part that is an output of a build event, they apply to *all* the outputs of that build event. For example, if a build event has two outputs, msg.exe and msg.map, then changing the part parameters to `"/Debug"` for either of the two parts has the same result. The next time the build event is processed, the `"/Debug"` parameter is used when invoking the build script that produces both msg.exe and msg.map.

You can also substitute the builder parameters into the file parameters by using the variable `$(BUILDERPARMS)`. For example, you might use the following command:

```
teamc part -build myfile.c -parameters "/Ti+ $(TC_BUILDERPARMS)" ...
```

At build time, the parameters specified in the builder for myfile.c are substituted for `$(TC_BUILDERPARMS)`.

As parameters of the part -build command

The part -build command parameters are not used the same way as the other two parameters. Instead, these parameters are used to set the values of environment variables that can be used for substitution into either the builder or part parameters. They are also set in the environment so they can be retrieved by the build script. In other words, they set up the environment used by the builder.

For example, if you issue a part -build command for msg.exe, you can specify `-parameters DEBUG=YES` and, inside of both the compile and link build scripts, retrieve the value of this variable from the environment, setting compiler or linker flags accordingly.

Writing a simple build script

This kind of build script is written into the **Script** attribute of the builder. When you create or modify the builder, you specify in this attribute the name of the transformation tool to be invoked.

For example, suppose you want to create a builder that compiles a C source file into a .exe file using IBM's VisualAge C++ compiler. You specify the following attributes for the builder:

Build script

```
icc
```

Parameters

```
"$(TC_INPUT) /Fe$(TC_OUTPUT)"
```

You can create this builder using the following command:

```
teamc builder -create c_builder -env OS2 -script icc -none  
-parameters "$(TC_INPUT) /Fe$(TC_OUTPUT)"
```

If you use this builder to create hello.exe from hello.c, the command actually issued during the build process is the following:

```
"icc hello.c /Fehello.exe"
```

Writing an executable file for a build script

Suppose you need to build a C application and you want to specify at build time whether to use debug information. To do this, you define in the builder parameters a variable called *debug* and set the variable when you start the build. In this case, you need a build script that is a separate executable file to pass the debug parameter after the variable substitution.

For a build script of this form, you first write a program or command file; this file is stored in the TeamConnection database when you create the builder. When a build is performed, this build script file is extracted from the database and run. It interprets the parameters passed to it and then invokes the actual transformation tool, such as the compiler.

Our earlier example describes a builder that compiles a C source file into a .obj file using IBM's VisualAge C++ compiler. Using this builder, you can specify at build time whether to use debug information. Here is the complete build script for such a builder, written in IBM's REXX language (it could just as easily have been written in C or COBOL).

```
/* sample C Build Script using debug flag */  
parse arg parms  
  
  environ = 'OS2ENVIRONMENT'  
  input   = VALUE('TC_INPUT',,environ)  
  output  = VALUE('TC_OUTPUT',,environ)
```

```

debug    = VALUE('DEBUG',,environ)

if debug = 'YES' then
do
    parms = parms || '/Ti+'
end
icc parms '/Fo' || output input

```

exit result



Windows NT and 95 build scripts must be able to return a value for a return code. Because *.bat command files provide little support for programming logic and cannot return a value, use a compiled executable for your build script. TeamConnection provides two sample Windows build scripts and their source files. These samples, fhbwcomp.exe and fhbwlink.exe, are C programs for the Microsoft Visual C++ compiler and linker, respectively. Because these samples are C programs, they can also be used with the OS/2 build server with only slight modifications.

You can create the builder that invokes this build script using the following command:

```

teamc builder -create c_builder2 -script c_compile.cmd -parameters "/c"
    -from d:\teamc\c_compile.cmd

```

Where d:\teamc\c_compile.cmd is the file to be stored in the TeamConnection database and c_compile.cmd is the name of the local file that the build process creates and runs during a build.

To build hello.obj using the debug option, you use the following command:

```

teamc part -build hello.obj -parameters "debug=YES" -pool os2pool

```

The command issued by the build server is the following:

```

c_compile.cmd /c

```

In turn, the build script inspects the contents of the parameters it received in its argument list and from the environment, and it forms this command:

```

"icc /c /Ti+ /Fohello.obj hello.c"

```

Testing a build script

The easiest way to test a build script is to write a simple driver program that sets the environment variables that the build script will expect and then runs the script against local files.

For example, to test the example build script in “Writing an executable file for a build script” on page 153, write a program that sets the TC_INPUT, TC_OUTPUT, and DEBUG parameters, and then runs the command file against a local copy of hello.c. If the test is successful, the script correctly builds hello.obj in the current directory, and DEBUG is interpreted correctly.

Modifying the contents of a build script

Sometimes you need to modify the contents of a build script. Remember that a build script is stored as part of the builder itself. Because builders are not versioned, you do not check them out as you would most TeamConnection parts. Instead, follow these steps:

1. Extract the builder (in which the build script is stored) from the TeamConnection database.
2. Make your changes at your workstation.
3. Store the contents back into the TeamConnection database by using the builder `-modify` command.

For example, to modify the build script in “Writing an executable file for a build script” on page 153, you first issue the following command:

```
teamc builder -extract c_builder2 -to d:\build\c_builder2
```

Then, you use an editor to update `d:\build\c_builder2`. To move the updated build script back into TeamConnection, you issue the following command:

```
teamc builder -modify c_builder2 -from d:\build\c_builder2
```

The builder is an implied dependency for any part that uses it. Therefore, the next time you build the application that uses the modified builder, all the parts that use it get rebuilt.

Putting a builder to work

For an application to use a builder, the builder must be attached to the TeamConnection parts that it will build.

For an existing part, do one of the following:

- **GUI:** From the Actions menu of the TeamConnection Tasks Window, select **Parts** → **Modify** → **Properties**. On the Modify Part Properties window, type the name of the builder in the **Builder** field.

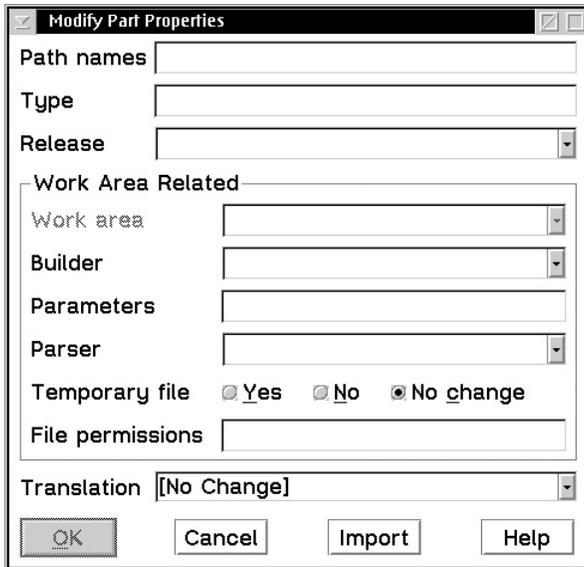


Figure 49. Modify Part Properties window

- From a command line, type the following and press Enter.
`teamc part -modify name -Builder name`

where the part *name* is the name of the output file to be created by this builder and the builder *name* is the name of the builder itself.

The complete syntax for this command is described in the *TeamConnection Commands Reference*.

You can also attach a builder to an output file when the part is created.

After you attach a builder to a part, the builder is ready for action. When the part is built, the builder invokes the build script, which in turn invokes a tool to transform the inputs of the part into the output.

For more information about attaching builders to the build tree, refer to “Creating the build tree for the application” on page 184.

Removing a builder from a part

If you no longer want to use a builder for a part, do one of the following:

- From the GUI, select **Parts** → **Modify** → **Properties** from the Actions menu of the TeamConnection Tasks window. On the Modify Part Properties window, type null in the **Builder** field.



Figure 50. Modify Part Properties window

- From a command line, type the following:
`teamc part -modify name -builder null -release name -family name`

Working with VisualAge C++ and Templates

When using VisualAge C++ and templates, template-include objects are saved in a subdirectory of the current directory called TEMPINC, so that subsequent builds can use them. When you start a build from TeamConnection, you need to specify the /Ft(dir) parameter with your builder or use PRAGMA statements to update the template-include objects for subsequent builds. This parameter suppresses resolution of files and imbeds them within the object file.

You can specify the /Ft(dir) parameter with a builder as follows:

```
teamc builder -create c_builder -script icc -parameters "/FtE:\template"
```

Chapter 13. Working with MVS build scripts and builders

A builder is an object that can transform one set of TeamConnection parts into another by invoking tools such as compilers and linkers. For example, one builder might transform a COBOL source file into an object file. Another might transform a set of object files into an executable file. Builders use build scripts to invoke the tools that actually transform TeamConnection parts.

For MVS, a build script is a text file that contains JCL-like statements with additional TeamConnection syntax and substitutable variables. TeamConnection parses these statements and does the necessary allocations and program calls for a build.

Note: The builder type cannot be binary.

Usually a build administrator creates build scripts and builders, but anyone with the proper authority can do so. For more information about the required authority, see Appendix I. Authority and notification for TeamConnection actions.

This chapter tells how to create MVS build scripts and builders. It assumes that you have read “Chapter 10. Basic build concepts” on page 129. The following table directs you to the tasks to be done. In some cases, if the instructions are the same for OS/2 and MVS, the table refers you to topics in “Chapter 12. Working with build scripts and builders” on page 147.

For more information about this task,	Go to this page.
Creating a builder for MVS builds	159
Writing an MVS build script	163
Testing a build script	154
Updating a builder	155
Putting a builder to work	155
Removing a builder from a part	156

Creating a builder for MVS builds

As with most other TeamConnection operations, there are two ways you can create a builder: using the graphical user interface (GUI) or the command line interface.

To create a builder using the GUI:

1. From the Actions pull-down menu on the Tasks window, select **Builders** → **Create**.
2. On the Create Builders window, type the requested information.

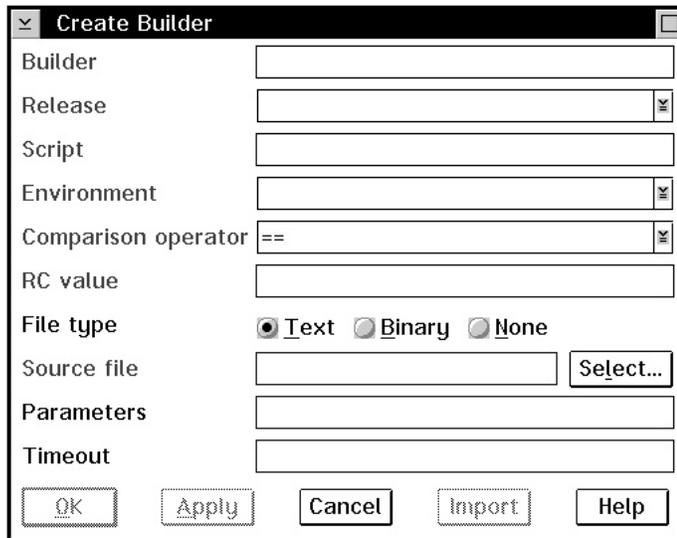


Figure 51. Create Builder window

To create a builder using the command line:

From an OS/2 command line, type the builder -create command and press Enter.

The complete command syntax is the following:

```
teamc builder -create name -condition RC_expression
               -environment name
               -from script_filespec
               -script name
               -value RC_value -release name
               -family nName
               [-text | -binary | -none]
               [-parameters parameters]
               [-timeout number] [-become user_name]
               [-verbose]
```

No matter which way you create a builder, you must specify a number of attributes for it. Together with the contents of the build script and the tools you use (the compilers, linkers, and so on), the following attributes define how a transformation takes place.

Builder

The name of the builder must be unique within a release. It can be anything you want; we recommend you establish and follow a meaningful naming convention. An example of a builder name is c370.

Release

This is the name of the release that contains the builder. Builders are release-specific objects. They are not versioned within a release; therefore you can have only one version of a builder at any time in a release.

To use the builder from a previous release, you can link to a part that uses it in the previous release. This action copies the builder to the new release. Otherwise, you must create it again in the new release.

Script, File type, and Source file

These fields work together to define the build script that the builder invokes to accomplish the transformation. (The **File type** field on the GUI corresponds to `-text`, `-binary`, and `-none` in the command. The **Source file** field on the GUI corresponds to the `-from` attribute in the command.)

You must first create a separate file containing the build script. All MVS build scripts must be written using JCL statements and the TeamConnection syntax described in “Writing an MVS build script” on page 163. You can store the build script one of two ways:

- **To store the build script as part of the builder:** specify the fully qualified path name of your build script file as the source file, and specify the file type as `text`. When the builder is created, this source file is stored as part of it in the TeamConnection database.

During a build, the build server creates and runs a local version of this file. Specify the name you want for this local file in the **Script** field. For example, you might specify these values:

File type

`text`

Script `fhbc`

Source file

`C:\build\script\fhbc.jcl`

When this builder is created, the contents of `C:\build\script\fhbc.jcl` are stored in the builder. When this builder is invoked, TeamConnection creates a file named FHBC in the data set referenced by the TEAMPROC ddname, writes the build script into this file, and then runs it.

- **To store the build script on MVS:** create the build script file and place it in the data set allocated to the TEAMPROC ddname in the RUNPGM JCL file. When you do this, specify the following attributes:

File type

`none`

Script `fhbc`

Do not specify a source file.

- If the builder is being used to only collect a set of build objects (for example, a VisualAge Generator collector part), specify these values:

File type

`none`

Script `null`

See “Synchronizing the build of unrelated parts” on page 196 for an example.

Environment

This is the name of the environment supported by the builder, such as MVS.

The value that you specify here can be anything you like, but it must exactly match the environment value specified in the command used to start the build server.

It is recommended that you follow a naming convention for this attribute, using values such as `os2` and `mvs`.

Comparison operator and RC value

Together, these two attributes make up a Boolean expression that defines the criteria used to decide whether a specific build event was successfully accomplished, when evaluated against the value returned by the build script.

The **Comparison operator** and **RC value** fields on the GUI correspond to the `-condition` and `-value` attributes in the command.

The values allowed for these operators:

- **Comparison operator** are as follows:
 - **EQ or ==** - Equals
 - **LT or <** - Less than
 - **LE or <=** - Less than or equals
 - **GT or >** - Greater than
 - **GE or >=** - Greater than or equals
 - **NE or !=** - Not equal to
- **RC value** can be any positive integer.

An example of a Boolean expression formed from these two attributes is `return_value LE 4`.

This expression means that the build event is considered a success if the build script returns a value less than or equal to four.

Parameters

This is a string passed to the build script as its argument.

For example, for a builder used for linking load modules, you might specify a parameter string of `list,test`.

Timeout

This attribute specifies the number of minutes that a build server will wait for an invoked build script to return before concluding an error has occurred and stopping the build event.

If the timeout value is reached, the build event fails.

Because MVS builds are processed in batch mode but the build is submitted to the build server in real time, consider writing a user exit to check the time of day before allowing a build request to be submitted. Another approach to

handling the timing of MVS builds is to start the MVS build server only at night and ensure that the MVS builders do not have short timeout values.

Writing an MVS build script

The best starting point for an MVS build script is an existing JCL fragment that is used for transforming inputs into outputs. For example, suppose you want to create a builder that compiles a C source file into an OBJECT file using IBM's C/370 compiler. You probably already have JCL that can be submitted as a batch job that does this.

When you create a build script for the MVS environment, you specify JCL statements with additional TeamConnection syntax. This build script is parsed by the build server. From the parsed results, TeamConnection allocates the specified ddnames and data sets; it then determines and executes the programs dynamically. The MVS build server uses the specialized TeamConnection syntax in the JCL to determine where to store the parts involved in an MVS build.

All statements in the MVS build script except for comments and inline data stream must start with two forward slashes (//).

Before you start writing your build script, refer to the manuals for the compiler, linker, or other transformation program to determine the data set requirements. Pay particular attention to the DCB attributes for LRECL, BLKSIZE, and RECFM.

Sample build scripts shipped with TeamConnection can be installed on MVS. Page 298 lists the sample build scripts. For instructions on installing these samples, refer to the *Administrator's Guide*

If you are debugging a build script, these manuals are also the first place to look for problems.

For more information about JCL syntax, refer to the *JCL User's Guide* and *JCL Reference* for your version of MVS. (These are listed in the bibliography at the back of this book.)

The following sample MVS build scripts are shipped with TeamConnection:

fhbasm.jcl

Calls the MVS assembler

fhbcobm.jcl

Calls the MVS COBOL compiler

fhbmpli.jcl

Calls the PL/1 MVS compiler

File name conversions for MVS

TeamConnection file names are modified by the MVS build server according to the following rules:

- The directory path of a file name is not used. All characters of a file name up to and including the rightmost slash (/ or \) are thrown away.
- Lowercase characters are converted to uppercase characters.
- The file extensions are stripped from the right, up to and including the leftmost period. The extension, minus the period, is used by the MVS build tool to direct the file to particular data sets according to user-specified syntax in the MVS build scripts.
- The remaining name is truncated from the left, to a maximum of 8 characters.
- Names must contain characters that are valid in MVS. MVS allows the following characters:

0123456789ABCDEFGHIJKLMNOPQRSTUVWXYZ\$@#

However, the name must begin with an alphabetic character.

- Underscore characters (_) in a base name are converted to at signs (@).

The following are examples of how a TeamConnection name is converted:

- A TeamConnection file name of src\build\fhblobj.C is converted to FHBLDOBJ on MVS.
- A TeamConnection file name of src/build/fhbtruncate.c is converted to FHBTRUNC on MVS.

In both of these examples, the .C or .c is split away. The MVS build server uses the resulting extension to resolve and possibly allocate the MVS data sets needed for the build process. The extensions are required for parts that participate in an MVS build.

A TeamConnection file name of src\build\fhbtest.c.old is converted to FHBTEST, and c.old becomes the extension.

Passing parameters to an MVS build script

To pass parameters to your build script, specify them in the **Parameters** attribute of the builder. These are passed to MVS through the combination of the PARM keyword parameter on an EXEC card and the &TCPARM variable.

Note: Take extra care to use no single or double quotes in the **Parameters** attribute of the MVS builder definition. This rule follows standard JCL syntax for parameter substitution in the PARM keyword parameter of an EXEC statement.

You can use the **&TCPARM** variable in your MVS build scripts. This variable is substituted with any parameters that were specified using the -parameter attribute of the builder command or the **Parameters** field on the Create Builder window when the builder was created.

You can also use the following TeamConnection variables in writing MVS build scripts:

&TCBLDUSR

On MVS, can be used in JCL scripts, which will be substituted with the userID before processing the script.

&TCINPUT

This variable is used for in-stream data. For each build input, the line where &TCINPUT appears is duplicated and the variable &TCINPUT substituted with the input name.

&TCOUTPUT

This variable is used for in-stream data. For each build output, the line where &TCOUTPUT appears is duplicated and the variable &TCOUTPUT substituted with the output name.

&TCWKAREA

The name of the work area in which the build is being performed.

&TCRELEAS

The name of the release in which the build is being performed.

Note: The &TCINPUT and &TCOUTPUT substitutable variables have limited scope in the MVS build scripts and should be used only within the in-stream data.

You can define other variables. You can set them by specifying a value for **Parameters** when you start a build. These variables are set in the parameters string passed to the build script.

Further, these variables can be used for variable substitution within MVS build scripts. Variable substitution works similarly to JCL variable substitution.

TeamConnection syntax for MVS build scripts

TeamConnection has extended the existing JCL syntax. The extended syntax tells the TeamConnection build server where to put the inputs, where to get the outputs, and where to get messages from the translators after an MVS build.

To direct inputs, outputs, and messages, add TCEXT=xxx to the data set attributes defined to a ddname, where xxx is one of the following:

- The base name extension from the TeamConnection part — for example, TCEXT=H, where H is the extension from A.H.
- One or more base name extensions from TeamConnection parts, surrounded by parentheses — for example, TCEXT=(H,HPP), where H is an extension from A.H or HPP is an extension from A.HPP.
- The string TCOUT, which declares that the contents of the data set assigned to the ddname will be sent back to TeamConnection. Users can view this information in one of these ways:
 - From a command line prompt, typing `teamc part name -viewmsg` and pressing Enter

- Selecting **Part** → **View** → **View build message** from the Actions pull-down menu on the Tasks window

Note: More than one ddname can specify TCOUT; the results are concatenated in the order of appearance.

When you add the TCEXT attribute for a ddname specification, you must also specify other attributes to allocate the data set through dynamic allocations:

- SPACE
- UNIT
- DCB, which includes the LRECL, BLKSIZE, and RECFM attributes

The UNIT attribute defaults to VIO unless the -U parameter is specified when the MVS build server is started.

For translation messages, you can allocate a data set to the ddname TC\$LIST and specify the attributes yourself. Otherwise, the build server allocates this data set with the following attributes by default:

```
//TC$LIST DD DCB=(RECFM=VB,LRECL=255,BLKSIZE=32640),  
// SPACE=(CYL,(2,1)),DISP=(NEW,DELETE),UNIT=VIO
```

Supported JCL syntax

The TeamConnection MVS build server supports only a subset of the available JCL syntax.

The following are not supported:

- A JOBSTEP statement
- DISP=(.,PASS)...

JCL procedures can be used on an EXEC statement. However, you must verify that any procedure called by the build script uses syntax that TeamConnection supports.

The following list indicates the positional and keyword parameters that are supported. You can verify the syntax in the *JCL Reference*.

EXEC statement

```
//label EXEC positional_parameter,keyword_parameter
```

The following parameters are supported.

Positional parameters:

- PGM=*program_name*, where *program_name* is an executable load module
- PROC=*procedure_name*, where *procedure_name* is an existing JCL procedure
- *procedure_name*, where *procedure_name* is an existing JCL procedure

Keyword parameters:

- PARM=*information*, where *information* is the parameter string passed to the load module.
- COND=(*code,operator [,stepname]*)
 - *code* is the value to test against the return code from a previous step
 - *operator* is the comparison to be made between the value for *code* and the return code
 - *stepname* is the step issuing the return code

All other keyword parameters are ignored and not used.

DD STATEMENT

```
//label DD keyword parameter
```

Positional parameters

The only supported positional parameter is [*], which begins an in-stream data set containing no JCL.

Keyword parameters

The following keywords are supported.

- DSN=*data_set_name* or DSNAME=*data_set_name*
- DISP=*status* or DISP=(*[status] [,normal-termination-disp] [,abnormal-termination-disp]*)
 - Valid values for *status* are NEW, OLD, SHR, or MOD.
 - Valid values for *normal_termination_disp* or *abnormal_termination_disp* are DELETE, KEEP, CATLG or UNCATLG.
- UNIT=*unit_type*, where *unit_type* is any value allowed in JCL. The default is VIO unless a different default is set when the MVS build server is started.
- SPACE=(*allocation_type*, (*primary* [, *secondary*] [, *directory*]) [,RLSE] [,CONTIG])
 - Valid values for *allocation_type* are TRK, CYL, or the block size.
 - *primary* is the primary number of the allocation type.
 - *secondary* is the secondary number of the allocation type.
 - *directory* is the number of directory blocks for a partitioned data set.
- DCB=(LRECL=*record_length*,BLKSIZE=*block_size*,RECFM=*record_format*)
Valid values for *record_format* are F, FB, V, VB, or U).
- DSORG=*data_set_organization*
Valid values for *organization* are the following:
 - P0 for a partitioned data set
 - PS for a sequential data set

- DDNAME=*label*, where *label* is the later ddname label reference. This parameter is supported only for simple cases.
- SYSOUT=*class*
This will always be allocated as a DUMMY DSN.

All other keyword parameters are ignored and not used.

Example of a build script for a C compile

The following JCL can be submitted as a batch job to do the following:

- Compile the source file member in the data set WELLSK.TEAMC.C
- Produce an object file member in the data set WELLSK.TEAMC.OBJ
- Produce a listing of the source file in the file member in the data set WELLSK.TEAMC.LISTING
- List the compiler messages in the file member in the data set WELLSK.TEAMC.ERROR

```

//COMPILE EXEC PGM=EDCCOMP,
// PARM='LO,SSCOMM,NOSEQ,NOMAR,LIS,FL(I),SO,DECK,TEST',
//      REGION=1536K
//STEPLIB DD DSN=SYS1.EDC.SEDCCOMP,DISP=SHR
//        DD DSN=SYS1.EDC.SEDCLINK,DISP=SHR
//        DD DSN=SYS1.PLI.SIBMLINK,DISP=SHR
//SYSMMSG DD DSN=SYS1.EDC.SEDCDMSG(EDCMSGE),DISP=SHR
//SYSIN   DD DSN=WELLSK.TEAMC.C(MEMBER),DISP=SHR
//USERLIB DD DSN=WELLSK.TEAMC.H,DISP=SHR
//SYSLIB  DD DSN=SYS1.EDC.SEDCHDRS,DISP=SHR
//SYSPUNCH DD DSN=WELLSK.TEAMC.OBJ(MEMBER),DISP=SHR
//SYSLIN  DD SYSOUT=*
//SYSPRINT DD DSN=WELLSK.TEAMC.ERROR(MEMBER),DISP=SHR
//SYSCPRT DD DSN=WELLSK.TEAMC.LISTING(MEMBER),DISP=SHR
//SYSUT1  DD UNIT=VIO,DISP=(NEW,DELETE),
//        SPACE=(32000,(30,30)),DCB=(RECFM=FB,LRECL=80,BLKSIZE=3200)
//SYSUT4  DD UNIT=VIO,DISP=(NEW,DELETE),
//        SPACE=(32000,(30,30)),DCB=(RECFM=FB,LRECL=80,BLKSIZE=3200)
//SYSUT5  DD UNIT=VIO,DISP=(NEW,DELETE),
//        SPACE=(32000,(30,30)),DCB=(RECFM=FB,LRECL=3200,BLKSIZE=12800)
//SYSUT6  DD UNIT=VIO,DISP=(NEW,DELETE),
//        SPACE=(32000,(30,30)),DCB=(RECFM=FB,LRECL=3200,BLKSIZE=12800)
//SYSUT7  DD UNIT=VIO,DISP=(NEW,DELETE),
//        SPACE=(32000,(30,30)),DCB=(RECFM=FB,LRECL=3200,BLKSIZE=12800)
//SYSUT8  DD UNIT=VIO,DISP=(NEW,DELETE),
//        SPACE=(32000,(30,30)),DCB=(RECFM=FB,LRECL=3200,BLKSIZE=12800)
//SYSUT9  DD UNIT=VIO,DISP=(NEW,DELETE),
//        SPACE=(32000,(30,30)),DCB=(RECFM=VB,LRECL=137,BLKSIZE=882)
//SYSUT10 DD SYSOUT=*
//

```

Figure 52. A JCL fragment for an MVS compile

The first step in converting the JCL fragment is to recognize the intent for each of the data sets and ddnames. For this C/370 compiler example, the SYSIN ddname needs to be associated with the source file, the SYS PUNCH ddname needs to be associated with the object file, and so on.

In each of these cases, the build script must tell the TeamConnection build server where to put or pick up the parts before and after the execution of the specified program (PGM=EDCCOMP).

Assume that your source files in TeamConnection have the extension .c, your object files have .obj, and your include files .h or .hpp. You allocate a data set to the SYSIN ddname to contain a source file with a .c extension. You specify the DCB, UNIT, DISP, and SPACE attributes to dynamically create this data set every time this build script is invoked. Notice that the attribute SPACE=(TRK,(10,5)) indicates a sequential data set organization.

You specify the output messages that will be returned to TeamConnection by using the TCOU attribute. This attribute tells the MVS build server to return the information in the data set associated with the TCEXT=TCOU attribute.

Note: The STEPLIB is renamed by the MVS build server to STEPLIBB for data set lookup of the program specified by the PGM parameter on an EXEC statement.

The following MVS build script is the result of converting the JCL fragment by adding the TeamConnection MVS JCL syntax.

```

//COMPILE EXEC PGM=EDCCOMP,
// PARM='LO,SSCOMM,NOSEQ,NOMAR,LIS,FL(I),SO,DECK,&TCPARM',
//      REGION=1536K
//STEPLIB DD DSN=SYS1.EDC.SEDCCOMP,DISP=SHR
//        DD DSN=SYS1.EDC.SEDCLINK,DISP=SHR
//        DD DSN=SYS1.PLI.SIBMLINK,DISP=SHR
//SYSMSGSGS DD DSN=SYS1.EDC.SEDCDMSG(EDCMSGE),DISP=SHR
//SYSIN DD TCEXT=(C, CPP),DISP=(NEW,DELETE),
//       UNIT=SYSDA,SPACE=(TRK,(10,5)),
//       DCB=(RECFM=VB,LRECL=150,BLKSIZE=3200)
//USERLIB DD TCEXT=(H,HPP),DISP=(NEW,DELETE),
//        UNIT=VIO,SPACE=(TRK,(5,10,10)),
//        DCB=(RECFM=VB,LRECL=50,BLKSIZE=3200)
//SYSLIB DD DSN=SYS1.EDC.SEDCHDRS,DISP=SHR
//SYSPUNCH DD TCEXT=OBJ,DISP=(NEW,DELETE),
//          UNIT=VIO,SPACE=(TRK,(10,5)),
//          DCB=(RECFM=FB,LRECL=80,BLKSIZE=3200)
//SYSLIN DD SYSOUT=*
//SYSRINT DD TCEXT=TCOUT,DISP=(NEW,DELETE),
//        SPACE=(32000,(30,30)),UNIT=VIO,
//        DCB=(RECFM=VB,LRECL=137,BLKSIZE=882)
//SYSCPRT DD TCEXT=TCOUT,DISP=(NEW,DELETE),
//        SPACE=(32000,(30,30)),UNIT=VIO,
//        DCB=(RECFM=VB,LRECL=137,BLKSIZE=882)
//SYSUT1 DD UNIT=VIO,DISP=(NEW,DELETE),
//        SPACE=(32000,(30,30)),DCB=(RECFM=FB,LRECL=80,BLKSIZE=3200)
//SYSUT4 DD UNIT=VIO,DISP=(NEW,DELETE),
//        SPACE=(32000,(30,30)),DCB=(RECFM=FB,LRECL=80,BLKSIZE=3200)
//SYSUT5 DD UNIT=VIO,DISP=(NEW,DELETE),
//        SPACE=(32000,(30,30)),DCB=(RECFM=FB,LRECL=3200,BLKSIZE=12800)
//SYSUT6 DD UNIT=VIO,DISP=(NEW,DELETE),
//        SPACE=(32000,(30,30)),DCB=(RECFM=FB,LRECL=3200,BLKSIZE=12800)
//SYSUT7 DD UNIT=VIO,DISP=(NEW,DELETE),
//        SPACE=(32000,(30,30)),DCB=(RECFM=FB,LRECL=3200,BLKSIZE=12800)
//SYSUT8 DD UNIT=VIO,DISP=(NEW,DELETE),
//        SPACE=(32000,(30,30)),DCB=(RECFM=FB,LRECL=3200,BLKSIZE=12800)
//SYSUT9 DD UNIT=VIO,DISP=(NEW,DELETE),
//        SPACE=(32000,(30,30)),DCB=(RECFM=VB,LRECL=137,BLKSIZE=882)
//SYSUT10 DD SYSOUT=*
//

```

Figure 53. A JCL fragment converted to a build script

Example of a build script for a COBOL compile

TeamConnection provides a sample build script program for compiling MVS COBOL programs. This sample is called `fhbcobm.jcl`. It invokes a JCL procedure called `IGYWC`, which needs to be in the system `PROCLIB` concatenation or in the data set identified by the `TEAMPROC` DD statement in the MVS build job. You may need to adjust the default parameters for the system. The following JCL should work with any IBM COBOL/II type of compiler such as the IBM COBOL/II compiler `IGYCRCTL`:

```

/*-----
/* PROGRAM:  cobolcmp.jcl
/* IBM COBOL for MVS
/* Compile Only
/*
/*-----
//COBOLCMP EXEC PGM=IGYCRCTL,PARM='&TCPARM'
/*
/* INPUT FILES WITH EXTENSION CBL
/*
//SYSIN DD TCEXT=CBL,DISP=(NEW,DELETE),
// SPACE=(32000,(30,10)),UNIT=SYSDA,
// DCB=(RECFM=FB,LRECL=80,BLKSIZE=6160)
/*
/* COPY FILES WITH EXTENSION CPY
/*
//SYSLIB DD TCEXT=CPY,DISP=(NEW,KEEP),
// SPACE=(32000,(30,30,30)),UNIT=SYSDA,
// DCB=(RECFM=FB,LRECL=80,BLKSIZE=6160)
/*
//SYSPRINT DD TCEXT=TCOUT,DISP=(NEW,DELETE),
// SPACE=(32000,(30,30)),UNIT=SYSDA,
// DCB=(RECFM=FBA,LRECL=133,BLKSIZE=3990)
//SYSLIN DD TCEXT=OBJ,UNIT=SYSDA,
// DISP=(NEW,DELETE),SPACE=(32000,(30,10)),
// DCB=(RECFM=FB,LRECL=80,BLKSIZE=3200)
/*
//SYSUT1 DD UNIT=SYSDA,SPACE=(CYL,(1,1))
//SYSUT2 DD UNIT=SYSDA,SPACE=(CYL,(1,1))
//SYSUT3 DD UNIT=SYSDA,SPACE=(CYL,(1,1))
//SYSUT4 DD UNIT=SYSDA,SPACE=(CYL,(1,1))
//SYSUT5 DD UNIT=SYSDA,SPACE=(CYL,(1,1))
//SYSUT6 DD UNIT=SYSDA,SPACE=(CYL,(1,1))
//SYSUT7 DD UNIT=SYSDA,SPACE=(CYL,(1,1))
//

```

Example of a build script for a link

Because MVS load modules are not easily transferable, TeamConnection provides a sample build script program that reads linkage editor SYSLIN control statements. This script produces a single file that can be returned from MVS and loaded into TeamConnection. You can later extract the file and transport it to MVS, where it can be link edited to produce an executable load module.

The next example shows this sample build script, named fhbtclnk.jcl, which is shipped with the TeamConnection client.

You can use either of the following for an INCLUDE control statement for the FHBTCLNK program:

- INCLUDE DDNAME(MEMBER)
- INCLUDE DDNAME

This syntax is a subset of the linkage editor INCLUDE card.

If the card is an INCLUDE ddname(MEMBER) control statement, the object code is copied into a sequential data set associated with the SYSMOD ddname. Otherwise, the control card is embedded in the data set associated with the SYSMOD ddname. This data set can be returned as the output from this build script.

```
//FHBTCCLNK EXEC PGM=FHBTCCLNK,  
// PARM='SIZE=(768K,192K),LIST,MAP,AMODE(31),RMODE(24),LET,XREF'  
//STEPLIB DD DSN=userid.teamc.LOADLIB,DISP=SHR  
//SYSMOD DD TCEXT=LOAD,DISP=(NEW,DELETE),  
// SPACE=(32000,(30,10)),UNIT=VIO,  
// DCB=(RECFM=U,LRE10CL=80,BLKSIZE=3200)  
//OBJ DD TCEXT=(OBJ,PRE),DISP=(NEW,DELETE),  
// UNIT=VIO,SPACE=(32000,(30,10,10)),  
// DCB=(RECFM=FB,LRECL=80,BLKSIZE=3200)  
//SYSPRINT DD TCEXT=TCOUT,DISP=(NEW,DELETE),  
// UNIT=VIO,SPACE=(TRK,(30,10)),  
// DCB=(RECFM=FB,LRECL=121,BLKSIZE=1210)  
//SYSLIN DD *  
INCLUDE OBJ(&TCINPUT)  
ENTRY CEESTART  
//
```

TCEXT attributes have been added to the following DD statements:

Data set	Purpose
SYSMOD	Return the output to check in to TeamConnection
OBJ	Receive the object files transported to MVS from TeamConnection
SYSPRINT	Return any FHBTCCLNK messages to TeamConnection

In the SYSLIN data stream, the statement INCLUDE OBJ(&TCINPUT) will be duplicated for all of the inputs to this build. The &TCINPUT variable will be replaced with the base name of the input without the extension.

To use the output of this build script as an MVS executable, do the following:

1. Extract the output from TeamConnection.
2. Transfer the output as a binary file from your workstation to MVS (for example, using FTP).
3. Link edit this output into a load module. Possible SYSLIN control statements for the link step include the following:

```
//SYSLIN DD *  
INCLUDE OBJECT(OUTPUT)  
NAME module(R)  
//
```

The output specified in `INCLUDE OBJECT(OUTPUT)` contains embedded control statements specified from the build script `FHBTCLNK`. The linkage editor recognizes these embedded statements and produces an executable load module from the output file. The `NAME` control statement cannot be embedded in the output data set.

Chapter 14. Working with parsers

This chapter describes how to create a parser. It assumes that you have read “Chapter 10. Basic build concepts” on page 129.

Consider the task of defining and maintaining a build tree. One of the more time-consuming, and error-prone, portions of this task is defining the dependencies that one TeamConnection part has on others.

For example, if `hello.c` includes `hello.h`, you need to define `hello.h` as a dependency of `hello.c` in the build tree. That sounds simple enough, but imagine a real application in which there are hundreds of dependencies and the dependencies have dependencies. Defining such a tree becomes very difficult; maintaining it, even more so.

To solve this problem and automate some of the work of defining and maintaining a build tree, you can instead use a parser object. The task of a parser is to inspect source code to determine dependencies. TeamConnection verifies all parser dependencies when the user creates or checks in the part and again during build. TeamConnection will add all parser dependencies that it can find and, for build, update them as needed. In the previous example, a parser can inspect `hello.c`, recognize that it has a dependency on `hello.h`, and create that dependency in the TeamConnection build tree.

Because parsers are language-dependent, you probably need a different parser for each language you use in a particular release. For example, you might have both a COBOL parser and a C parser in a release. Many parts in the release can use the same parser.

Usually a TeamConnection administrator defines parsers, but anyone with the proper authority can do so. For more information about the required authority, see Appendix I. Authority and notification for TeamConnection actions.

Creating a parser

As with most other TeamConnection operations, there are two ways you can create a parser: using the graphical user interface (GUI) or the command line interface.

To create a parser using the GUI:

1. From the Actions pull-down menu on the Tasks window, select **Parsers** → **Create**.
2. On the Create Parser window, type the requested information.

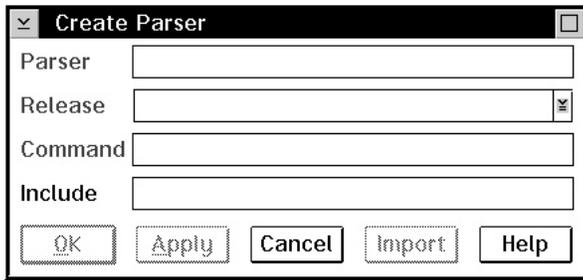


Figure 54. Create Parser window

From a command line, type the parser `-create` command and press Enter. The complete command syntax looks like the following:

```
teamc parser -create name -command name -release name -family name
             [-include paths]
             [-become user_name] [-verbose]
```

No matter which way you create a parser, you must specify a number of attributes for it. Together with the contents of the parser command file, the following attributes define how a parser determines the dependencies for a TeamConnection part.

Parser The name of the parser must be unique within a release. It can be anything you want, but for best results, establish and follow a meaningful naming convention. An example of a parser name is `c_parser`.

Release

This is the name of the release that contains the parser. Parsers are release-specific objects. They are not versioned within a release; therefore you can have only one version of a parser at any time in a release.

To use the parser from a previous release, you can link to a part that uses it in that release. This action copies the parser to the new release. Otherwise, you must create the parser again in the new release.

Command

This is the name of the command file that the parser invokes to determine the dependencies. It can be any file name that exists in the execution path of the family server at the time a build is performed. The parser command is run as a subprocess on the machine where the family server is located.

The task of the command file is to inspect the source file and return a list of dependencies. The syntax for invoking this command is discussed in “Writing a parser command file” on page 177.

Include

This is a concatenated set of paths that define where the parser looks for parts when processing the set of dependencies returned from the command file. These dependencies come in two types:

- A dependency in which the file is stored in the TeamConnection database. For example, `hello.c` includes `hello.h`, and both files are stored in the TeamConnection database. During a build, these dependencies must be extracted to a path accessible by the build server. Because a build extracts parts from TeamConnection, anyone requesting a build needs to have PartExtract authority to all parts involved in the build.
- A dependency on a file that is not stored in the TeamConnection database. An example of such a dependency is `stdio.h`, which is typically stored in a compiler's include path and not in the TeamConnection database.

Each path named in **Include** is queried in the TeamConnection database to see if it contains a part matching the dependency name. For example, suppose you define a parser named `c_parser` with an include path as follows:

```
src\include;src\package;.;src\comm\include;
```

One of the parts to which this parser is attached, `src\example.cpp`, contains the statement `#include "example.hpp"`. Thus the command file for `c_parser` reports `example.hpp` as a dependency of `src\example.cpp`. The parser concatenates each path listed in `c_parser`'s include path with the name `example.hpp`, then inspects the contents of the TeamConnection database to see if a part with that name exists. So the TeamConnection database is queried first to find `src\include\example.hpp`, then `src\package\example.hpp`.

The period (`.`) in the include path tells TeamConnection to concatenate the path of the part to which the file is a dependent with the dependent's file name. In this example, that means the TeamConnection database is queried to find a part named `src\example.hpp`.

Writing a parser command file

A parser command file accepts two parameters as input:

- *source file*—the name of the file that contains the source to be parsed.
- *dependency list file*—the name of a file into which the names of the dependent files should be written, one per line. For example, the contents of the file might look like this:


```
hello.h
stdio.h
```
- *family*—the family that contains the source to be passed.
- *release*—the release that contains the source to be passed.
- *workarea*—the workarea that contains the source to be passed.

Both the source file and the dependency list file are created by the TeamConnection family server. They are erased after the parse is complete.

To write a command file, write a program, in any language, that does the following:

1. Reads the source file

2. Determines which other files are used by it
3. Writes out the list of such files into the dependency list file

For example, for a C source file, the program could report a list of all the files included by the source file (using `#include` statements). For a COBOL program, `COPY` statements would be the cue. TeamConnection ships a sample of a command file named `fhbopars.cmd`. It is written in REXX.

Putting a parser to work

For an application to use a parser, the parser must be attached to the TeamConnection parts that it will check for dependencies. Unlike a builder, a parser is attached to the *input* part rather than the output.

To attach a parser to a part, do one of the following:

- From the GUI, select **Parts** → **Modify** → **Properties** from the Actions menu of the TeamConnection Tasks window. On the Modify Part Properties window, type the name of the parser.

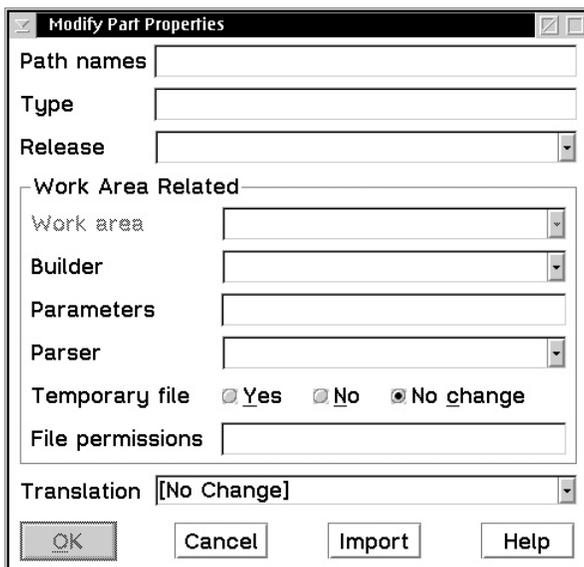


Figure 55. Modify Part Properties window

- At a command prompt, type the following and press Enter:

```
teamc part -modify part -parser name -release name
      -family name
```

The complete syntax for this command is described in the *Commands Reference*.

You can also attach a parser to a part when the part is created.

After you attach a parser to a part, it is ready for action. The next time the part is checked in, when a part is created, or when the parser is attached, the parser will invoke its command file, which will report back a list of dependencies.

Using a parser does not keep you from defining dependencies manually by using the GUI or the part -connect command. If you explicitly define a dependency in this way, the dependency is not deleted unless you delete it, regardless of whether the parser would recognize it as such.

For more information about attaching parsers to the build tree, refer to “Creating the build tree for the application” on page 184.

Removing a parser from a part

If you no longer want to use a parser to determine dependencies for a part, do one of the following:

- From the GUI, select **Parts** → **Modify** → **Properties** from the Actions menu of the TeamConnection Tasks window. On the Modify Part Properties window, type null in the **Parser** field.



Figure 56. Modify Part Properties window

- From a command line, type the following:

```
teamc part -modify name -parser null  
-release name -family name
```

Chapter 15. Building an application: an example

This chapter uses an extended example to describe in more detail how each of the components of the build function work together. All commands used in this example are described in detail elsewhere in this publication. This example walks through the control flow for a sample application, explaining what happens at each step.

These are the tasks involved in building our sample application, msgcat.exe:

Task	Page
Starting build servers	182
Creating builders and parsers	183
Creating the application build tree	184
Starting the build	189
Monitoring the build	192
Building in spite of errors	193
Forcing a build of all parts	193
Finding out which parts will be built	194
Canceling a build	194

We will use a simple example build tree that looks like the following:

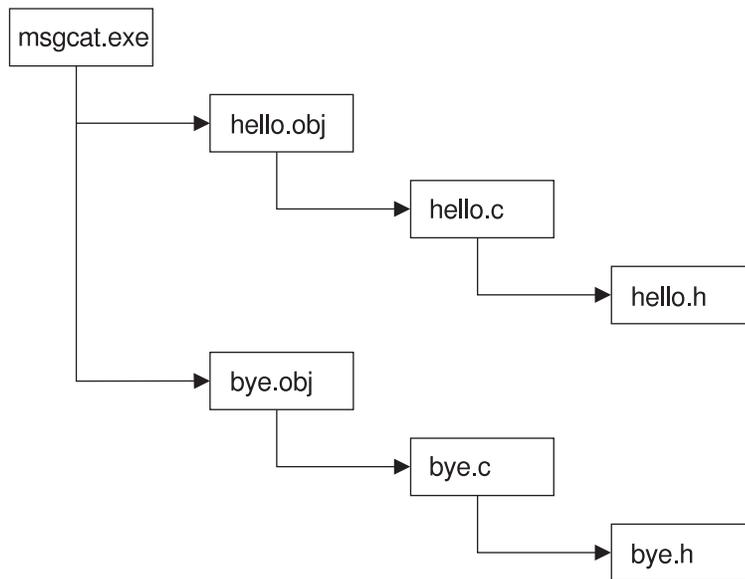


Figure 57. Sample build tree

For more examples of build trees, see “More sample build trees” on page 195.

In terms of the build object model, the objects that make up this tree look like this:

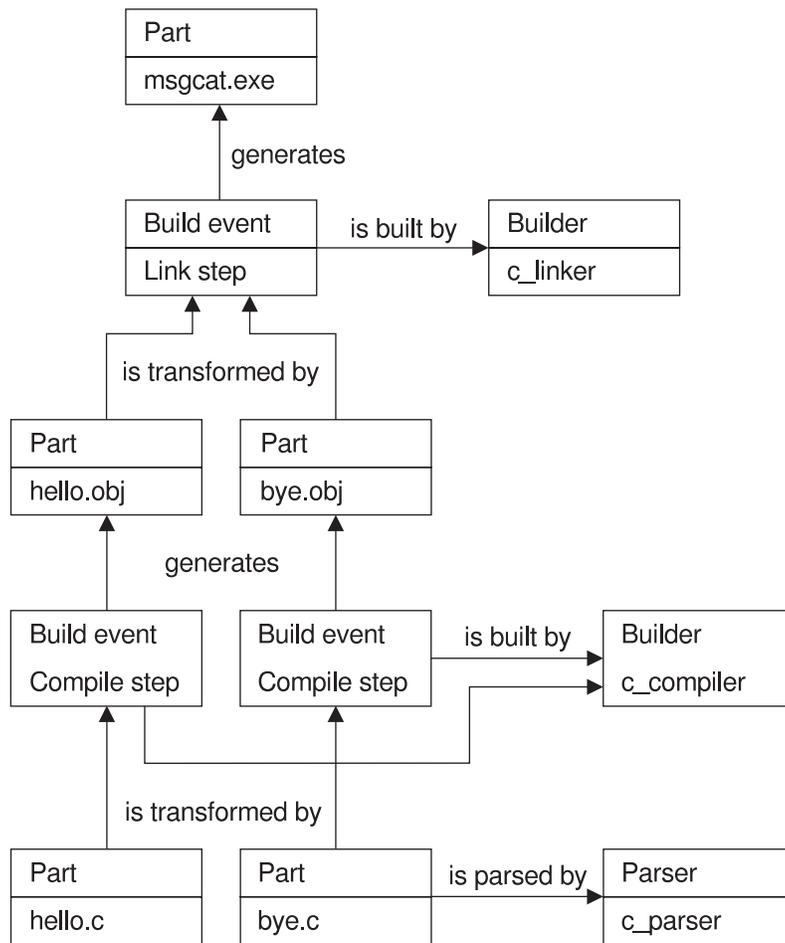


Figure 58. Sample build object model for msgcat.exe

Starting the build servers

The software development team in our example is building large applications using a family named testfam, so they set TC_FAMILY to testfam. They plan to spread the work across several build servers, taking advantage of TeamConnection’s ability to perform multiple build events simultaneously.

Mark, the build administrator, has installed a number of build servers on the team's machines, for building OS/2 and MVS applications. As he starts them (in pairs), he groups them into *pools*, according to the work he expects to use them for.

Mark plans for the following pools:

mvspool

For MVS builds

pool1 For normal OS/2 builds

pool2 For fast, high-priority OS/2 builds

Each pool is formed as Mark starts build servers and assigns them to it. He starts the following build server (bldserv2):

```
teamcbld -b bldserv2 -p pool1 -e os2
```

The parameters specify the following:

- p** The build server is assigned to the pool named pool1.
- e** The environment is os2.

Use the teamcbld command to start the build server when the family server has already been started. To start the family server along with the build server, you can use the teamcd command.

Creating builders and parsers

For the parts of the application that are written in C language, Mark creates the following:

- A builder named `c_compiler`, to do the compiles
- A builder named `c_linker`, to do the links
- A parser named `c_parser`, to check for dependencies

For both builders Mark specifies `os2` as the **Environment**, the same as that of the build server (bldserv2) started earlier. Build events that use these builders (`c_compiler` and `c_linker`) can take place on this build server.

After he creates the builders and parsers for the applications, Mark spreads the following information to the programmers who will be using them:

- The names of the build pools
- The names and purposes of the builders and parsers

Creating the build tree for the application

At this point, Greg begins defining the build tree for his portion of the application, as shown in Figure 57 on page 181. He has already created the files `hello.c`, `hello.h`, `bye.c`, and `bye.h` in the `TeamConnectiondatabase`. Now he does the following:

1. Creates a place-holder part for the output of the link step. This file, `msgcat.exe`, is the target for the entire build, the output of linking `hello.obj` and `bye.obj` using the builder `c_linker`, and the parent of `hello.obj` and `bye.obj`. Because the file has no contents initially, he selects **No source** (or specifies `-empty` on the command line), to identify it as a place holder.

Using the GUI, he can create this file by selecting **Create** from the **Actions** → **Parts** menu of the Tasks window, and completing the fields as shown in the following illustration:

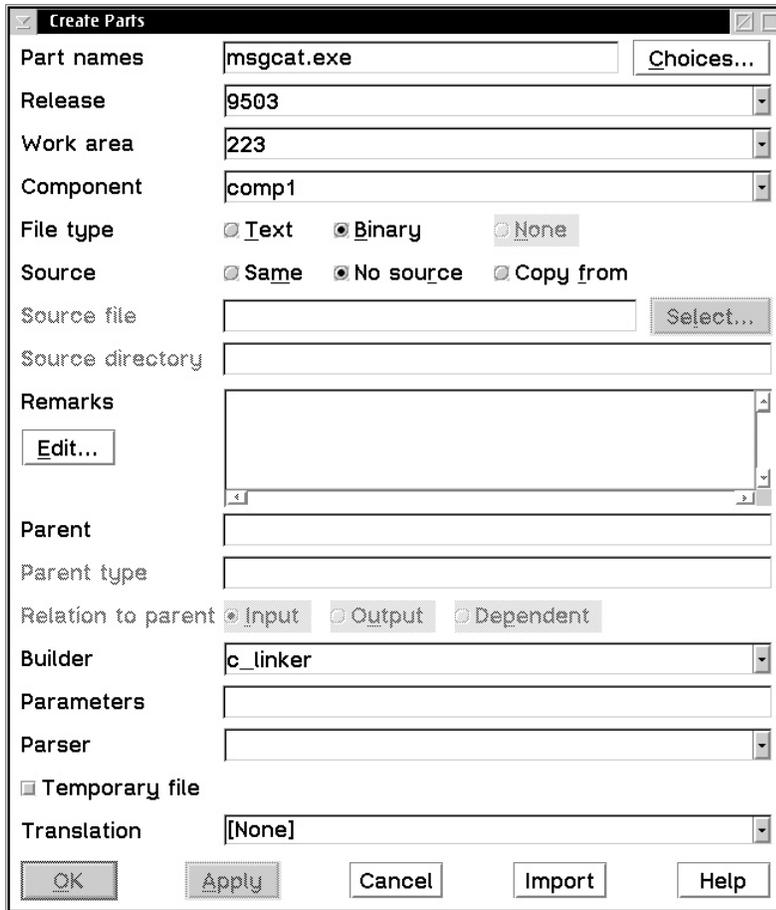


Figure 59. Create Parts window

Using the command-line interface, he can create the part by issuing the following command:

```
teamc part -create msgcat.exe -builder c_linker -binary -empty
-release 9503 -workarea 223 -component comp1
```

- Creates two place-holder parts for the output of compiling the .c files. These parts are the output of the compile step; c_compiler, the builder that manages that step, is attached to both of them. He indicates that they are input to their parent file, msgcat.exe.

Using the GUI, he can create these files by selecting **Create** from the **Actions** → **Parts** menu of the Tasks window, and completing the fields as shown in the following illustration:

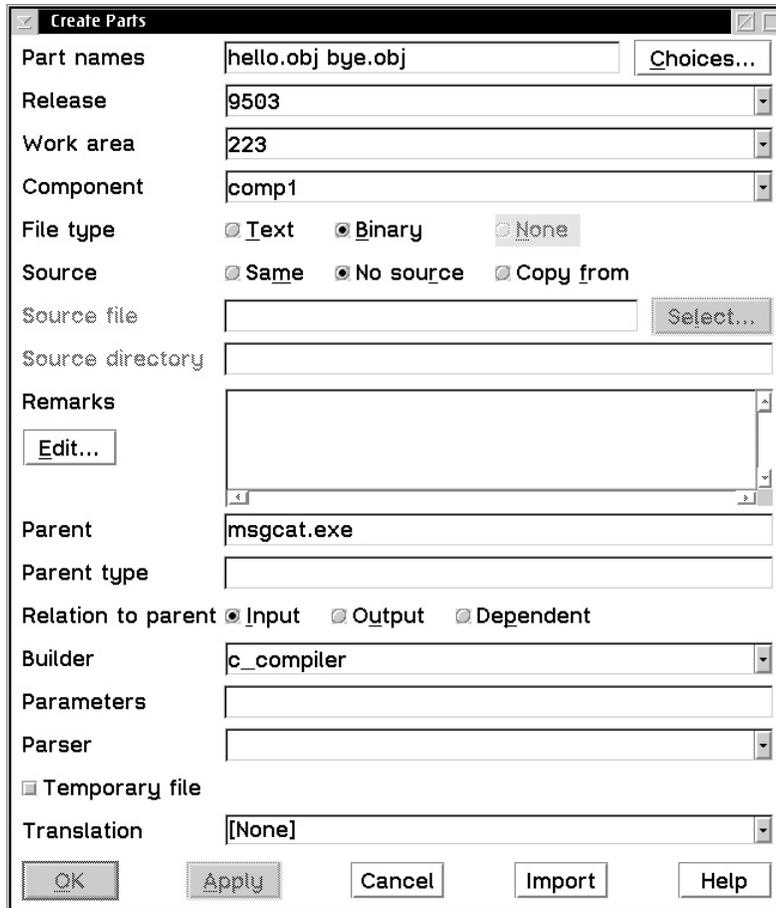


Figure 60. Create Parts window

Using the command-line interface, he can create the parts by issuing the following command:

```
teamc part -create hello.obj bye.obj -builder c_compiler -binary -empty
-release 9503 -workarea 223 -component comp1 -parent msgcat.exe -input
teamc part -create hello.h bye.h -release 9503 -workarea 223
```

3. Attaches the parser `c_parser` to the `.c` files.

Using the GUI, he can attach the parser to these files by selecting **Modify** → **Properties** from the **Actions** → **Parts** menu of the Tasks window, and completing the fields as shown in the following illustration:



Figure 61. Modify Part Properties window

Using the command-line interface, he can attach the parser to these parts by issuing the following command:

```
teamc part -modify hello.c bye.c -parser c_parser -release 9503
-workarea 223
```

Remember, the parser is attached to an *input* file. The part's contents will be parsed and dependency information will be stored.

4. Connects the .c files into the build tree.

Using the GUI, he can connect these files by selecting **Connect** from the **Actions** → **Parts** menu of the Tasks window, and completing the fields as shown in the following illustration. He needs to execute this function twice: once to connect hello.c to hello.obj and once to connect bye.c to bye.obj.

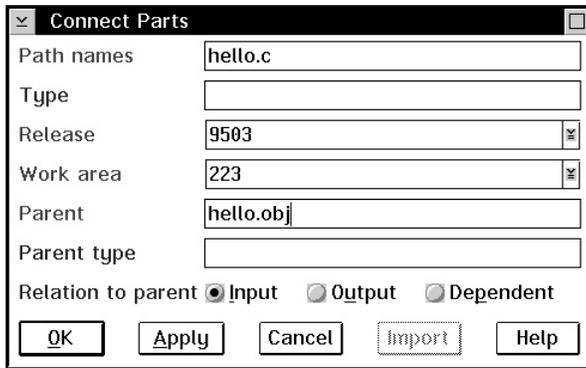


Figure 62. Connect Parts window

Using the command-line interface, he can connect these parts by issuing the following commands:

```
teamc part -connect hello.c -parent hello.obj -input -release 9503
        -workarea 223
teamc part -connect bye.c -parent bye.obj -input -release 9503
        -workarea 223
```

The .h parts are not connected because he expects the parser on hello.c and bye.c to find the correct dependencies.

5. Now, Greg can see the build tree in the GUI. From the **Objects** pull-down menu on the Tasks window, he selects **Parts** → **View build tree**. The BuildView Filter window is displayed; from here he can bring up the build tree.

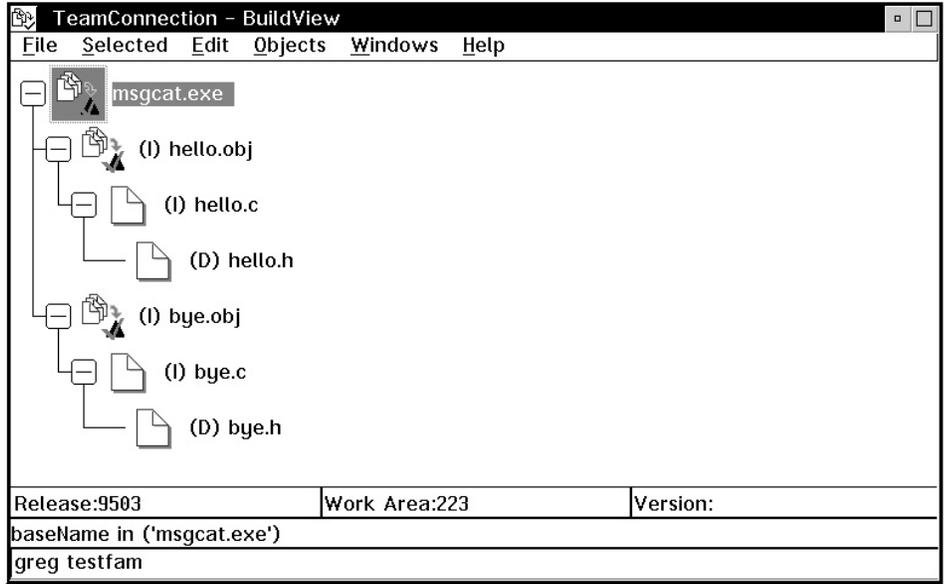


Figure 63. The build tree display

Starting the build on the client

After much hard work on his source code, Greg is ready to start building his application.

Using the GUI, he can start the build by selecting **build** from the **Actions** → **Parts** menu of the Tasks window, and completing the fields as shown in the following illustration:

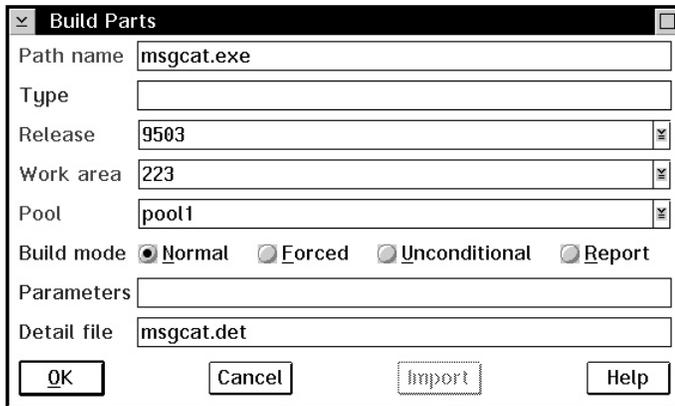


Figure 64. Build Parts window

Using the command-line interface, he can start the build by issuing the following command:

```
teamc part -build msgcat.exe -release 9503 -workarea 223
-pool pool1 -normal -detail msgcat.det
```

This command specifies the following:

Build target

The name of the part at the top of the build tree, msgcat.exe, which is the final output of this build. TeamConnection uses the build target to determine the scope of the build.

Work area

The version of the TeamConnection parts and build tree to be used when performing this build. This version is completely specified by naming the family, release, and work area: in this case, -release 9503 -workarea 223. The output of the build is placed in this work area.

Build pool

The set of build servers that should be used to process the build request, as defined when the build servers are started. The pool pool1 includes the build server started in “Starting the build servers” on page 182.

Build mode

How the build takes place. Possible values for this build option include the following:

Normal

Builds only the parts that are out-of-date. Processing stops after the first error is returned.

Force Builds all parts, even if they are not out-of-date. Processing stops after the first error is returned.

Unconditional

Builds only parts that are out-of-date but continues processing even if errors are returned. Note that outputs are not rebuilt for inputs that have failed.

Report Gives a preview of what would be built if you invoked a build. The report identifies what steps would occur without any translations taking place.

In our example, Greg specifies `-normal`, which is the default. In this mode, only the parts that are stale with respect to their inputs are rebuilt. In other words, only the minimum amount of work to bring everything up-to-date is performed. “Running a build in spite of errors” on page 193 and following sections provide examples of using the other build modes.

In normal mode, the build is halted if an error is found. Any remaining build events in the build scope are canceled, but any build events already performed are not undone.

Detail file name

The name of an output file in which TeamConnection stores the collected stdout and stderr of the build scripts.

When Greg starts the build, the information from the command is passed to the testfam family server over TCP/IP. At this point, Greg’s TeamConnection client waits to receive confirmation from the family server that the build request was received and is being processed. Included in this processing; the family server does the second phase of parsing, where the validation of dependency information is done, and then it determines what needs to be rebuilt and adds this to the job queue.

Note: Only one build is allowed in a work area at one time (though the build events that make up the build might be distributed to different build agents on a number of machines). So if Greg is sharing work area 223 with Barbara, she cannot issue a `part -build` command in that same work area until Greg’s build is complete.

TeamConnection handles the next parts of the build process automatically.

Putting the build scripts to work

At this point, the build server looks at the description of the event it has been asked to perform, then checks its cache for each part and the build script it needs. If it does not find the parts there, or if the cached parts are out of date, it invokes the build script, passing to it the names of the input and output parts and the parameters specified on the builder. The parts created by the build script and the return code generated by it are sent back to the build server, which then updates the contents of the TeamConnection database.

In our example, each of the build servers receives a compile event to perform. Each extracts the .c source files it needs from the TeamConnection database and the contents of the build script for the c_compiler builder. The build servers then run their build scripts.

The results (the .obj files and the return code) are sent back to the build servers. After updating the TeamConnection database, the build servers re-enter the polling loop to see if any more build events await their attention.

Because the compile steps are performed in parallel, Greg can build this application a little more quickly than if they had happened in serial mode. In this simple example, the difference is hardly noticeable; but in a large build of hundreds of parts, with multiple build servers available on a local area network, the performance improvement can be enormous.

Finishing the job and reporting the results to the user

The processing described by the previous two steps is repeated until there are no more build events remaining. The results of the build are displayed in the Build Progress window or in stdout. At this point the build is complete.

To complete our example, the previous two steps are repeated to complete the link step, using either of the build servers in pool1. Greg now can extract the resulting executable from TeamConnection, using the part -extract msgcat.exe command, and run it.

Monitoring the progress of a build

During the course of a build, you can monitor its progress in several ways:

- If the build was started from the command line, by issuing the report -view partview command against the work area in which you are building. From this report, you can determine the states of the parts. Use the part -viewmsg command to see the build messages issued because of a failed build. For complete syntax of these commands, refer to the *Commands Reference*
- If the build was started from the GUI, in the Build Progress window. You can find the same information by looking at stdout.

Greg can see how the build is progressing by checking the Build Progress window. For example, he might see these messages:

```
6021-700 Number of distinct build events for this build: 3.
Build of 'hello.obj' started at '15:33:47 1995-08-10'
via a build agent on the host 'OCTOFVT'.
Build of 'hello.obj' successfully completed at '15:34:45 1995-08-10'.
Completed Jobs: 1
Remaining Jobs: 2
Build of 'bye.obj' started at '15:34:49 1995-08-10'
```

```
via a build agent on the host 'OCTOFVT'.
Build of 'bye.obj' successfully completed at '15:35:22 1995-08-10'.
Completed Jobs: 2
Remaining Jobs: 1
Build of 'msgcat.exe' started at '15:35:26 1995-08-10'
via a build agent on the host 'OCTOFVT'.
Build of 'msgcat.exe' successfully completed at '15:35:56 1995-08-10'.
Completed Jobs: 3
Remaining Jobs: 0
Processing Completed for 'msgcat.exe'.
```

To see the commands that TeamConnection issued during the build, you can look at the detail file specified in the part `-build` command.

Running a build in spite of errors

If you find that a build is stopping because of errors, you can check the build detail file or the Build Progress window for the cause. If the error is minor, you might decide to run the build despite the errors — for example, when you are debugging. To do this, specify that you want the build to complete unconditionally.

In our example, when Greg builds `msgcat.exe` for the first time, he wants to find and correct any errors that occur during the build, so he uses the following command:

```
teamc part -build msgcat.exe -release 9503 -workarea 223 -unconditional
```

As in normal mode, only the parts that are stale with respect to their inputs are rebuilt; only the minimum amount of work to bring everything up-to-date is performed.

However, even if an error is found, the build continues if possible. As with normal mode, if the build is halted, any build events remaining are canceled. Any build events already performed are not undone.

Building all parts, regardless of build times

To make sure that **all** parts in the build tree get built, whether or not they are stale, you specify the `-force` parameter on the part `-build` command.

In this mode, all parts that are descendants of the build target are rebuilt.

In our example, Greg can force TeamConnection to build all parts in the `msgcat.exe` build tree using the following command:

```
teamc part -build msgcat.exe -release 9503 -workarea 223 -force -pool pool1
```

If an error occurs, the build is halted. Any remaining build events are canceled, but any build events already performed are not undone.

Finding out which parts will be built

Before running a build of a large application, you might want to find out exactly which parts will be built. If you specify that you want to run in report mode, TeamConnection determines what will be built in a normal build and produces a report showing the results.

If Greg really wants to see which parts of msgcat.exe will be built before he runs the actual build, he can issue the following command:

```
teamc part -build msgcat.exe -release 9503 -workarea d410 -report -pool pool1
```

He sees the following report:

```
6021-700 Number of distinct build events for this build: 3.
6021-407 The builder c_compiler will be invoked.
6021-406 The builder parameters consist of:
        command:  compC.cmd
        input:    hello.c
        output:   hello.obj
        dependent: hello.h
6021-407 The builder c_compiler will be invoked.
6021-406 The builder parameters consist of:
        command:  compC.cmd
        input:    bye.c
        output:   bye.obj
        dependent: bye.h
6021-407 The builder c_linker will be invoked.
6021-406 The builder parameters consist of:
        command:  linkC.cmd
        input:    hello.obj bye.obj
        output:   msgcat.exe
        dependent:
```

The report shows that bye.obj and msgcat.exe must be rebuilt.

Canceling a build

To cancel a build that is in progress, do one of the following:

- If the build was started from the GUI, on the Build Progress window select the **Cancel Build** push button.
- If the build was started from the command line, type the following command and press Enter:

```
teamc part -build name -cancel
```

Where *name* is the part that you are building. Be sure to specify the same part name that you specified when starting the build, rather than a part that is lower in the build tree.

This command will stop all queued and in_progress builds. Any build events already performed for that build are not undone.

For example, if Greg cancels the build of msgcat.exe when the compile steps have been completed, then the link step is not performed. However, the newly compiled hello.obj and bye.obj are left in the database, with their build times updated.

More sample build trees

The msgcat.exe example is just one possible build tree. Here are some others.

Defining multiple outputs from a single build event

Figure 65 shows part of the build tree for robot.dll:

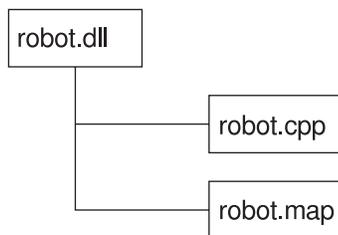


Figure 65. The build tree for robot.dll

Because the build tree shows the relationships between parts hierarchically, robot.map is a child of robot.dll, even though it is actually built from the same input part, robot.cpp. But robot.map is defined as an *output* of robot.dll. Here are the commands to set up this relationship.

First come the commands to create the parts:

```
teamc part -create robot.dll -builder dll_builder -binary -empty  
-release 9503 -component robot  
teamc part -create robot.cpp -release 9503 -component robot
```

```
teamc part -create robot.map -builder dll_builder -binary -empty  
-release 9503 -component robot
```

Next are the commands to connect the parts into the build tree:

```
teamc part -connect robot.cpp -parent robot.dll -input -release 9503  
teamc part -connect robot.map -parent robot.dll -output -release 9503
```

You might use this command to start the build:

```
teamc part -build robot.dll -workarea 915 -release 9503
```

The output of this build would be both robot.dll and robot.map. Any parameters specified in the teamc part `-build robot.dll` command would also apply to the build of robot.map.

Synchronizing the build of unrelated parts

An entire application can require multiple separate builds. For example, in the robot application, there might be one build to create the .dll parts, another to create the .exe parts, and so on. To ensure that the entire application gets built together, you can create a part that acts as a collector, with the .dll and .exe parts as input to it.

For example, Tim creates this build tree for the robot application:

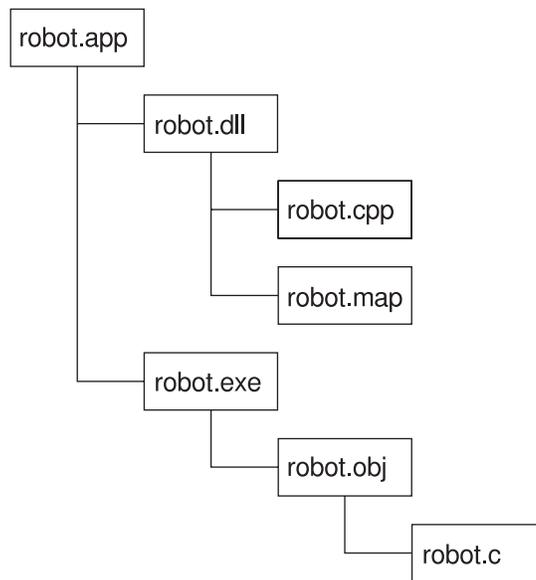


Figure 66. The build tree for robot.app

Assuming he already has the build trees for robot.dll and robot.exe set up, here is how he sets up the collector part:

1. He creates a null builder with no contents:

```
teamc builder -create nullBuilder -script null -none -environment os2  
-condition == -value 0
```
2. He creates the collector part:

```
teamc part -create robot.app -builder nullBuilder -none -release 9503  
-component robot
```

The `-none` flag identifies this as a part that will never have any contents.

3. Tim connects the other parts to the collector:

```
teamc part -connect robot.dll robot.exe -parent robot.app -input  
-release 9503
```

When Tim builds robot.app, the result is a build of both robot.dll and robot.exe.

Part 5. Using TeamConnection to package products

Chapter 16. Using TeamConnection to package a product	201
Setting up your build tree for packaging	201
Setting up a build tree for the gather tool.	202
Chapter 17. Using the Gather tool	205
Using the teamcpak command for the Gather tool.	205
Command line flags.	206
Examples of the teamcpak gather command	208
Writing a package file for the Gather tool.	208
Syntax rules for a Gather package file.	208
Keywords for a Gather package file	209
Using exit keywords in the DATA clause	212
Using exit keywords in the RULE clause	212
Using exit keywords: an example	213
Chapter 18. Using the Tivoli Software Distribution packaging tool	215
Using the teamcpak command with Tivoli Software Distribution	215
Command line flags.	216
Example of the teamcpak softdist command.	217
Writing a package file for Tivoli Software Distribution	217
Syntax rules for a Tivoli Software Distribution package file	217
Keywords for a Tivoli Software Distribution package file	218
Problem determination for the Tivoli Software Distribution tool	220
Sample package file	221

This section describes how to use the TeamConnection packaging function, which helps you automate the packaging and distribution of your applications. This section is written for the person in your organization who is responsible for software distribution.

Chapter 16. Using TeamConnection to package a product

After you have built an application to your satisfaction, it is time to distribute it to users. This chapter describes how you can use TeamConnection to help automate the packaging and distribution steps.

TeamConnection provides the following:

- Two electronic software distribution tools:
 - **Gather**, which moves an application's parts into a single directory in preparation for distribution.
 - **Tivoli Software Distribution**, a bridge tool that automates the installation and distribution of software or data using Tivoli software distribution tools as the distribution vehicle.
- Two sample build scripts for connecting the Gather and Tivoli Software Distribution tools with TeamConnection user-defined builders.
 - Gather - **gather.cmd** which specifies the `teamcpak gather` command.
 - Tivoli - **softdist** (on UNIX platforms) or **softdist.exe** (on Windows NT) which specifies the `teamcpak softdist` command.

To use TeamConnection in packaging a product, you might do any of the following tasks:

Task	Page
Setting up your application's build tree for packaging	201
Using the <code>teamcpak gather</code> command	205
Writing a package file for the gather tool	208
Using the <code>teamcpak softdist</code> command	"Using the <code>teamcpak</code> command with Tivoli Software Distribution" on page 215
Writing a package file for the Tivoli Software Distribution tool	"Writing a package file for Tivoli Software Distribution" on page 217

Setting up your build tree for packaging

When TeamConnection builds an application, the application's build tree identifies the parts to be built and the tools to use in building it. Similarly, when you use TeamConnection for packaging the application, the build tree can define the parts to be packaged and the tools to do it.

The output of a packaging step might be any of the following:

- The application's parts gathered into a new directory structure
- The distribution of the application using NVBridge
- The distribution of the application using some other distribution tool

Setting up a build tree for the gather tool

To gather the parts of your application into a single directory for distribution, you create an output part whose builder calls the gather tool, and you make this output part the top level of the build tree.

For example, for the robot control application, `robot.app`, the build tree might look in part like this:

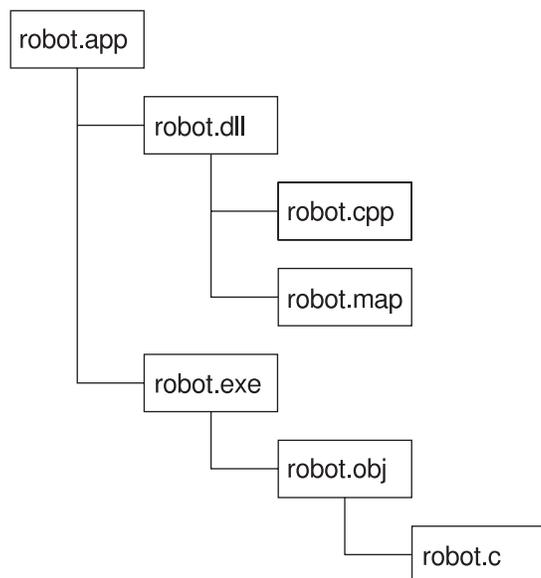


Figure 67. Part of the build tree for `robot.app`

After the application is built, the programming team needs to get it to the test team. They could extract the application, but doing a simple extract would preserve the existing structure, with parts contained in directories according to their application component. A better structure might be to place all of the `.dll` files in one directory, all of the `.exe` files in another, and so on. To move the parts into this structure, the test team does a different kind of build, using the gather tool.

To make this happen, Annmarie does the following:

1. She creates the top-level part for the new build tree. The name of this part is the same as the directory in which the gathered parts are to be placed. In this example, e:\robot is the output file from the gather step. Annmarie uses the following command:

```
teamc part -create e:\robot -none -builder gather1 -family octo
-release 9503 -workarea 410
```

2. She writes a package file that contains instructions for the gather tool and creates this file as a TeamConnection part:

```
teamc part -create robot.pkf -text -parent e:\robot -input -family octo
-release 9503 -workarea 410
```

For more information, see “Writing a package file for the Gather tool” on page 208.

3. She creates a builder, gather1, that calls the gather tool:

```
teamc builder -create gather1 -script gather.cmd
-parameters "-o -x" -release 9503
-environment os2 -condition == -value 0 -family octo
```

gather.cmd is a sample build script that is shipped with TeamConnection. It specifies the teamcpak gather command.

4. She connects robot.exe and robot.dll to e:\robot as inputs:

```
teamc part -connect robot.exe -parent e:\robot -family octo
-release 9503 -workarea 410
```

```
teamc part -connect robot.dll -parent e:\robot -family octo
-release 9503 -workarea 410
```

5. She also connects a readme file for the application:

```
teamc part -connect read.me -parent e:\robot -family octo
-release 9503 -workarea 410
```

As a result of Annmarie’s work, the build tree for e:\robot looks like this:

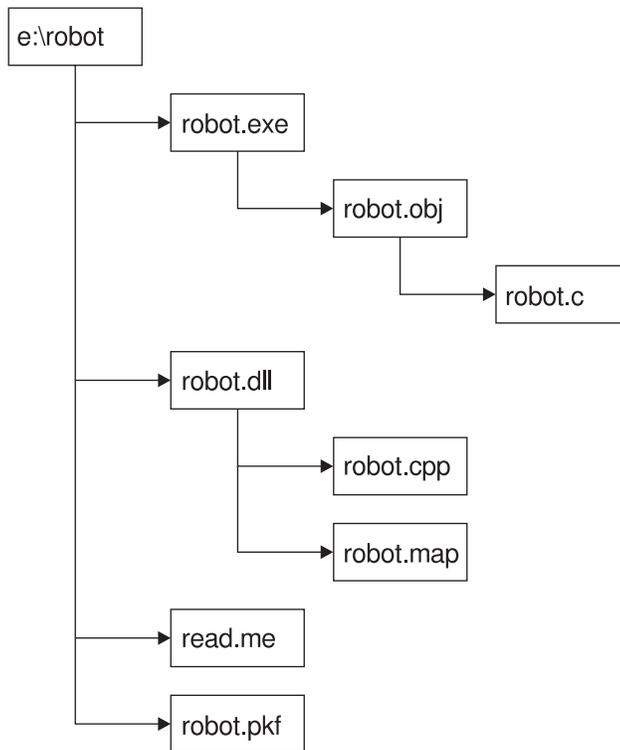


Figure 68. Adding the gather step to the build tree

The package file, robot.pkf, specifies the directories into which the robot files are gathered, with e:\robot as the target root directory. When Annmarie builds e:\robot, the .dll files are placed in e:\robot\dll; the .bin files are placed in e:\robot\bin. Instead of extracting the built application from TeamConnection, the test team can pull the application from e:\robot.

If Annmarie wants to gather the same files into a different target directory, all she needs to do is write a different package file and connect the parts to a different parent.

Chapter 17. Using the Gather tool

The Gather tool automates the movement of software and data from one directory to another on the same machine to prepare a package for electronic distribution. It can copy or erase files; it can create or delete directories.

This tool takes a list of input files and moves them into a directory structure as directed by a package specification file. You specify the target root directory path in this file, along with a collection of rules that instruct which files to copy to which directories. How these files and directories are actually handled is controlled via option flags.

By writing different package specification files, you can take the same input files and transfer them into different target directory structures.

Take the robot application as an example. We previously showed one possible directory structure, with each subdirectory containing files with the same extension:

```
e:\robot
  \dll
    hand.dll
    optics.dll

  \exe
    hand.exe
    optics.exe
```

By writing a different package file, you might put both .dll and .exe files in the same target directory:

```
f:\robot
  \bin
    hand.dll
    optics.dll
    hand.exe
    optics.exe
```

You can build both target directories concurrently.

Using the teamcpak command for the Gather tool

To start the Gather tool, use the teamcpak command. This command is found in the directory where the TeamConnection family server is installed. If it is started from a build script, it does not need to be in the execution path of the machine from which the build is started.

The complete command syntax for teamcpak gather looks like the following; you must supply a value for the words that start with a capital letter, such as String. You must specify the command parameters in the order shown.

```
teamcpak [-i] [-o "String"] gather Input_file...
```

Where

-i Specifies that only one *Input_file* is specified in the command: an include file containing the list of input files. This parameter is optional.

If you specify `-i`, it must precede the `gather` flag.

-o "String"

Specifies that the string listed in quotes be passed to the Gather tool. The opening quote must be followed by a blank. For a list of possible flags to be passed, see "Command line flags".

This parameter is optional. If you do not specify `-o`, the default settings for the tool are used.

If you specify `-o`, it must precede the `gather` flag.

gather Specifies the tool to be invoked. If you specify `-i` or `-o`, they must precede this value.

Input_files

Specifies the files to be copied and the name of the package specification file. You can specify this parameter in these ways:

- Specify the name of an include file, whose contents is a list of input files. One of these input files must be a package specification file with the extension `.pkf`. In this case, you must also specify the `-i` parameter.
- Specify a list of two or more files. One of the files must be a package specification file with the extension `.pkf`.
- Specify the directory from which the files are to be copied and the name of the package specification file.

If more than one package file is listed, the first package file on the command line or in the include file is used, and the others are treated as ordinary files.

Command line flags

You can specify the following flags in the `teamcpak` command, using the `-o` parameter. All of these flags are optional. If you do not specify a flag, the `teamcpak` command runs using defaults.

-a Assume that the target tree structure might not exist. If a required directory does not exist, create it and continue processing.

This flag cannot be specified if the `-t` flag is specified.

If neither `-a` nor `-t` is specified, the default is to assume that the desired tree structure already exists. No verification is performed to confirm that the directories exist. If they do not, the condition is detected while the package file rules are being processed. If you stop the `teamcpak` command, some target directories might contain updated files.

- f** Force deletion from the root: if this is used in combination with `-t` or `-c` and `TARGETROOT` is a root directory (for example, `e:\, \, /`).
- t** Ensure that the target tree is exactly the tree specified in the package file. If a directory of the same name exists, the Gather tool does the following:
- Erases the entire contents of the directory and all of its subdirectories
 - Destroys the directory and all subdirectories
 - Performs a `mkdir` command to create the entire tree structure again as specified in the package file
- This flag cannot be specified if the `-a` flag is specified.
- If an `rmdir` command fails during processing, the `teamcpak` command stops.
- If neither `-a` nor `-t` is specified, the default is to assume that the desired tree structure already exists. No verification is performed to confirm that the directories exist. If they do not, the condition is detected while the package file rules are being processed. If you stop the `teamcpak` command, some target directories might contain updated files.
- m** Accept missing source files.
- If this flag is not specified, the default is to ensure that at least one file matches each source specification in the package file. If a match is not found, the Gather tool stops processing.
- d** Accept duplicate files. If a file is found on the target directory that matches the source file specification, it is overwritten by the source file.
- If this flag is not specified, the default is to ensure that no files on the target match the source file specification. For example, if the source specification is `g*.c`, and `greg.c` is found on the target, the Gather tool stops processing.
- c** Clean up the target directories. Erase all files on all target directories that existed before writing source files to these directories. No confirmation messages are issued, and permission errors are ignored.
- If this flag is not specified, the default is to write the source files into the target directories without erasing existing files.
- e** End with delete. This action removes all source files and directories after the Gather tool successfully completes.
- If this flag is not specified, the default is to end without deleting source files and directories.
- x** Abort without recovery. If the program does not end successfully, no attempt is made to restore the file system.
- If this flag is not specified, the Gather tool attempts to restore the file system if the program does not end successfully. To do this, the tool first backs up the file system. The backup directory is the value of the `TMP` environment variable.

Examples of the teamcpak gather command

The following are examples of the teamcpak gather command.

```
teamcpak gather d:\demoapp demoapp.pkf
teamcpak gather a.exe b.exe \help\*.hlp demoapp.pkf
```

In the first example, an input source directory is specified. In the second example, a list of files is specified. In both cases, the files are to be copied into target directories as specified in the demoapp.pkf file.

```
teamcpak -i -o " -t -m -x" gather myfiles.lst
```

The file myfiles.lst contains a list of files to be transformed by the Gather tool, and the name of the package file to be used in the gather. The `-o "-t -m -x"` parameter passes three flags to the Gather tool:

- `-t` specifies that, if the target directories already exist, they be destroyed and recreated.
- `-m` specifies that processing continues even if a source file cannot be found.
- `-x` specifies that, if the program does not end successfully, the file system is left as is, with no attempt to restore it.

Writing a package file for the Gather tool

Use the package file to specify the target directories and the rules for copying files for a gather operation. You can also specify user exit programs to run before, during, or after the gather operation.

A sample package file named gather.pkf is shipped with TeamConnection. You can customize it for your own gather operations.

Syntax rules for a Gather package file

Follow these syntax rules when you write a package file:

- Package files are free format. Text is not positional, and many statements can exist on the same line.
- Comments can appear anywhere within the file. Use the characters `#|` and `|#` as delimiters, as shown in the following example:
`#| This is a comment |#`
- Package file keywords must be prefixed with a left parenthesis and must have a corresponding balanced right parenthesis to end the scope of the keyword.
- If the value for a keyword is a string that contains blanks or parentheses, enclose the string in double quotes.

The following shows the syntax of a package file for the Gather tool. Keywords must appear in the order shown. The first letter of an argument is capitalized; you must supply these values.

```

(DATA
  (PACKAGEFORMAT gather)
  (TARGETROOT Filename)
  (RULE
    (SOURCE Filename...)
    (TARGET Path)
    [(EXITPRIOR String... | EXITREPLACE String... | EXITPOST String...)] )
  .
  .
  [(EXITPRIOR String...)]
  [(EXITPOST String...)]
)

```

Keywords for a Gather package file

DATA This keyword is required. It must be the first keyword in the package file, and it can be specified only once.

All other keywords are nested within the DATA clause.

PACKAGEFORMAT gather

This keyword is required. It can be specified only once. It tells the teamcpak command that this package file is for Gather.

TARGETROOT target_root_path

This keyword is required. It can be specified only once.

Use this keyword to identify the target root directory. Source files are copied to this directory as specified by the RULE statements.

Follow these guidelines when you select your TARGETROOT values:

- Include the drive letter along with the target directory.
- Specify a directory that contains few if any subdirectories that are unrelated to the data you are moving.
- If you specify a drive's root directory (*drive:*), run the teamcpak command using the defaults or only the `-x` or `-x -a` flags.
- Do not set the value of TARGETROOT to *drive:* under the following circumstances:
 - The TARGETROOT drive is the same as the drive from which the teamcpak command is run, and you have recovery set (that is, you have not specified `-o "-x"`).
 - The logical drive for the TARGETROOT has less than 50% free space, and you have recovery set (that is, you have not specified `-o "-x"`).

RULE This keyword is required. You can use one or more RULE keywords within a Gather package file.

Each RULE clause represents a set of Gather operations targeted for one target subdirectory. A RULE clause must contain one SOURCE and one

TARGET keyword. The files in the SOURCE directory are copied to the TARGET path. The target path is derived by concatenating the value of TARGETROOT with a backslash (\), followed by the value of the TARGET keyword specified in the RULE clause.

A RULE clause can also contain one user exit clause: EXITPRIOR, EXITPOST, or EXITREPLACE. For a description of the exit keywords, go to page 212.

The following example copies all *.exe, *.cmd, and *.hlp files to target directory f:\demoapp\bin.

```
(DATA
.
.
(TARGETROOT f:\demoapp )
.
.
(RULE
(SOURCE *.exe *.cmd *.hlp)
(TARGET bin )
)
.
.
)
```

SOURCE <list of file specifications>

This keyword is required once for each RULE clause. It must be the first keyword within the RULE clause.

This keyword specifies the files to be copied to the path specified by the TARGET keyword. Specify a list of file specifications separated by blanks. You can use the wildcard characters supported by OS/2 or Windows NT.

The directory from which these files are copied depends on how the input files are specified in the teamcpak command:

- If the teamcpak command specifies a source directory, the files specified in the SOURCE keyword come from that directory or subdirectories of it. The full path of the source files is constructed by concatenating the directory from the teamcpak command with a backslash (\), followed by the file specifications found in the SOURCE keyword. You can specify subdirectories in the SOURCE file specifications.
- If the teamcpak command specifies a list of files, these files are first copied to a temporary directory, then copied from there to the TARGET directories. In this case, you can use OS/2 or Windows NT wildcards to specify multiple file names in the SOURCE file specifications, but you cannot specify subdirectories.

In the following example, directory d:\demoapp is specified on the teamcpak command:

```
teamcpak -o "-x -t -m" gather d:\demoapp demoga.pkf
```

The resulting source path is the concatenation of d:\demoapp with the SOURCE file specifications. Therefore, all of the .exe files in the directory d:\demoapp\bin are copied to the target directory e:\demoapp\bin.

```
(DATA
  (TARGETROOT e:\demoapp)
  .
  .
  (RULE
    (SOURCE bin\*.exe)
    (TARGET bin)
  )
  .
  .
)
```

In the following example, a list of input files is specified on the teamcpak command:

```
teamcpak -o "-x -m" gather c:\a.exe c:\b.exe d:\rexx\*.cmd demoga.pkf
```

The resulting source path for the files in the SOURCE clause is the concatenation of the teamcpak temporary directory with the SOURCE file specifications. Therefore, the source for the *.exe files is d:\teamcpak.@@@*.exe. The input files d:\teamcpak.@@@a.exe and d:\teamcpak.@@@b.exe are copied to the directory e:\demoapp.

```
(DATA
  (TARGETROOT e:\demoapp)
  .
  .
  (RULE
    (SOURCE *.exe )
    (TARGET targetroot)
  )
  .
  .
)
```

TARGET Target_path

This keyword is required one time in each RULE clause. It must follow the SOURCE keyword.

The value specified by this keyword is used to construct the target path into which the files specified by the SOURCE keyword are copied. The value of the TARGETROOT keyword is concatenated with a backslash (\), followed by the value of the TARGET keyword.

If you specify targetroot as the value, files are copied directly to the target root directory, not to a subdirectory.

In the first RULE clause of this example, files are copied to the target directory f:\demoapp\bin\files. In the second RULE clause, the target directory is f:\demoapp.

```

(DATA
  (TARGETROOT f:\demoapp )
  .
  .
  (RULE
    (SOURCE *.bin *.dll )
    (TARGET bin\files )
  )
  (RULE
    (SOURCE *.hlp )
    (TARGET targetroot )
  )
  .
  .
)

```

EXITPRIOR, EXITPOST, and EXITREPLACE String...

These keywords are optional. They specify a user exit program to run as part of the gather operation.

To specify an exit that is global to the Gather operation, specify EXITPRIOR or EXITPOST in the DATA clause. You can specify each of these keywords only once in the DATA clause. These keywords must come after all of the RULE clauses. EXITREPLACE cannot be used in the DATA clause.

You can also specify an exit that is specific to one RULE clause. Only one exit keyword is allowed in each RULE clause.

These keywords accept a list of strings separated by spaces. The first string is the name of the program to execute. The strings that follow are its parameters.

Using exit keywords in the DATA clause

When used within a DATA clause, these keywords identify a program or command to be executed within a command shell. EXITPRIOR executes before all RULE statements have been processed; EXITPOST, after all RULE statements.

The exit keywords accept any executable file or command. The exit program must return an integer return value, with zero meaning the exit was successful.

Using exit keywords in the RULE clause

EXITPRIOR, EXITPOST, and EXITREPLACE are optional within a RULE clause. Only one can be specified in any given RULE clause.

When used within a RULE clause, these keywords identify a program or command to be executed within a command shell before, after, or in place of processing of each Gather copy operation. The exit program is called once for each SOURCE specification entry within the SOURCE clause. Parameters are separated by spaces and passed to the exit in this order:

- Any parameters included in the invocation string

- The resolved SOURCE file specifications
- The resolved TARGET specification

The exit keyword accepts any executable file or command. The exit program must return an integer return value, zero meaning successful; it must also accept or ignore the additional Gather parameters added to the end of the invocation string.

When used in the context of the RULE clause, exit keywords must follow the TARGET keyword.

Using exit keywords: an example

In the following example, the first EXITPRIOR statement relates to the DATA clause and specifies a user backup exit program, which executes before performing Gather copy operations. This backup exit is passed two flags. The command stream executed in an OS/2 shell is:

```
"e:\util\backup.cmd \i \t"
```

The second occurrence of the keyword illustrates how to use it in the context of a RULE clause. In this example, an encryption program will run against each source file specification. The exit program is passed the \k:347867 key option, the value for the source specification, and the value for the target specification. In this example, the command stream executed in an OS/2 shell is:

```
"encrypt \k:347867 d:\demoapp\a.exe f:\demoapp\bin":
```

The package file looks like this:

```
(DATA
  (PACKAGEFORMAT gather)
  (TARGETROOT d:\tcws)
  (RULE
    (SOURCE *.exe *.cmd)
    (TARGET exe)
    #|this program will be run for each source file|#
    (EXITPRIOR encrypt \k:347867 )
  )
  (EXITPRIOR "e:\util\backup.cmd \i \t" )
)
```

Chapter 18. Using the Tivoli Software Distribution packaging tool

The Tivoli Software Distribution packaging tool supports automated distribution between a single Tivoli Software Distribution server and its TCP/IP-connected clients. The Tivoli Software Distribution tool works either by itself or in conjunction with TeamConnection's Gather tool to enable you to distribute files through Tivoli Software Distribution. Use of this tool requires you to be familiar with Tivoli configuration and system administration so that TeamConnection can start Tivoli Software Distribution to distribute file packages.

The Tivoli Software Distribution distribution tool must be run on a Tivoli managed node running on any of TeamConnection's UNIX platforms or Windows NT.

The Tivoli Software Distribution distribution tool includes a sample build script named `softdist` (on UNIX platforms) or `softdist.exe` (on Windows NT). It can be run from within a TeamConnection builder. This build script maps TeamConnection build parameters to the command line syntax for the Tivoli Software Distribution tool through the `teamcpak` command line interface.

You can use Tivoli Software Distribution as a builder for packaging in two ways:

- Integrate it with the gather step, so that the Gather tool leaves the package files in a directory from which Tivoli Software Distribution picks them up.
- Use it without the gather step. In this case, the build script for Tivoli Software Distribution must set up the directory and move files into it to interface correctly with the `teamcpak` command.

To simplify the interface, the Tivoli Software Distribution tool uses a select set of options. If you want to take full advantage of Tivoli Software Distribution features, you can import a Tivoli Software Distribution package specification. Importing a package specification provides you access to all Tivoli Software Distribution functions.

The Tivoli Software Distribution tool produces a Tivoli File Package, which is used for distribution.

Using the `teamcpak` command with Tivoli Software Distribution

To start the Tivoli Software Distribution tool, use the `teamcpak` command. This command is found in the directory where the TeamConnection family server is installed. If it is started from a build script, it needs to be in the execution path of the build server.

The complete syntax of the `teamcpak softdist` command follows. You must specify the command parameters in the order shown.

```
teamcpak [-i] [-o "string"] softdist inputFile
```

- i Specifies that only one *inputFile* is specified in the command: an include file containing the list of input files. This parameter is optional.

-o "string"

Specifies that the string listed in quotes be passed to the Tivoli Software Distribution tool. For a list of possible flags to be passed, see "Command line flags".

inputFile

Specifies the files to be copied and the name of the package specification file. You can specify this parameter in these ways:

- Specify the name of an include file, whose contents is a list of input files. One of these input files must be a package specification file with the extension .pkf. In this case, you must also specify the -i parameter.
- Specify a list of two or more files. One of the files must be a package specification file with the extension .pkf.
- Specify the directory from which the files are to be copied and the name of the package specification file.

If more than one package file is listed, the first package file on the command line or in the include file is used, and the others are treated as ordinary files.

The following are examples of specifying input files.

```
teamcpak -i softdist myInputFile
teamcpak softdist d:\inputDir\myPkfFile.pkf inputFile1 inputFile2 . . .
```

Command line flags

You can specify the following flags in the teamcpak command, using the -o parameter. All of these flags are optional.

- a** Create directories on the target.
- c** Clear the target (delete all specified files and directories) before the apply. If you use this option, do not use the -x option.
- t** Overwrite existing files (delete specified files on the target prior to distribution).
- m** Accept input errors, such as missing files and directories from the SOURCE keyword.
- n** Send no notices to Tivoli. If you want to post Tivoli notices, you must configure Tivoli Notices before using this packaging tool.
- p** Preview only; do not actually distribute files.
- r** Reboot the target after distribution.
- x** If an error occurs, leave any distributed files on the target; do not clean up. If you use this option, do not use the -c option.
- k** Keep the Tivoli file package. To enable the Tivoli Software Distribution tool to perform more efficiently, the Tivoli Software Distribution package file is created when the package part is created and then destroyed and recreated whenever the part is modified. Use the -k option to prevent the package file from being destroyed.

Example of the teamcpak softdist command

The following is an example of the teamcpak softdist command.

```
teamcpak -i -o "-a -n -t" softdist Client.lst
```

The -i parameter specifies that the input file Client.lst is to be used. The -o parameter passes the following options to Tivoli Software Distribution:

- -a creates directories on the target.
- -n indicates that no error notices are to be sent to Tivoli Software Distribution.
- -t indicates that any existing files on the target are to be overwritten.

Writing a package file for Tivoli Software Distribution

This section describes the Tivoli Software Distribution package file keywords and their effect on normal processing behavior.

A sample package file named client.pkf is shipped with TeamConnection. You can customize it for your own use.

Syntax rules for a Tivoli Software Distribution package file

Follow these syntax rules when you write a package file:

- Package file keywords must appear in the order shown below.
- Package file keywords must be prefixed with a left parenthesis and must have a corresponding balanced right parenthesis to end the scope of the keyword.
- If the value for a keyword is a string that contains blanks or parentheses, enclose the string in double quotes.
- Default options are supplied for all Tivoli Software Distribution required Tivoli Software Distribution options. Specific options can be set in your TeamConnection package file.
- Comments can appear anywhere within the file. Use the characters #| and |# as delimiters, as shown in the following example:

```
#| This is a comment |#
```

The following shows the syntax of a package file for Tivoli Software Distribution. The keywords must appear in the order shown here. You must supply the values for the strings that are shown in *italics*.

```
(DATA
(PACKAGEFORMAT softdist)
(TARGETROOT filename)
(MANAGER ProfileManager)
(NODES "ManagedNode... PCManagedNode...")
(IMPORT filename |
  [(DISTRIBUTE [FULL | CHANGED])
  [(INSTALLPGM filename)]
```

```

    [(LOGNODE ManagedNode)]
    [(LOGFILE directory)]
  )
)

```

Keywords for a Tivoli Software Distribution package file

DATA This keyword is required. It must be the first keyword in the package file, and it can be specified only once.

All other keywords are nested within the DATA clause.

Example:

```

(DATA
  .
  .
  other keywords go here
  .
  .
)

```

PACKAGEFORMAT softdist

This required keyword must be the first keyword within the DATA clause. It can be specified only once. It tells the teamcpak command that this package file is for Tivoli Software Distribution.

Example:

```

(DATA
  .
  .
  (PACKAGEFORMAT softdist)
  .
  .
)

```

TARGETROOT

This keyword specifies the directory path to which files are to be distributed on the target systems. You can specify only one target root. All target systems use identical target roots.

Example:

```

(DATA
  .
  .
  (TARGETROOT /usr/local/teamc/images)
  .
  .
)

```

MANAGER

This keyword specifies a Tivoli Software Distribution profile manager that you have already created in the Tivoli Software Distribution system.

Example:

```
(DATA
.
.
(MANAGER Distrib1)
.
.
```

NODES

This keyword specifies the nodes to which the files are to be distributed. These must already have been defined to the profile manager as subscriber ManagedNodes or PCManagedNodes. To distribute files to non-subscribers, you need to use Tivoli Software Distribution options set in an import file package definition.

Example:

```
(DATA
.
.
(NODES "tcaix01 tcaix02")
.
.
```

IMPORT

Use this keyword to select Tivoli Software Distribution options not supported in the `-o` parameter of the `teamcpak softdist` command. The *filename* parameter is the name of a Tivoli Software Distribution import file package definition. You can generate an import file using the Tivoli Software Distribution user interface. If you use the IMPORT keyword, then instead of calling the standard Tivoli Software Distribution packaging command the Tivoli Software Distribution tool will call `wimprftp` to get all of the Tivoli Software Distribution configuration options. Using the IMPORT keyword disables other options and causes errors if they are specified.

If you specify the IMPORT keyword, do not specify the DISTRIBUTE, INSTALLPGM, LOGNODE, or LOGFILE keywords.

If you use the INCLUDE option in an import file, it is overridden by the list of files provided to the `teamcpak` command.

Example:

```
(DATA
.
.
(IMPORT importFile)
.
.
```

DISTRIBUTE

Specify FULL to distribute all files or CHANGED to distribute only those changed since the last distribution. The default is FULL.

Example:

```
(DATA
.
.
(DISTRIBUTE CHANGED)
.
.
```

INSTALLPGM

Use this keyword to specify an installation script to be run during distribution on each node that receives files. Specify the full file path name of the script.

Example:

```
(DATA
.
.
(INSTALLPGM /tivoli/fpTeamcAIX/tcinst1.ksh)
.
.
```

LOGNODE

This keyword specifies the system on which the log file is located. The node name you specify must be a managed node. The default is the current build machine or a machine running teamcpak.

Example:

```
(DATA
.
.
(LOGNODE tcaix04)
.
.
```

LOGFILE

This keyword specifies the directory path and file name of the log file on the log node. If you use the LOGNODE keyword, this keyword is required.

Example:

```
(DATA
.
.
(LOGFILE /tmp/softdist.log)
.
.
```

Problem determination for the Tivoli Software Distribution tool

If you are having trouble distributing files using the Tivoli Software Distribution distribution tool, you can use the following tools or teamcpak options to determine what the problem is:

Log file

Check the file name you specified in the LOGFILE keyword for error messages.

Mail Check Tivoli mail messages generated during the distribution.

-k option

Run the teamcpak command with the -k option to keep the package file after the distribution has been run. This allows you to reprocess a distribution from the Tivoli GUI and test variations.

-x option

Run the teamcpak command with the -x option to leave any distributed files on the target.

Trace facility

Run teamcpak with the trace facility. Use this facility only under guidance of an IBM service representative. See the *Administrator's Guide* for more information.

The following message displays when a Tivoli Software Distribution command fails during a distribution.

```
6022-303 Tivoli/Software Distribution %s command failed with return code: RC.
To correct problem use:
- package file parameters LOGNODE and LOGFILE to record Tivoli output,
- packaging option "-k" to keep Tivoli File Package and teamcpak log file
  or "-m" to ignore input errors,
- packaging option "-x" to not clean up files that are distributed,
- TeamConnection Trace facility (see TeamConnection Administration Guide)
- or Tivoli Trace facility (see Tivoli documentation)
```

Sample package file

The following is an example of scripts and items required to automatically execute packaging, distribution, and installation of files in a AIX-based system.

- The **teamcpak** command syntax that will execute subcommands or scripts for the package, distribute, and install functions.

```
teamcpak -i -o "-a -n -t" softdist Client.lst
```

- The **Client.pkf** file you create containing keywords and parameters for distributing and packaging functions.

```
(DATA
(PACKAGEFORMAT softdist)
(TARGETROOT /user/local/teamc/images)
(MANAGER Distribi)
(NODES perlovrs tcaix02)
(INSTALLPGM /tivoli/fpTeamcAIX/tcinstall.ksh)
(LOGNODE tcaix00)
(LOGFILE /tmp/fpTeamcAIX.log)
)
```

- The **Client.lst** file you create containing the list of files passed to **teamcpak**. The first line contains the package file by convention. The example also contains customized installation files (**tcinstall.ksh**), TeamConnection tar files, and an installation script (**tcinst.ksh**).

```

/usr/teamc/tivoli/Client.pkf
/usr/teamc/tivoli/tcinstall.ksh
/tcinstall/v208/fullpak/aix4/tar/client.tar
/tcinstall/v208/fullpak/aix4/tar/msgen_us.tar
/tcinstall/v208/fullpak/tcinst.ksh

```

The following presents an example of a Tivoli installation script (**tcinstall.ksh**) that is copied to the target along with the tar files and the TeamConnection installation script (**tcinst.ksh**), then executed on each NODES entry.

```

#!/bin/ksh
# Clear existing log
#
# You can easily update these.....
#
INST_DIR=/usr/local/teamc
INST_TMP=${INST_DIR}/tcinstl.tmp
INST_OUT=${INST_DIR}/tcinstl.out
INST_ERR=${INST_DIR}/tcinstl.err
INST_LOG=${INST_DIR}/tcinstl.log
IMAGE_DIR=${INST_DIR}/images

rm -f $INST_ERR $INST_OUT $INST_LOG $INST_TMP >/dev/null 2>&1
mkdir -p ${IMAGE_DIR}

exec 1>${INST_ERR}
exec 2>&1
exec 3>${INST_LOG}

# Install VisualAge TeamConnection using responsefile
print -u3 Starting VisualAge TeamConnection installation at 'date'
print -u3 User id= 'id'
print -u3 Input: $*

# Set up installation environment
# - assumes bourne or korn shell
grep DB2_ROOTDIR './.profile'
if (($? != 0))
then
    print -u3 Updating /.profile
    exec 4>>/.profile
    cd /
    print -u3 'DB2 and TeamConnection settings'
    print -u4 'DB2_ROOTDIR=/usr/lpp/ODI/DB24.0'
    print -u4 'export DB2_ROOTDIR'
    print -u4 'PATH=$PATH:$DB2_ROOTDIR/cset/bin'
    print -u4 'export PATH'
    print -u4 'LIBPATH=$LIBPATH:$DB2_ROOTDIR/common/lib'
    print -u4 'export LIBPATH'

```

```

    . /.profile
else
    print -u3 /.profile already updated
fi

# Set up error logging
# - if *.warning is in file (preceeded by spaces and tabs only
grep "^[ ]*\*.warning" /etc/syslog.conf
if (($? != 0))
then
    print -u3 'Updating /etc/syslog.conf'
    touch /var/spool/syslog
    chmod 666 /var/spool/syslog
    exec 4>> /etc/syslog.conf
    print -u4 '*.warning /var/spool/syslog'
    stopsrc -s syslog
    startsrc -s syslog
else
    print -u3 /etc/syslog.conf already updated
fi

# Update services file for tcocto family
grep "tcocto" /etc/services
if (($? != 0))
then
    print -u3 Updating /etc/services
    exec 4>> /etc/services
    print -u4 'tcocto 8888/tcp'
else
    print -u3 /etc/services already updated
fi

# Generate response file
###
### You can change to use enviroment variables!!
###
print -u5 '1'
print -u5 '5'
print -u5 '/usr/local/teamc/images'
print -u5 '/usr/local/teamc'
print -u5 '/usr/local/teamc/nls'
print -u5 'en_US'
print -u5 '/usr/local/teamc/X11'
print -u5 ''
print -u5 'i'

# Run provided TeamC install script
ls -laR ${IMAGE_DIR} >> ${INST_TMP} 2>&1
cd ${IMAGE_DIR}
${IMAGE_DIR}/tcinst.ksh < ${IMAGE_DIR}/tcinst1.response
if (($? != 0))
then
    # Failed installation
    print -u3 TeamC installation failed

```

```
        exit 1
else
    # Clean up installation directory after listing contents
    print -u3 We have successfully copied TeamC installation files
    print -u3 Installation directory contents:
    ls -laR ${INST_DIR} >> ${INST_TMP} 2>&1
fi

cd /
# Remove installation stuff
print -u3 TeamC cleaning up temporary installation directory
rm -rf ${IMAGE_DIR}
cat ${INST_TMP} >> ${INST_LOG}
rm -rf ${INST_TMP}
exit 0

# end of file
```

Part 6. Appendixes

The following appendixes contain various pieces of information that you can refer to as you plan for and use TeamConnection.

Part 7. Appendixes

Appendix A. Environment Variables

You can set environment variables to describe the TeamConnection environment in which you are working. You are not required to set your TC_FAMILY environment variable for the TeamConnection client command line interface. However, if the TC_FAMILY environment variable is not set, the -family must be specified for every client command. See "Setting environment variables" on page 235 for more information about setting environment variables.

The names of the TeamConnection environment variables, the purpose they serve, the equivalent TeamConnection flag, the equivalent Settings notebook field, and the TeamConnection component that uses the environment variable are listed in the following table.

You can override the value you set for an environment variable by using the corresponding flag in a TeamConnection command. When an environment variable has a Settings notebook equivalent, TeamConnection uses the two as follows:

- The environment variable controls the command line interface.
- The Settings notebook controls the graphical user interface.

If there is no Settings notebook equivalent for the environment variable, then the environment variable takes effect regardless of the interface you are using.

To see some of your client settings, you can issue the following command from a command prompt:

```
teamc report -testServer
```

This command returns information like the following:

```
Connect to Family Name:      ptest
Server TCP/IP Name:         amachine.company.com
Server IP Address:          9.1.23.45
Server TCP/IP Port Number:  9999
```

```
Server Specific Information -----
Product Version:            3.0.0
Operating System:           AIX
Message catalog language:   English
Server Mode:                non-maintenance
Authentication Level:       HOST_ONLY
```

Table 3. TeamConnection environment variables

Environment variable	Purpose	Flag	Setting	Used by
LANG	Specifies the language-specific message catalog.			Client, family server

Table 3. TeamConnection environment variables (continued)

Environment variable	Purpose	Flag	Setting	Used by
NLSPATH	Specifies the search path for locating message files.		NLS path	Client, family server
PATH	Specifies where tcadmin is to search for the family create utilities.			Client, build server, family server
TC_BACKUP	Controls whether or not the following commands create backup files. If this environment variable is set to off or OFF, the commands do not create backup files. <ul style="list-style-type: none"> • builder -extract • part -checkout • part -extract • part -merge • part -reconcile 			Family server
TC_BECOME	Identifies the user ID you want to issue TeamConnection commands from, if the user ID differs from your login. You assume the access authority of the user ID you specify.	-become	Become user	Client, build server (except mvs)
TC_BUILDENVIRONMENT	Specifies the build environment name, such as OS/2 or MVS. The value you specify here can be anything you like, but it must exactly match the environment specified for a builder in order for the builder to use this build agent. This value is case-sensitive.	-e		Build server
TC_BUILDMINWAIT	Minimum amount of time to wait (in seconds) between queries for new jobs. Default setting is 5, minimum setting is 3.			Build server
TC_BUILDMAXWAIT	Maximum amount of time to wait (in seconds) between queries for new jobs. Default setting is 15, maximum setting is 300.			Build server

Table 3. TeamConnection environment variables (continued)

Environment variable	Purpose	Flag	Setting	Used by
TC_BUILDOPTS	<p>Specifies build options for sending build log file messages to the screen, and setting the logging level. If you do not specify any of these options, then the build server writes build messages to the build log file (teamcbld.log), and writes a minimum level of messages to the log file. Some possible values are:</p> <ul style="list-style-type: none"> • TOSCREEN (-s) sends the teamcbld.log file to the screen in addition to sending it to a file. • USEENVFILE (-n) <ul style="list-style-type: none"> – writes the changed environment variables to a file called tcbldev.lst instead of setting them in program's environment. The format of the file is variable=value. – writes the list of input files to a file called tcblidin.lst. One file per line, format is pathName type. – writes the list of output files to a file called tcblidout.lst. One file per line, format is pathName type. 	-s, -n		Build server
TC_BUILDPOOL	Specifies the build pool name.	-p	Pool	Build server
TC_BUILD_RSSBUILDS_FILE	Specifies the name of startup files to be used to provide information about build servers to the build process.			Build server
TC_CASESENSE	Changes the case of the arguments in commands, not in queries.		Case	Client

Table 3. TeamConnection environment variables (continued)

Environment variable	Purpose	Flag	Setting	Used by
TC_CATALOG	<p>Specifies a specific file for the TeamConnection message catalog. Sometimes, depending upon the operating system environment, the catalog open command will only look in a particular directory for the catalog. If the host is running multiple versions of TeamConnection, this variable may be used. To set this environment variable, specify the file path name of the message catalog as in the following example:</p> <pre>TC_CATALOG= "/family/msgcat/teamc.cat"</pre>			Family server, oe build server
TC_COMPONENT	Specifies the default component.	-component	Component	Client, make import tool
TC_DBPATH	Specifies the database directory path. Family specific database files reside here.			Family server
TC_FAMILY	Identifies the TeamConnection family you work with.	-family	Family	Build server, client, family server, make import tool
TC_MAKEIMPORTRULES	Specifies the name of the rules file that TeamConnection uses when importing the makefile data into TeamConnection. If you set this environment variable, then you do not have to use the /u option with the fhomigmk command. Specify the full path name of the rules file. If neither this environment variable nor the /u option is used, TeamConnection uses default rules.			Make import tool

Table 3. TeamConnection environment variables (continued)

Environment variable	Purpose	Flag	Setting	Used by
TC_MAKEIMPORTTOP	Strips off the leading part of the directory name when importing parts into TeamConnection. For example, you have parts with the following directory structure: g:\octo\src\incl. To create these parts without the g:\octo structure, you can set TC_MAKEIMPORTTOP=g:\octo before you invoke the make import tool. The parts created in TeamConnection have the directory structure of src\incl.			Make import tool
TC_MAKEIMPORTVERBOSE	Causes the -verbose flag to be added to part commands created by fhomigmk.			Make import tool
TC_MIGRATERULES	Specifies the name of a file containing the rules to be applied for migration of makefiles if the name is not supplied on the fhomigmk command line as a parameter.			Client
TC_MODPERM	Controls whether or not the read-only attribute is set after a part is created, checked in or unlocked in TeamConnection. To cause the read-only attribute to be set, specify TC_MODPERM=ON. To prevent the read-only attribute from being set, specify TC_MODPERM=OFF. The default is TC_MODPERM=ON.			Client
TC_NOTIFY_DAEMON	An alternate way of starting notifyd with the teamcd command. If you set this environment variable, then you do not have to use the -n option with the teamcd command. Specify the full path name of the mail exit to use with notifyd.			Family server
TC_RELEASE	Specifies a release.	-release	Release	Client, make import tool
TC_TOP	Specifies the source directory.	-top	Top	Client

Table 3. TeamConnection environment variables (continued)

Environment variable	Purpose	Flag	Setting	Used by
TC_TRACE	Specifies the variable that lets the user designate which parts should be traced. You should modify this only when directed to do so by an IBM service person. Otherwise it is set to null. To trace all parts, specify TC_TRACE=*			Client, family server, build server
TC_TRACEFILE	Specifies the output (part path and name) of the trace that the user designates using TC_TRACE. The default trace file name is tctrace. For the MVS build server, the default trace file is stdout.			Client, family server, build server
TC_TRACESIZE	Specifies the maximum size of the trace file in bytes. If the maximum is reached, wrapping occurs. The default is one million bytes.			Client, family server, build server
TC_USER	Specifies the user login ID for single-user environments OS/2 and Windows 95 (if not using the login facility). This environment variable is not used in multiuser environments AIX, HP-UX, Solaris, MVS, MVS/OE, and Windows NT. If a user is using the Windows 95 login facility, this environment variable is not used.		User ID	Client, build server
TC_WORKAREA	Specifies the default work area name.	-workarea	Work area	Client, make import tool
TC_WWWPATH	Specifies the path for the HTML helps and image files for Web client.			Client, family server
TC_WWWDISABLED	Disables the Web client.			Family server

The following environment variables are dynamically set by the teamcblid command processing before the build script is invoked:

Table 4. TeamConnection dynamically set build environment variables

Environment variable	Purpose	Flag	Setting	Used by
TC_BUILD_USER	Login of user who initiated the part -build command.			Build server

Table 4. TeamConnection dynamically set build environment variables (continued)

Environment variable	Purpose	Flag	Setting	Used by
TC_INPUT	List of input files (separated by spaces).			Build server
TC_INPUTTYPE	List of input file types (such as TCPart).			Build server
TC_OUTPUT	List of output files.			Build server
TC_OUTPUTTYPE	List of output file types.			Build server
TC_LOCATION	Directory where build script is invoked.			Build server (except MVS build server)

Setting environment variables

For methods of setting your environment variables, refer to your operating system documentation. For example, you can use the following command to set the TC_FAMILY environment variable:

- OS/2 - SET TC_FAMILY=familyName@hostname@portnumber
- UNIX - export TC_FAMILY=familyName@hostName@portNumber

Appendix B. Importing makefile information into TeamConnection

TeamConnection provides a command to help you import makefile information into the TeamConnection database. The `fhomigmk` command reads a makefile and creates the parts in it. Build tree connections are created based on a rules file. The command syntax of the `fhomigmk` command is:

```
fhomigmk /m [makefile]  
         /f [family]  
         /r [release]  
         /w [work area]  
         /c [command file]  
         /u [rules file]  
         /x  
         /s  
         /k
```

You can precede the parameter with either a slash (/) or a dash (-).

The parameters are defined as follows:

/m [*makefile*]

The name of the makefile you want to import into TeamConnection. If you do not specify this parameter, TeamConnection uses `makefile`.

/f [*family*]

The name of the TeamConnection family into which the makefile data will be imported. If not specified, TeamConnection uses the value of the `TC_FAMILY` environment variable. If the value of `TC_FAMILY` is not defined, the value `none` is used.

/r [*release*]

The name of the TeamConnection release into which the makefile data will be imported. If not specified, TeamConnection uses the value of the `TC_RELEASE` environment variable. If the value of `TC_RELEASE` is not defined, the value `none` is used.

/w [*work area*]

The name of the TeamConnection work area into which the makefile data will be imported. If not specified, TeamConnection uses the value of the `TC_WORKAREA` environment variable. If the value of `TC_WORKAREA` is not defined, the value `none` is used.

/c [*command file*]

The name of the command file that will be produced and saved. If this file already exists, commands created by the specified makefile are appended to the existing contents.

/u [*rules file*]

The name of the rules file that TeamConnection will use when importing the makefile data into TeamConnection. If not specified, TeamConnection uses the value of the `TC_MAKEIMPORTRULES` environment variable. If no rules file is

found, TeamConnection uses default rules. "Creating a rules file" on page 239 explains the rules, the format of this file, and the default rules.

- /x** Specifies that you want to run the command file that was produced by the `/c` parameter.
- /s** Specifies that the build tree is to be displayed after the command is issued. If specified, the command file is run even if the `/x` parameter is not specified.
- /k** Specifies that you want TeamConnection not to erase the intermediate files it uses to process this command. This might be useful in debugging problems that arise during the import. However, in general, you will not specify this parameter. When specified, the following intermediate files are saved:

modified makefile

A modified form of the imported makefile. The command invocations (of things like linkers and compilers) are replaced by calls to a TeamConnection command that captures dependency data. To find the cause of import errors, type the following command at an OS/2 command line:

```
nmake -f mod_make
```

create file

A list of all the objects referenced by the makefile that should be created in the TeamConnection database.

connect file

A list of all the objects referenced by the makefile that should be connected to other objects in the TeamConnection database. Each line contains one dependency relationship in the format `<child>` `<parent>`.

TeamConnection provides an environment variable, `TC_MAKEIMPORTTOP`, that when set strips off the leading part of the directory name. For example, you have parts with the following directory structure: `g:\octo\src\inc\`. Because you want the parts created without the `g:\octo` structure, you set `TC_MAKEIMPORTTOP=g:\octo` before you invoke the make import tool. The parts created in TeamConnection have the directory structure of `src\inc\`.

Another environment variable provided by TeamConnection, `TC_MAKEIMPORTVERBOSE`, when set causes the `-verbose` flag to be added to part commands.

The following is an example of invoking the make import tool:

```
fhomigmk /m Mymak /w mywork /s /u myrules
```

In this example, the makefile called `mymak` is used to create a temporary command file containing TeamConnection commands. The commands are formed based on the rules defined in the file `myrules`. The family and release used in the commands are those specified in the environment variables `TC_FAMILY` and `TC_RELEASE`. The work area

used in the commands is `mywork`. After the commands are issued, the resulting build tree is shown using the TeamConnection GUI.

Creating a rules file

The import rules file is a text file that describes how you want TeamConnection to create and connect parts. In this file you supply a set of rules, one per line, using the following syntax:

file mask

The mask specifying the names of the files to which this rule applies. The `*` and `?` wildcards are supported. For example, you could specify file names such as `*.cbl`, `abc*.cpp`, or `foo\src*.obj`.

type The type of contents of the files to which the rule applies when they are stored in TeamConnection as a part. Allowed values are `binary`, `text`, `none`, or `ignore`. If you specify `ignore` as the file type, then all files that match the file mask are bypassed.

builder The name of the TeamConnection builder to be associated with the part. The builder is not created for you. If you specify a builder, it must exist in TeamConnection before you run `fhomigmk`. A value of `none` means that no builder will be associated with the part.

parser The name of the TeamConnection parser to be associated with the part. The parser is not created for you. If you specify a parser, it must exist in TeamConnection before you run `fhomigmk`. A value of `none` means that no parser will be associated with the part.

connect

How the part will be connected to other parts in TeamConnection. The following values are allowed:

- `input`
- `output`
- `dependent`
- `none`

When `none` is specified, the part is not connected to another part even though a dependency was found for the part in the make file. For example, when you indicate `none` for a file mask of `*.h` files, the `*.h` files are created in TeamConnection, but not connected to the files that include them. The value you will use most often is `input`.

content

Where the initial content of the part can be found:

- `none` indicates that the part is initially created as *empty*.
- `directory\` indicates to concatenate with the name of the file in the makefile. This is where the contents are expected to be found.

- * indicates to use the name in the makefile, relative to the current working directory.

For example, if a makefile specifies a file src\abc.cbl and the makefile specifies f:\mysrc\, the content is expected in f:\mysrc\src\abc.cbl. For a file of *.cbl, the content is expected in src\abc.cbl relative to the current working directory.

parameters

The build parameters to be attached to the part. Enclose the parameter in double quotes if it has spaces. Use the value none to indicate no parameters.

component

The TeamConnection component that will contain the part. If none is specified, the value of the TC_COMPONENT environment variable is used.

As TeamConnection processes each part referenced in the makefile, it looks for a rule that matches the part name. If a match is found, the rule is used. The rules are searched from top to bottom. The first matching rule is used.

Comments are denoted by a pound sign (#) in the first column.

Columns are separated by spaces.

A sample rules file, called fhomigmk.rul, is supplied with TeamConnection. Use this file to help you create a rules file that is appropriate for your development environment.

The following is a simple example of an import rules file:

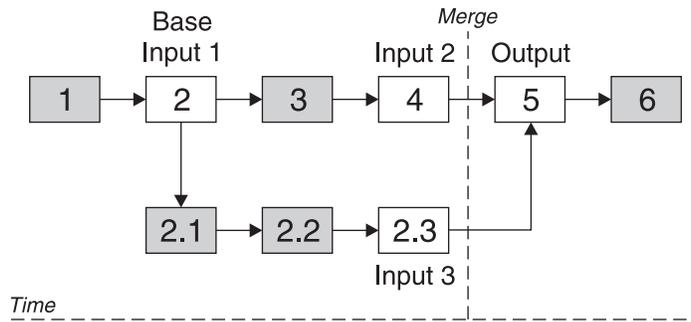
```
<top of file>
#-----
# file mask  type    builder  parser  connect  content  parameters  component
#-----
#
# *.exe      binary linker  none    input  none     /Debug   ship
# *.obj      binary icc     none    input  none     /Ti+    objects
# *.cpp      text   none    cplus  input  *        none     source
# *.h*      text   none    cplus  none    *        none     source
<end of file>
```

If you do not specify a rules file in the /u parameter of the fhomigmk command, TeamConnection uses the value of the TC_MAKEIMPORTRULES environment variable. If no rules file is found, TeamConnection uses the following default rules:

```
<top of file>
#-----
# file mask  type    builder  parser  connect  content  parameters  component
#-----
#
# *.*       text   none    none    none    none     none     root
```

Appendix C. TeamConnection Merge

The TeamConnection VisualMerge provides a way for you to merge two or three selected files together to make one single file. You can select options for viewing differences and collisions, as well as view the composite output of the merged files. Both a command line (`tcmerge`) and a graphical user interface are provided. The figure below describes the merge of three files into a single file. Items one through six represent the main stream of development for a particular part or file. Item 2 represents the base or common file from which two streams proceed. At the time of the merge, the most recent version in the two streams are items 4 and 2.3. In this scenario, the user selects the base or common file, item 2, as input 1. The user also uses item 4, the latest version in development, as input 2. The latest version in the branch is represented as item 2.3, or input 3. Item 5 represents the outcome of the merge. Once the merge is complete a window appears with the three original files and the final merged file.



The following is the command line syntax for TeamConnection VisualMerge with the parameter abbreviations shown. The first set of brackets encloses the abbreviation that may be used; the remainder shows the full parameter name. For example, `[-re[place]]` indicates that `-re` is the abbreviation for the **replace** parameter.

```
tcmerge <file1|directory1> <file2|directory2> [<file3|directory3>]
[-ti[tles] <titlenames>] [-out <file|directory>]
[-prime[out] <file|directory>] [-re[place]]
[-ignoreco[lumns] <list of ranges to ignore>]
[-ignoreb[lanks] <|l|t|a|l|t>] [-ignoreca[se]
[-ignoreu[nique] <file|directory>]
[-ignoreb[lanklines]] [-auto[merge]] [-nologo]
```

If you do not specify input parameters at the command prompt, a **Merge Files** GUI is presented to assist you in entering required and optional parameters. If only input and output files or directories are provided, the following command defaults will be in effect:

- The output file or directory is not primed with one of the input files or directories.
- No columns will be ignored in difference calculation.
- No blanks will be ignored in difference calculation.

The following lists the parameters which may be used with the `tcmerge` command.

Parameter	Description
<file1 directory1>	First file or directory to merge. A combination of filenames and directory names is not valid.
<file2 directory2>	Second file or directory to merge.
[<file3 directory3>]	Third file or directory to merge. (Optional)
[-ti[itles] <titlenames>]	Names of files or directories being merged.
[-out <file directory>]	The file or directory into which the merged differences will be written. This flag is optional.
[-prime[out] <file directory>]	Primes the specified input file or directory to the output window.
[-re[place]]	The filename with the <code>-out</code> flag will be replaced. If not specified, and the filename specified with the <code>-out</code> flag exists, the <code>-replace</code> flag must be specified.
[-ignoreco[lumns] <list of ranges to ignore>]	List of column ranges to ignore during the compare. Entries must use the format <startColumn,endColumn> with no blanks within an entry. Specify one column <Column> where start and end would be the same number.
[-ignoreb[lanks] <l t a lt>]	Specifies blanks to ignore. Specify one: l (leading), t (trailing), lt (leading trailing), or a (all blanks). Specify: l for blanks at the beginning of a line; t for blanks at the end of a line; lt for blanks at the beginning and end of the line; a for all blanks on a line.
[-ignoreca[se]]	Case of the characters is ignored. For example, there is no difference when an uppercase "C" and a lowercase "c" are compared.
[-ignoreb[lan]klines]]	Any blank lines are ignored during the compare.
[-ignoreu[nique] <file directory>]	An input file or directory is ignored during the compare.
[-auto[merge]]	Automatically merge differences between files or directories. When conflicts found, a GUI is invoked for you to resolve the conflict. Use this with the <code>-out</code> parameter to keep input source files unchanged. Note: Using this parameter is equivalent to using the <code>TeamConnection automrg</code> command. See the Commands Reference for additional information.
[-nologo]	Hides the TeamConnection logo.

Appendix D. Enabling an OS/2 Workframe project for TeamConnection

TeamConnection lets you create a Workframe version 3 project that has TeamConnection options as well as a set of TeamConnection actions. For each project, you specify on the Project Options window the values for these options. By doing this, you logically connect a Workframe project with a set of TeamConnection parts. This makes it easier for you to perform TeamConnection actions, such as checking parts in and out, directly from the Workframe.

Creating a TeamConnection-enabled Workframe project

Follow these steps to create a Workframe project that is enabled for TeamConnection.

1. On an OS/2 command line, type the following command and press Enter:

```
fhotcini.cmd
```

This command creates a TeamConnection **Project Smarts** catalog on your desktop. (If you have already created this catalog, there is no need to perform this step again for additional projects.)

2. Open the TeamConnection **Project Smarts** catalog. Select the TeamConnection project, and select the **Create** pushbutton.
3. Specify the location for the TeamConnection project you want to create; then select **OK**.

When the action completes, you will see a TeamConnection Project on your desktop.

Setting up your project options

Options are provided so that you can set up each TeamConnection Workframe project. To set the options, do the following:

1. Select **Tools Setup** from the project's Views pull-down menu.
2. Select the **Project Options** or **File Options** menu from any of the TeamConnection actions.

The following options are provided:

Family The TeamConnection family.

Release

The TeamConnection release.

Work area

The TeamConnection work area in which you will perform TeamConnection actions.

Query mask

Any valid TeamConnection -where clause for parts. Leave blank to see all parts. (This is used in the project's **Show Parts** action.)

Show filter

Check this if you want to display the PartFull Filter window instead of using the query mask in the **Show Parts** action.

Profile Names the rules file to use for the **Import Make** action. Specify the fully qualified name unless you are sure it will be found in your path correctly. Select the **Find** push button if you need help.

Using your TeamConnection Workframe project

You can perform a set of TeamConnection actions from within your project:

- “Project actions” lists the actions you can perform without selecting a part.
- “Part actions” on page 245 lists the actions you can perform against a selected TeamConnection part.

Project actions

OS/2 has three project actions. They are:

1. Invoke the TeamConnection GUI
2. Show all parts from current context
3. Build All (build project target)

Extract part

Displays an unprimed Extract Parts window.

Checkout part

Displays an unprimed Check Out Parts window.

Checkin part

Displays an unprimed Check In Parts window.

Unlock part

Displays an unprimed Unlock Parts window.

Lock part

Displays an unprimed Lock Parts window.

Create part

Displays an unprimed Create Parts window.

Build part

Displays an unprimed Build Parts window.

View part contents

Displays an unprimed View Part Contents window.

View part information

Displays an unprimed View Part Information window.

Edit part

Displays an unprimed Edit Part window.

Show parts

If the project attribute **Show filter** is not set, issues a query based on the project attribute's query mask. If the project attribute **Show filter** is set, displays the PartFull Filter window.

Part actions**Extract part**

Displays the TeamConnection Extract Parts window to extract the selected part.

Checkout part

Displays the TeamConnection Check Out Parts window to check out the selected part to the work area specified in the project options.

Checkin part

Displays the TeamConnection Check In Parts window to check in the selected part to the work area specified in the project options.

Unlock part

Displays the TeamConnection Unlock Parts window to unlock the selected part.

Lock part

Displays the TeamConnection Lock Parts window to lock the selected part.

Create part

Displays the unprimed TeamConnection Create Parts window.

Build part

Displays the TeamConnection Build Parts window to start a build of the selected part.

View part contents

Displays the TeamConnection View Part Contents window for the selected part.

View part information

Displays the TeamConnection View Part Information window for the selected part.

Edit part

Displays the TeamConnection Edit Part window for the selected part.

Import makefile

Imports the information contained in the selected makefile into the TeamConnection database.

The **Import makefile** action is restricted to files with the extension .mak. The other actions in this list apply to files of all types.

Using your project: a simple scenario

Suppose you are working on a defect in the family FAMILY1, release REL1_1. You have created a TeamConnection work area called SANDBOX to work in. You want to use the Workframe to access your TeamConnection parts. Here is what you might do.

1. Create a TeamConnection Workframe project called DefectABC.
2. Open the project. Select **Tools setup** from the View pull-down menu.
3. Select any of the actions. Press mouse button 2 to display the context menu; then select **Project options** or **File options** from the context menu. The result is a window in which you can specify the TeamConnection information about the project.
4. Specify the family FAMILY1, the release REL1_1, and the work area SANDBOX. Check the **Show filter** check box. Select **OK**.
5. Specify the general Workframe attributes of the project using the project's Settings notebook. These attributes include information such as the location of the OS/2 files for the project. For example, in this scenario, you specify that you want this project to contain all files in the directory c:\defect_abc, which is initially empty.
6. Select **TeamConnection** → **Show files** from the Project context menu. The PartFull Filter window is displayed. Specify the filter criteria; then select **OK**. For example, specify that you want to see all the parts with extensions .cpp, .exe, and .hpp. The Parts window is displayed.
7. Select the parts client.exe, server, client.cpp, client.hpp, server.cpp, and server.hpp. Select **Extract** from the context menu.
8. On the Extract Parts window, type c:\defect_abc in the **Target directory** field and select **OK**. Now you can interact with these parts directly from the Workframe project.
9. In the Workframe project, run the ipmd.exe debugger until you determine the cause of the problem. Suppose you find the bug is in client.cpp.
10. Go back to the Parts window. Select **client.cpp** from the list of parts, and select **TeamConnection** → **Checkout part** from the context menu for the object. The part is checked out to the SANDBOX work area.
11. Edit the file to fix the problem; then select **TeamConnection** → **Checkin part** to check the part back into SANDBOX, the TeamConnection work area from which it was checked out.
12. Build the part by selecting **TeamConnection** → **Build part** on the context menu for the file client.exe.
13. When the build completes, extract the resulting executable by selecting **TeamConnection** → **Extract part** from the file's context menu.
14. Run the executable to verify that the problem has been fixed.

Appendix E. Enabling a Workframe/NT project for TeamConnection

TeamConnection lets you create a Workframe/NT project that has TeamConnection options as well as a set of TeamConnection actions.

Workframe/NT doesn't use the TeamConnection dialogues. Instead, WorkFrame/NT uses monitored commands to invoke TeamConnection.

You must perform the following steps to integrate TeamConnection with WorkFrame/NT:

1. Install VisualAge C++
2. Install the TeamConnection Client
3. Make a backup copy of the WorkFrame Configuration File which is named vacpp.iws. This file is located in the MAINPRJ directory where VisualAge C++ is installed.
4. Replace or merge the WorkFrame Configuration File with the TeamConnection version of this file.

If you have made changes to the WorkFrame Configuration File named vacpp.iws, you can use the TeamConnection Merge tool to view the differences between the two files and selectively merge the text into a single file. You can run the TeamConnection merge tool from a Windows command line. The syntax is:

```
TCMERGE FILE1 FILE2
```

Note: It is very important that you make a backup copy of the WorkFrame Configuration File (vacpp.iws) before making any changes to the file. Any changes made to the configuration files will not be migrated to a future version. Errors in the configuration file can prevent WorkFrame/NT from operating correctly.

5. Reboot to activate the WorkFrame/NT changes.

Setting up your project options:

All TeamConnection commands are monitored from inside an editor session. Environment variables are used to specify valid parameters for the project actions.

To set or change environment variables within WorkFrame/NT, from the Project's View pull-down menu, select Settings → Environment Variables.

You can use the following environment variables:

TC_FAMILY

The TeamConnection family. Required for all commands.

TC_USER

The current TeamConnection user. Required for all commands.

TC_BECOME

The current TeamConnection user. Required for all commands.

TC_RELEASE

The current TeamConnection release. Required for all commands.

TC_COMPONENT

The current TeamConnection component. Required for all commands.

TC_WORKAREA

The name of the current TeamConnection work area. NULL is the default if a work area is not specified.

TC_TOP

Root of the work directory. Required for all commands.

TC_MAKEIMPORTRULES

The path and file name of the import rules file. Required for all Import Makefile.

DPATH Includes all work directories for the WorkFrame project. Required for Import Makefile only if the project has more than one work directory. DPATH may already be set in the environment.

Each time you change the TC command parameters, you must also change the appropriate environment variable before selecting that TC action.

If you are using multiple directories in the WorkFrame project, you must set the TC_TOP environment variable to specify the top directory structure for each directory that you use.

Using your TeamConnection WorkFrame project

You can perform a set of TeamConnection actions from within your project:

- “Project actions” lists the actions you can perform without selecting a part.
- “Part actions” on page 249 lists the actions you can perform against a selected TeamConnection part.

Project actions

WorkFrame/NT has three project actions. They are:

1. Invoke the TeamConnection GUI
2. Show all parts from current context
3. Build All (build project target)

Part actions

The current release and work area are specified with environment variables. See "Setting up your project options:" on page 247.

TC Extract Part

Extract the specified part.

TC View Part Information

View information about the specified part.

TC Build Part

Build the specified output part.

TC Create Part

Create the specified part in TeamConnection.

TC Unlock Part

Unlock the specified part.

TC Import Makefile

Imports the information contained in the selected makefile into the TeamConnection database.

TC Checkin Part

Check in the specified part.

TC View Part Contents

Displays the contents of the specified part.

TC Lock Part

Locks the specified part.

TC Checkout Part

Checks out the specified part.

WorkFrame/NT uses selective part options:

- You can only checkin, checkout, or view contents of a non-binary part.
- Makefiles can only be imported.
- For WorkFrame/NT all TeamConnection actions are prefaced with "TC". For example, TC Checkout Part.

Appendix F. Enabling and Using the VisualAge TeamConnection Enterprise Server Bridge

Overview of the VisualAge TeamConnection Enterprise Server Bridge

VisualAge TeamConnection for Smalltalk provides a repository with support specifically for highly-interactive, prototyping environments that emphasize iterative development, such as VisualAge for Smalltalk or VisualAge Generator. A bridge from VisualAge TeamConnection Enterprise Server to VisualAge TeamConnection for Smalltalk (and to VisualAge Generator through VisualAge TeamConnection for Smalltalk) provides access to the software configuration management (SCM) support provided by VisualAge TeamConnection for Smalltalk, along with the scalable, enterprise-level support provided by VisualAge TeamConnection Enterprise Server. VisualAge TeamConnection Enterprise Server's ability to manage all development artifacts (not just source code), to share information in a common model, and to integrate multiple tools and multiple languages across the enterprise on a single baseline extends the capabilities of software development groups. The VisualAge TeamConnection Enterprise Server Bridge (also referred to as *the bridge* in this documentation) provides essential integration for VisualAge tools that use VisualAge TeamConnection for Smalltalk as their day-to-day operational library.

VisualAge Generator Version 3.0 has access to the bridge through its use of VisualAge for Smalltalk Pro Version 4.0 or later, which can interface directly with VisualAge TeamConnection Enterprise Server library support. The bridge supports the import and export of VisualAge Generator objects (*parts*) to and from VisualAge TeamConnection Enterprise Server.

VisualAge TeamConnection for Smalltalk provides a collaborative component development environment for application development and integration using fine-grained object languages, such as Smalltalk. The VisualAge TeamConnection for Smalltalk repository is designed for languages that run on the universal virtual machine (uVM). The repository includes persistence, versioning, and configuration management.

VisualAge TeamConnection Enterprise Server can be used to manage artifacts (*parts*) that need to be shared with non-uVM based languages or tools for purposes of build management, problem tracking, and other configuration management functions. These artifacts can be exported from VisualAge TeamConnection for Smalltalk to VisualAge TeamConnection Enterprise Server through the bridge and stored as TeamConnection parts.

Objects stored in a TeamConnection database can be queried and retrieved back into the VisualAge TeamConnection for Smalltalk development environment as needed. The units of storage in TeamConnection include exported VisualAge for Smalltalk and VisualAge Generator components (such as applications and configuration maps) and large grained objects (files). Small-grained objects, such as VisualAge Generator data items, are imported and exported as constituents of applications. The data items in an

application are exported to VisualAge TeamConnection Enterprise Server in a form that makes their definitions available to other tools through the data model.

Scope of this documentation

This documentation is intended for users and administrators installing and using the bridge. It is assumed that you familiar with both the VisualAge for Smalltalk and TeamConnection products.

The following subsections describe the mechanics of enabling the bridge for VisualAge for Smalltalk Pro (Version 4.0 or later), the process of exporting VisualAge for Smalltalk and VisualAge Generator components to TeamConnection, and the process of importing these components back into VisualAge TeamConnection for Smalltalk. See the VisualAge Generator documentation for tool-specific details. TeamConnection information related to change tracking and build processing are addressed in the TeamConnection documentation.

Many terms used by VisualAge for Smalltalk and TeamConnection are problematic because the tools may define these terms differently. *Release* and *component* are typical examples. To avoid any ambiguity, such terms may preceded by the name of the tool they are applied to, such as *TeamConnection release*.

Description of the bridge

Basic functionality

It makes sense to describe the functionality of the bridge from the perspective of a Smalltalk developer, because it is through the Smalltalk image that the user drives the bridge. The bridge is an import/export facility for three types of entities:

- Smalltalk configuration maps
- Smalltalk applications
- Files residing on local and networked file systems that are accessible through the image

The bridge allows a Smalltalk developer to store any of these entities in a TeamConnection database and retrieve them at a later time. Configuration maps and applications must be versioned before they can be exported. This enforces the notion that the Smalltalk developer uses the bridge and TeamConnection to maintain baselines rather than for managing work-in-progress.

Developers use VisualAge TeamConnection for Smalltalk's fine-grained support to facilitate the process of shared development in open editions of components on a daily basis. At appropriate junctures, components are versioned and promoted to VisualAge TeamConnection Enterprise Server, where together with other project elements, they form a baseline across an entire project. The resulting baseline may contain objects such as program elements, files, and metadata.

From the perspective of the VisualAge TeamConnection Enterprise Server user or administrator, the bridge allows the Smalltalk image to function as a TeamConnection client, storing and retrieving parts in a TeamConnection family database.

How the bridge communicates with TeamConnection

The bridge functions are initiated from within the VisualAge for Smalltalk environment. Each operation that interacts with VisualAge TeamConnection Enterprise Server runs for some time in the Smalltalk image, but at some stage will make use of functions built into an appropriate version of the TeamConnection client and server. The bridge itself is implemented in Smalltalk and in one of the DLLs in TeamConnection.

The unit of transfer used by the bridge for Smalltalk components is a VisualAge TeamConnection for Smalltalk library. Each library stored in TeamConnection contains one of the following:

- a Smalltalk application and its released subapplications (and their released subapplications, and so forth)
- a configuration map without any of its released applications

Note: Subapplications cannot be exported through the bridge without an enclosing application.

VisualAge TeamConnection for Smalltalk libraries are stored in VisualAge TeamConnection Enterprise Server databases as TeamConnection parts. When the bridge exchanges a library, the target in a TeamConnection database is specified by a TeamConnection context. A TeamConnection context is comprised of the following parameters:

- Family name
- Release name
- Work area name

Note: Each TeamConnection context can contain only one version of any named application or configuration map. This is unlike VisualAge TeamConnection for Smalltalk libraries, in which multiple versions of a named Smalltalk component can co-exist.

The bridge is aware of the various relationships between VisualAge TeamConnection for Smalltalk components. When an application is transferred through the bridge, all of its released subapplications are transferred with it. When a configuration map is transferred through the bridge, the bridge will also transfer the released applications in separate operations. Depending on a user-specified setting, the bridge can also transfer required maps of configuration maps.

Preparing to use the VisualAge TeamConnection Enterprise Server Bridge

The bridge is delivered as a configuration map suitable for loading into a VisualAge for Smalltalk Version 4.0 (or later) image. The library, TCEMBR.DAT, will contain the configuration map **ENVY/Manager-TeamConnection Bridge** and, for VisualAge Generator build support, **VAGen ENVY/TC Bridge**.

These configuration maps should be imported into your development library so that it can be loaded by all of the users sharing that library. The step-by step instructions are described in "Installing and activating the bridge" on page 255.

Usually, the **Library Supervisor** or the first user to use the bridge will perform this operation and then inform other users that the tool is available in the library.

The sections that follow describe the steps necessary to set up the bridge and verify that it is functional.

Setting up the bridge environment

The following information is especially pertinent to the individual(s) responsible for bridge setup and administration.

Prerequisites

Before the bridge will work, you must have the following:

- VisualAge for Smalltalk Pro Version 4.0 or later installed. See the VisualAge for Smalltalk documentation to confirm that you have the appropriate hardware and software prerequisites available.
- A VisualAge TeamConnection Enterprise Server family that is running.
- A TeamConnection GUI client installed on the machine where you are running your Smalltalk image.

Note: This release of the bridge only runs on OS/2 and Windows platforms.

You should verify that you are able to communicate with the relevant VisualAge TeamConnection Enterprise Server family by using the TeamConnection GUI client. If you cannot communicate with the TeamConnection server in this manner, the bridge will definitely not function correctly.

Environment variables

The bridge relies on the user to specify the various parameters that make up the TeamConnection context. By default, the bridge will query the variables in the environment that the image is running. These variables, **TC_FAMILY**, **TC_RELEASE**, and **TC_WORKAREA**, are used as initial values for the default TeamConnection context.

There are two additional environment variables that can be defined for the bridge, as follows:

- **TC_COMPONENT** is used as the default TeamConnection component for parts stored in TeamConnection through the bridge. If **TC_COMPONENT** is not defined or is empty, the value `root` is used.
- **TC_RELATIVE** is used to specify the initial destination path for files retrieved from TeamConnection through the bridge. If **TC_RELATIVE** is not defined or is empty, the current directory according to the image is used.

It is not necessary to define any of these variables for the bridge to work. Defining them only makes setting up the default bridge configuration in the image easier for a bridge user.

A system administrator may want to have the environment variables automatically defined in a network login script. When a user logs into a LAN and then uses the bridge, the user will be provided with the defined values as hints for setting up the default TeamConnection configuration.

Installing and activating the bridge

The bridge is loaded into the image like any other configuration map using the **Load** option from the **Editions** menu of the **Configuration Maps Browser**.

Once the bridge is loaded, the submenu **TeamConnection Bridge** will appear on the **Tools** menu of the **System Transcript** window. This submenu is referred to as the *bridge menu*. The bridge menu is the launching point for all of the bridge operations.

“Importing and loading the bridge” provides step-by-step instructions for the bridge loading process.

Importing and loading the bridge

Follow these steps to load the bridge:

1. Open the VisualAge for Smalltalk Pro - Client.
2. Go to the **System Transcript** window and select **Browse Configuration Maps** from the **Tools** pulldown menu.
3. In the Configuration Maps Browser window, select **Import** from the **Names** pulldown menu.
4. A dialog will prompt you to enter the full path name of the library that you want to import. For purposes of activating the bridge, you will need to supply a TeamConnection pathname (determined by where you have installed TeamConnection) for the file called TCEMBR.DAT. Select the **OK** pushbutton to continue and display the **Selection Required** window.
5. In the **Selection Required** window, select ENVY/Manager-TeamConnection Bridge in the **Names** list, which will prime the **Versions** list with a version number.

6. Select the version in the **Versions** list and move it to the **Selected Versions** list using the right-arrow pushbutton.

Note: If you are using a Version 3 VisualAge TeamConnection Enterprise Server, you should select version R1.02. If you are using a Version 2 TeamConnection server, you should select version R1.01.

7. For the additional interoperability with VisualAge Generator described in “Using the bridge: a simple scenario for VisualAge Generator developers” on page 264, you must also import the configuration map called VAGen ENVY/TC Bridge, as described in the previous steps.
8. Select the **OK** pushbutton to initiate the import process. During the process of importing the TCEMBR.DAT file into the VisualAge for Smalltalk Pro manager.dat file, the **System Transcript** window will issue a message stream that confirms the success of the import.
9. In the **Configuration Maps Browser** window, select ENVY/Manager-TeamConnection Bridge from the **Names** list.
10. Select the item (there should only be one available) in the Editions and Versions list.
11. Click mouse button 2, and select **Load** from the pop-up menu.

Note: For the additional interoperability with VisualAge Generator described in “Using the bridge: a simple scenario for VisualAge Generator developers” on page 264, you must also load VAGen ENVY/TC Bridge, as described in the two previous steps. ENVY/Manager-TeamConnection Bridge *must be loaded first*.

12. After the application loading progress dialog completes without errors, the bridge should be functional. You can close the **Configuration Maps Browser** window at this time.

Testing the bridge

To verify that the bridge is active and ready for VisualAge TeamConnection for Smalltalk component export/import functions, follow these steps:

1. Go to the **System Transcript** window and select the **Tools** pulldown menu. Then select **Default Properties** from the **TeamConnection Bridge** cascade menu. This will display the **Default Properties** notebook.
2. Verify that the TeamConnection family in the **Family** field on the **Context** page of the **Default Properties** notebook is appropriate for your project. You may need to coordinate your access to the family with your family administrator.
3. Select the **Test Server** pushbutton. If the bridge is properly configured, the server connection test will return an information window that provides server-specific information. Select the **OK** pushbutton to dismiss the server information window.

You are now ready to export VisualAge TeamConnection for Smalltalk components to a VisualAge TeamConnection Enterprise Server.

Note: When exiting VisualAge for Smalltalk Pro - Client, you should save your image so that the bridge will be preserved for future use.

Using the VisualAge TeamConnection Enterprise Server Bridge

You can perform TeamConnection functions on VisualAge TeamConnection for Smalltalk components, provided that you supply parameters necessary to identify a *bridge configuration*. Bridge configuration parameters are defined by the **Default Properties** notebook, as described in “Setting default properties”.

Each time the bridge interacts with TeamConnection, it uses the parameters in a bridge configuration to ensure that the behavior of the operation is in accordance with the users’ specifications. Because specifying a configuration for each operation would be time-consuming and most operations would use the same configuration, you can specify a default configuration. Each time the user initiates an operation, you can use the default configuration or modify it.

The default configuration is stored in the image so that once it is setup, it will be maintained until the bridge is reloaded from the library.

Setting default properties

To set properties for import and export actions across the bridge, open the **Default Properties** notebook as described in “Testing the bridge” on page 256. The **Default Properties** notebook contains four pages of settings, as follows:

- Context page
- Operations page
- Import page
- Export page

Each page of the **Default Properties** notebook includes the following controls:

Show this dialog when exporting and importing checkbox

The **Show this dialog when exporting and importing** checkbox specifies whether the dialog should be shown each time an import/export operation for the bridge is initiated by the user. If the dialog is shown, it gives the user the opportunity the default configuration for a particular operation only.

push buttons

- OK** Saves the current settings as default setting. This option may not be available if some fields are left incomplete or contain invalid values.
- Cancel** Closes the **Default Properties** notebook and ignores any changes made in the dialog.

Defaults

Updates the dialog fields with the values in the current default bridge configuration.

Reset Updates the dialog fields with the initial values that are set when the bridge is first loaded into the image.

Context page

The context page is used to specify the TeamConnection family, release, and work area used as the context for the default bridge configuration.

Family Use this field to input the name of your TeamConnection family server. Select the **Test Server** pushbutton to return an information window that provides server-specific information. If you cannot successfully communicate with the TeamConnection server, you may have specified an invalid family name. Your TeamConnection family administrator may be of some assistance at this point.

Release

The TeamConnection release. By selecting the **Query releases** pushbutton, you can prime the **Release** field drop-down menu with valid release choices based on the **Family** field value.

Work area

The TeamConnection work area in which you will perform TeamConnection actions. By selecting the **Query work areas** pushbutton, you can prime **Work area** field drop-down menu with valid work area choices based on the **Release** field value.

Note: Any communication with a TeamConnection server takes time. Querying the available releases and work areas typically takes a few seconds, which is the reason that this data is not automatically used to populate the dialog.

Operations page

The **Operations** page determines how operations are performed in TeamConnection, including whether operations are forced and how parts in the database are locked.

Force The **Force** and **Don't force** radio buttons are mutually exclusive.

In TeamConnection terms, *force* is an indication that changes should be forced into the VisualAge TeamConnection Enterprise Server repository, possibly breaking links with the part in other version contexts. Its intent is to indicate that, although the specified version might not match the current set of versions applicable to the object in the persistent store, the changes in those versions specified in the version string are to be made, breaking the links to those current versions not specified.

The force option is important only if you specify that a part version is to be locked. If you want to retrieve or store a locked part in a particular release or work area that is linked to another release or work area, you might want to specify the force option when you are checking in or checking out the part,

even if someone else might have the part checked out in another context. See the discussion of locking below for a description of TeamConnection locking options.

Locking

These mutually-exclusive radio buttons enable you to instruct TeamConnection on how to manage the locking behavior for parts that you are exporting to or importing from the TeamConnection repository.

Obtain and release

Also known as *optimistic* locking, TeamConnection will attempt to check out the part(s) before checking in changes that you have made in the VisualAge TeamConnection for Smalltalk environment. If this action is successful, the part(s) will not be locked in TeamConnection after the export.

Obtain and retain

TCCS will attempt to check out the part(s) before checking in changes that you have made in the VisualAge TeamConnection for Smalltalk environment. If this action is successful, the part(s) will remain locked in TeamConnection after the export.

Retain For parts already locked in TeamConnection, after changes are exported from VisualAge TeamConnection for Smalltalk, the locked parts should remain locked (i.e., the lock is *retained* by the original owner).

Release

For parts already locked in TeamConnection, after changes are exported from VisualAge TeamConnection for Smalltalk the locked parts should be unlocked, and therefore available to other developers in that context.

Import page

The **Import** page provides default settings options when importing VisualAge TeamConnection for Smalltalk components or files previously exported to a TeamConnection database.

Configuration Maps

If the **Import all required maps too** checkbox is checked, it specifies that when a configuration map is imported, its required configuration maps (along with any other required configuration maps, recursively) should be retrieved from TeamConnection as part of the import action. If this option is enabled, and a configuration map being imported does have required maps, the maps can only be imported if they actually exist in the TeamConnection database.

Note: The checkbox is checked as the default.

Destination for Files

The **Destination path for files** field identifies the target (base) directory for imported files.

Replacing Existing Files

These mutually-exclusive radio buttons enable you to select a desired default method for overwriting files (or not) in your working target directory.

Ask user

This choice enables you to choose which files are to be overwritten.

Do not replace existing files

Files that currently exist in the target directory will not be overwritten.

Replace existing files

Files that currently exist in the target directory are automatically overwritten.

Export page

The **Export** page provides default settings options for exporting VisualAge TeamConnection for Smalltalk components or files to a TeamConnection database.

Storage in TeamConnection

The **Component** field identifies the TeamConnection target component for your export action. This component designation, along with TeamConnection family, release, and work area values supplied in the **Context** page of the **Default Properties** notebook, is necessary to define the context for any new TeamConnection parts created by an export action.

Configuration Maps

If the **Export all required maps too** checkbox is checked, it specifies that when a configuration map is exported, its required configuration maps (along with any other required configuration maps, recursively) should be exported TeamConnection. If this option is enabled, and a configuration map being exported does have required maps, the maps can only be exported if they actually exist in the TeamConnection database.

This option is used to prevent version mismatches when a configuration map requires other configuration maps, as in the following case:

1. For a configuration map that requires other configuration maps, you do an export *with* the required maps.
2. At some later time, you export again *without* the required maps.
3. When you attempt to import *with* the required maps, the import may fail, because a configuration map level in TeamConnection does not match the level previously exported from VisualAge TeamConnection for Smalltalk.

Note: The checkbox is checked as the default.

Exporting VisualAge TeamConnection for Smalltalk components to TeamConnection

The **TeamConnection Bridge** cascade menu provides an **Export** choice, which offers the following choices:

- Configuration Maps
- Applications
- Files

Note: You must have the appropriate authority to update all parts associated with the configuration maps, applications, or files to be exported.

As a general rule, it is advisable to export applications and configuration maps along with any configuration maps required by these VisualAge TeamConnection for Smalltalk components to avoid version mismatches. If you make a change to an application, it is important to update all the exported configuration maps that contain the application and to export all of the configuration maps again.

Note: The **Export all required maps too** checkbox located on the **Export** page of the **Default Properties** notebook defaults to this behavior.

The following describes two simple cases in which a mismatch might occur:

1. Export a configuration map that contains several applications.
2. Make a change to one of the contained applications.
3. Export the changed application only.
4. Attempt to import the configuration map.

or

1. Export two configuration maps that contain the same application.
2. Make a change to the common application and export only one of the configuration maps that contains the application.
3. Attempt to import the second configuration map.

As the number of programmers authorized to version components and the complexity of your applications increase, so does the possibility for these types of problem to occur. Therefore, it is important to coordinate update authority in such a way that all affected parties are notified about configuration changes, and that someone in the development group has authority over all levels of components. It may also be advisable to limit export actions to higher levels of authority than you have previously.

Exporting components is a substantial operation that typically takes at least ten to twenty seconds (possibly minutes for a large collection of components). Such an operation begins with the bridge exporting the components to temporary VisualAge TeamConnection for Smalltalk libraries and then generating detailed descriptions of the library contents for the benefit of TeamConnection. To guarantee atomicity and minimize the number of times that the bridge must communicate with TeamConnection (thus avoiding unnecessary overheads), all components are transferred in one primitive operation.

Even a single configuration map usually counts as more than one component, because it typically contains at least one release application. Once the primitive operation is

invoked, control of the process is in the TeamConnection client code, which is effectively blocked against the TeamConnection server. Because the Smalltalk image is blocked waiting for the primitive to return, the user interface will not update, and the user cannot halt the operation.

Exporting configuration maps and applications

The process for exporting VisualAge TeamConnection for Smalltalk-based configuration maps and applications to a TeamConnection family database includes the following steps:

1. Select **Configuration Maps** or **Applications** from the **Export** cascade menu. You will be prompted to select an appropriate version of the configuration map or application that you want to export.

Note: VisualAge TeamConnection for Smalltalk components must be versioned in VisualAge for Smalltalk before being exported to TeamConnection.

2. In the **Selection Required** window, select the desired configuration map or application in the **Names** list, which will prime the **Versions** list with a version number.
3. Select the version in the **Versions** list and move it to the **Selected Versions** list using the right-arrow pushbutton. Because only one version of any named configuration map can exist in a TeamConnection context, it is only possible to choose one version for any particular name.
4. Select the **OK** pushbutton to initiate the export process.
5. If the **Show this dialog when exporting and importing** option has been set in the default bridge configuration, you will be presented with the **Export Properties** notebook, which is primed by values in the **Default Properties** notebook. If you are satisfied with the current values in the **Export Properties** notebook, select the **OK** pushbutton to initiate the export process.

If the export succeeds without errors, a message is logged to the **System Transcript** window. Users are informed of any errors with a message box.

Exporting files

The process for exporting VisualAge TeamConnection for Smalltalk-based files to a TeamConnection family database includes the following steps:

1. Selecting **Files** from the **Export** cascade menu.
2. You will be prompted to select the files that you want to export. You can add to or delete files from the list using the **Add**, **Remove**, or **Remove All** pushbuttons.
3. Select the **OK** pushbutton to initiate the export process.
4. If the **Show this dialog when exporting and importing** option has been set in the default bridge configuration, you will be presented with the **Export Properties** notebook, which is primed by values in the **Default Properties** notebook. If you are satisfied with the current values in the **Export Properties** notebook, select the **OK** pushbutton to initiate the export process.

If the export succeeds without errors, a message is logged to the **System Transcript** window. Users are informed of any errors with a message box.

Importing VisualAge TeamConnection for Smalltalk components from TeamConnection

The **TeamConnection Bridge** cascade menu provides an **Import** choice, which offers the following choices:

- Configuration Maps
- Applications
- Files

The process for importing any of these VisualAge TeamConnection for Smalltalk components is essentially the same, and includes the following steps:

1. Select **Configuration Maps**, **Applications**, or **Files** from the **Import** cascade menu.
2. If the **Show this dialog when exporting and importing** option has been set in the default bridge configuration, you will be presented with the **Import Properties** notebook, which is primed by values in the **Default Properties** notebook.
3. When you are satisfied with the current values in the **Import Properties** notebook, select the **OK** pushbutton.
4. You will be prompted to supply a query pattern to further reduce the number of candidates for import. Select the **OK** pushbutton to launch the query of the TeamConnection context that you have specified up to this point.

Note: Use the wildcard characters (* and ?) as delimiters for your queries.

5. A list of VisualAge TeamConnection for Smalltalk components matching your query is returned. Each of these components exists in a TeamConnection database specified by the TeamConnection context in the configuration used for this operation. Select the objects you want to import from this list and select the **OK** pushbutton to initiate the import action.
6. In the case of configuration maps and applications, the selected components will be imported into the default VisualAge TeamConnection for Smalltalk library that the image is connected to. For files, the selected files will be written to the path specified by the **Destination path for files** option in the bridge configuration used for this operation. If any of the files already exist, they may be overwritten, or the user may be prompted, depending on the value of the **Replace existing files** option.

If the import succeeds without errors, a message is logged to the **System Transcript** window. Users are informed of any errors with a message box.

Using the bridge: a simple scenario for VisualAge Generator developers

The following scenario is a generalized case used to illustrate the way that VisualAge Generator developers might use the bridge to accomplish change tracking and build processing. An actual implementation requires that a Smalltalk development team begin with versioned VisualAge TeamConnection for Smalltalk components and a plan for sharing common applications.

After you have installed the bridge, you can export a versioned configuration map to a TeamConnection family database. For VisualAge Generator developers, this means that you can generate programs, tables, and map groups using the TeamConnection build interface. See the VisualAge Generator *Generator's Guide* and "Part 4. Using TeamConnection to build applications" on page 127 in this document for details.

Scenario assumptions

For purposes of describing the scenario, the following assumptions are established:

- A development team using VisualAge Generator wants to perform problem tracking and build generation.
- A family is created in TeamConnection with a release r1, defined with a track-driver process (i.e., all part changes are made in reference to work areas).
- A build server has been started to handle build requests for generation, and similarly for preparation.
- Data item definitions and records used to access data in a database are kept in a "common" application, while programs and their other associates are kept in a separate application.

Note: This assumption enables the scenarios to include application prerequisites.

Exporting VisualAge TeamConnection for Smalltalk components to TeamConnection

To prepare for exporting the VisualAge TeamConnection for Smalltalk components to TeamConnection, perform the following activities:

1. **In TeamConnection:**
 - a. Create a feature called f1 and accept the feature.
 - b. Create a work area called wa1 for implementation of the feature.
2. **In VisualAge TeamConnection for Smalltalk:**
 - Create applications for the common data and for other VisualAge Generator parts.
3. **In VisualAge Generator Developer:**
 - Create programs and their associates.
 - Create generation option, linkage table, resource association, link edit, and bind parts as necessary.

4. In VisualAge TeamConnection for Smalltalk:

- a. Create a configuration map that gathers the common data application and the application containing all the other parts.
- b. The class developers version their classes.
- c. The class owners release versioned classes into the two applications and the application managers version the applications.

If more than one developer has been working on the feature, each may have opened a new edition of a part's class extension, so a merge of the method editions will have to be performed.

- d. The configuration map manager releases the versioned applications into the configuration map and versions the configuration map.
- e. An administrator uses the bridge to export the configuration map to the work area associated with feature f1. See "Exporting VisualAge TeamConnection for Smalltalk components to TeamConnection" on page 260 for bridge export instructions.

After the bridge export action is completed, there will be a part for the configuration map and for each of its applications, and parts for each class and method in each application.

Object mapping in TeamConnection

After a bridge export action, the following parts are created in TeamConnection in wa1 for f1 in the component and release specified as context for the export action:

- For each application of the configuration map there will be an application part.
- For each class and method, a part with a name qualified by the application name for uniqueness, as described in Table 5.

The bridge must map VisualAge TeamConnection for Smalltalk components to part names in TeamConnection in such a way that the parts can be retrieved in a reusable form when they are imported from TeamConnection back into the VisualAge TeamConnection for Smalltalk environment.

Table 5. Name generation mapping for the VisualAge TeamConnection Enterprise Server Bridge

Class Type	Naming Convention	Mapped Name Example
EmLibrary	<class_name_of_blob_object>.<name_of_blob_object>	EmApplication.MyApp, EmConfigurationMap.MyConfigMap
EmConfigurationMap	<config_map_name>	MyConfigMap
EmApplication	<application_name>	MyApp
EmSubapplication	<app_name>.<subapp_name>	MyApp.MySubapp
EmClass OR EmClassExtension	<app_name>.<subapp_name>.<class_name>	MyApp.MyClass, MyApp.MySubapp.MyClass
EmInstanceMethod OR EmClassMethod	<app_name>.<class_name>.<method_name> OR <app_name.subapp_name>.<class_name>.<method_name>	MyApp.MyClass.MyMethod, MyApp.MySubapp.MyClass.MyMethod

See the VisualAge Generator *Generator's Guide* for additional information related to generation part output names in TeamConnection.

Build generation

The VisualAge Generator *Generator's Guide* provides detailed VisualAge Generator build generation instructions. The following overview is provided to place these activities in the context of the bridge:

1. In VisualAge Generator Developer:

- For each program, the Options Override (OVR) part that has been exported to TeamConnection creates an initial build tree for VisualAge Generator applications in TeamConnection.

Refer to the VisualAge Generator *Generator's Guide* for more details on the OVR part.

2. In TeamConnection (build function):

- a. The build administrator selects the EZEPREP collector of the initial build tree of a program proxy in wa1 for f1, and requests a build.
- b. TeamConnection places the generator build event on the build queue, and the generator build server detects a new build event that it can service.
- c. The build server invokes the generator build script, which parses the name of the generation configuration map name from the generated part's build parameters.
- d. The build script invokes the generator, which imports the configuration map and its references to the generation manager. The VisualAge TeamConnection for Smalltalk manager used by generation is identified by a VisualAge INI file on the generation build server.

The bridge determines whether each application referenced already exists in the generation VisualAge TeamConnection for Smalltalk manager, and only imports an application if that version of the application is not already in the manager.

Note: For VisualAge Generator builds, you can use the environment variable TC_ENVY_REFRESH to control when VisualAge Generator builders will import the required configuration map from TeamConnection.

TC_ENVY_REFRESH can be used to affect the following behaviors:

- If TC_ENVY_REFRESH=null, the configuration map will not be imported from TeamConnection if that version of the configuration map is already in the connected VisualAge TeamConnection for Smalltalk manager.
- If TC_ENVY_REFRESH=null, the configuration map will be imported from TeamConnection if that version of the configuration map is not already in the connected VisualAge TeamConnection for Smalltalk manager.
- If TC_ENVY_REFRESH=notnull, the configuration map will always be imported from TeamConnection. This setting is important if you use VisualAge DataAtlas to modify data elements that are dependents of

the configuration map. In that case, if the data elements have been modified since the configuration map was exported to TeamConnection, the builder will warn you that the configuration map is not synchronized with its dependent data elements *only if* TC_ENVY_REFRESH=notnull. Such a warning allows you to import the changed data elements into a new edition of the configuration map and export the resulting new version to TeamConnection, before trying the build again

Setting TC_ENVY_REFRESH is only relevant in the environment of the TeamConnection build server that performs the VisualAge generator builds.

- e. The generator uses the bridge to update the outputs and dependencies in the build tree.
- f. If there are tables and/or map groups used by the program, the generator determines whether there is already a build tree for them. If not, initial build trees are added for them using the program's OVR part.
- g. TeamConnection re-examines the build tree of the EZEPREP collector and determines that new build events have been added to the build scope for preparation of the generation outputs, and possibly for generation and preparation of tables and map groups. Build events are started to complete the preparation of generation outputs, and generation and preparation of tables and map groups if necessary.

Note: See the VisualAge Generator *Generator's Guide* for greater detail on this process.

3. In TeamConnection (change control):

- A project administrator completes the fix record(s) for the feature f1 and adds the work area wa1 to a system test driver. Eventually the driver is committed to the release and the feature is completed.

Making a change to a member

1. In TeamConnection:

- a. Defect d1 is created and accepted in TeamConnection
- b. Work area wa2 is created for the implementation of the defect d1.

2. In VisualAge TeamConnection for Smalltalk:

- a. An application manager creates new edition of an application that requires a change.
- b. A developer makes a change to one or more parts.
- c. The class developer of the changed parts versions the class, the class owner releases the class into the new edition of the application, and the application manager versions the application. If more than one defect is in progress, the class owner must release only the versions that apply for defect d1.

- d. The configuration map owner opens a new edition of the configuration map used to generate the program being changed, and releases the new application version into the configuration map. The configuration map owner versions the configuration map.
 - e. The administrator uses the bridge to put the configuration map back into the work area wa2 for defect d1.
3. **In TeamConnection (build function):**
- a. Build administrator builds the program(s) affected by the change. This can be done by selecting the preparation collector for each program and requesting a build, or by selecting a collector for a subsystem, and building the subsystem collector. Only programs, tables, or map group affected by the changes to proxy members will be rebuilt.
 - b. The generation process continues as it did for the initial build (see “Build generation” on page 266 for details), except that there should be no need to add new build trees unless a new table or map group was added to a program being built
4. **In TeamConnection (change control):**
- A project administrator completes the fix record(s) for the defect d1 and adds the work area wa2 to a system test driver. Eventually the driver is committed to the release and the feature is completed.

Appendix G. Source Code Control User's Guide

Differences between other source code control providers and TeamConnection

The purpose of this document is to help Visual Basic, Visual C++, and PowerBuilder users, make TeamConnection their Visual environments source code control provider. This document assumes the reader is a new user of TeamConnection, but has some familiarity with source code control.

Note: This component is for a client operating in a Windows 95 or NT environment. However, it may accept files from other environments providing the Win platforms can access them.

The following shows the version level supported by TeamConnection:

- PowerBuilder 5.0.04 or higher
- Visual Basic version 5.0 with SP2 applied
- Visual C++ version 5.0

Projects vs Families

Most source code control providers group all code into projects. TeamConnection uses an object oriented approach that provides much more control over the software product while allowing greater flexibility. Projects have one dimension of control. Development environments like Visual Basic group all of their files into projects. Using projects to group source code has several limitations. First, the source code control system is limited to providing just version control. While version control is useful, once the enterprise-size organization is reached, it is often not sufficient to control just versions of the source code. TeamConnection provides not only versioning but defect and feature tracking, build and driver management, access control, and much more. TeamConnection uses families, releases, components, and work areas for management and control.

TeamConnection uses several layers of control. The highest level is the family. The family is the name of the data base, where TeamConnection stores all of the code, the versions, and all other information related to the code. A family represents a complete and self-contained collection of TeamConnection users and development data. Data within a family is completely isolated from data in all other families. One family cannot share data with another. It is important to know the name of the family where TeamConnection will store your code and associated information.

A part in TeamConnection is a collection of data that is stored by the family server. This can include files, text, objects, binary objects, or modeled objects. Parts can be stored by a user, a tool, or generated from other parts, such as when a linker generates an executable file.

Components are used to organize the data in a family. Components are arranged in a hierarchical tree structure, with a single top component called root. The component owns the parts that may be in it, and controls access to the parts. Once you are given access to a component, you have access to all the parts and subcomponents in that component. The component also controls the process that TeamConnection uses, for example, to report and fix a defect. Within each family, development data is organized into groups called components. The component hierarchy of each family includes a single top component, initially called root, and descendants of that root. Each child component has at least one parent component; a child can have multiple parents.

The release is somewhat analogous to a project. A release is a logical grouping of the components that make up a product. An application is likely to contain parts from more than one component. Because you probably want to use some of the same parts in more than one application, or in more than one version of an application, TeamConnection groups parts into releases. A release is a logical organization of all parts that are related to an application; that is, all parts that must be built, tested, and distributed together. Each time a release is changed, a new version of the release is created. Each version of the release points to the correct version of each part in the release. Each part in TeamConnection is managed by at least one component and contained in at least one release. One release can contain parts from many components; a component can span several releases. Each time a new development cycle begins, you can define a separate release. Each subsequent release of an application can share many of the same parts as its predecessor. You need to know the name of the release.

A work area is basically a view of a release. For example, a work area can be opened for each defect that needs to be fixed. More than one programmer can work in the same work area at the same time. A programmer can have more than one work area active at a time. A release contains the latest integrated version of each of its parts. As users check parts out of the releases, update them, and then check them back in, TeamConnection keeps track of all these changes, even when more than one user updates the same part at the same time.

You need to know the name of the work area in which you will be working. A good practice is to create and name a work area after the defect being addressed in the work area. For example, name work area W1557 for defect 1557. You can create a work area if you have the authority in TeamConnection, but this must be done through the TeamConnection GUI.

For more information about families, releases, components, work areas, parts, and what you can do with them, see your TeamConnection Documentation.

Installing the TeamConnection source code control DLL

Before you can use the integrated support from your development environment you must install TeamConnection and the TeamConnection Source Code Control DLL. If you are using TeamConnection Version 2.0.8 or later, the source code control DLL is already installed.

Note: If you have not already done so, follow the directions and install the TeamConnection client for your workstation. The following directions assume that you have successfully installed the TeamConnection client.

Connecting TeamConnection to an IDE

In order for TeamConnection to be your source code control provider, both of the following Registry structure conditions must be in place:

- **HKEY_LOCAL_MACHINE\SOFTWARE\SourceCodeControlProvider** contains the key name **ProviderRegKey**. If TeamConnection is your source code control provider, the value (Data) for this key name should be **SOFTWARE\IBM\VisualAge TeamConnection**. The key name may currently point to another provider if you have another provider installed.
- **HKEY_LOCAL_MACHINE\SOFTWARE\SourceCodeControlProvider\InstalledSCCProviders** contains the key names for the available (installed) source code providers. The key name for TeamConnection is **IBM VisualAge TeamConnection Enterprise Server** and the data is **SOFTWARE\IBM\VisualAge TeamConnection** when the TeamConnection client is installed.

Note: There may be other source code control providers currently designated (e.g., **Microsoft Visual SourceSafe** with data **SOFTWARE\Microsoft\SourceSafe**), depending on which other SCC providers have been installed.

Removing the TeamConnection Source Code Control DLL

To change the default source code control system for Visual Basic, Visual C++, or PowerBuilder, change the value in the ProviderRegKey to the registry key of another provider.

To remove TeamConnection, leave the value in the ProviderRegKey blank.

Using TeamConnection as your source code control provider

Once the installation procedure is complete, starting your development environment automatically links the TeamConnection Source Code Control DLL.

Before you start

There are several things you must know before you can start using TeamConnection as your source code control provider. If you are not sure of this information, contact your administrator. Your family administrator can help you find the following information:

- Family, defined by the following attributes:
 - Name
 - TCP/IP Host

- Port Address
- Component
- Release
- WorkArea

You also need to know the project name. The project name is used by your development tool to relate to the TeamConnection attributes of family, release, work area, and component by the Source Code Control DLL.

Opening a project

One of the few differences you see when using TeamConnection as your source code control provider occurs when you open or save a project (depending on the IDE you are using). When you open or save the new project, the **TeamConnection Settings** window opens. At the top of this window is a field with your development project name. In addition to the project name field, there are fields for family attributes, work area, release, and component. If this is a new project, these fields are blank. If this is not a new project, the fields contain the previous values. You can change these values only when this window is open. If at anytime you decide to change any of these values, you might need to first close the project and reopen it.

Once all the fields are filled in, select **OK**. The project will open. If you select **Cancel**, the source code control system might disconnect from the development environment until another project is opened.

Under some versions of Visual Basic, for example, projects automatically close and open after certain operations. This causes this TeamConnection Settings window to open at times when it may appear unnecessary. When this occurs, select **OK**. If you select **Cancel**, you might be in a state that requires shutting down and restarting Visual Basic to reconnect the source code control system.

Integrated features

Once you open a project you can use the integrated features of the development environment to access your files in TeamConnection. The development environment keeps track of the files that are known to TeamConnection, and the checkout status of each file. For example, the development environment keeps track of files checked out to other users.

The exact steps necessary to perform each of the following actions depend on the development environment being used. However, for a given IDE, the steps are the same regardless of the source code control provider. For example, if you check out a file in the Visual C++ development environment when it is connected to Visual SourceSafe, the steps you use are exactly the steps you use when Visual C++ is connected to TeamConnection.

Check-in

The steps to check-in a file vary by the development environment. In most cases pressing mouse button 2 when the mouse pointer is over a file icon of a file checked out to you, brings up a menu that includes the file check-in option. Selecting the file checkin option opens the **Check-In** window. Checking the **keep checked out** box on the **Check-In** window sets the keep locked flag, TeamConnection saves the file, but keeps it checked out to you. Selecting **OK** causes the TeamConnection part check-in function to execute and the file is checked in.

Note: Depending on development environment and software levels, some features that follow may be unavailable.

Check-out

Similar to check-in, the check-out action can be started by pressing mouse button 2 on the file icon of a file not already checked-out. Check-out calls the TeamConnection Part Check-out function.

Uncheck-out

A checked out file can be unchecked out. Again this action can often be started by right clicking the file icon of a file that is checked out. Uncheck-out calls the part unlock function in TeamConnection.

Get Version

Rather than check out a file, you can also get the latest version of the file. Get Version calls the TeamConnection Part Extract function.

Adding Files to source code control

Adding a file that is not already under source code control places the selected file into the source code control system. Add calls the TeamConnection part create function.

Properties

Selecting **Properties** from a pull-down menu opens the properties GUI. Information that TeamConnection needs to correctly check out and check in parts is provided here. For example, the work area field changes each time an existing work area is integrated and a new work area is created.

Full features of TeamConnection

Most development environments allow you to evoke TeamConnection from the pull-down menus. In Visual Basic, TeamConnection appears as an option in the Add-Ins pull-down menu. In Visual C++, TeamConnection appears in the Source Code Control option of the Tools pull-down menu. PowerBuilder uses the PowerBuilder Library Icon to

designate TeamConnection as the Source Code Control provider. From the TeamConnection GUI you can create new work areas (if you have the correct authority), retrieve previous versions of a part, open or process defects, and perform many other actions against parts.

Migrating project data bases

One key issue for programmers and project managers moving from another source code control system to TeamConnection is how to migrate the database of projects. The following describes one way to bring the current level of code for a small to medium sized project into TeamConnection.

Migrating an existing project: The following example illustrates the simplest way to migrate an existing Visual Basic source code control database into TeamConnection. Lets say we were using the ABC source code control system, and we are going to migrate our project, Baylout, to TeamConnection. The idea is to extract all the files in Baylout using the ABC source code control system, and then add them as parts in TeamConnection. Follow the steps below to perform this migration.

1. Make ABC the default source code provider. To do this, set the registry key **ProviderRegKey** to point to the registry entry for source code control provider ABC. See "Installing the TeamConnection source code control DLL" on page 270 for more information on how to perform this step. Once you complete this step, ABC will be the Source Code Control provider when we open Visual Basic.
2. Start the Visual Basic development environment.
3. Open project Baylout.
4. Extract all the files to your system.
5. Exit the Visual Basic development environment.
6. Edit the registry key **ProviderRegKey** to be:
SOFTWARE\IBM\TeamConnection

See "Installing the TeamConnection source code control DLL" on page 270 for more information on how to perform this step. TeamConnection is now the default source code control provider and is attached when the development environment starts.

7. Restart the Visual Basic development environment.
8. Open project Baylout again. Save the Baylout project. The **TeamConnection Settings** window will display.
9. When the **TeamConnection Settings** window opens it will have Baylout listed as the project. Fill in the values for family attributes, release, component, and work area, then select **OK**.
10. Add the files to **TeamConnection** following the steps in the Visual Basic development environment.
11. Repeat these steps until all of your projects are migrated to TeamConnection.

Starting a new project

Starting a new project in Microsoft Visual Basic, Visual C++, or PowerBuilder is essentially the same regardless of the source code provider. The only operational difference is that the **TeamConnection Source Code Control Settings** window opens at some point. When the **TeamConnection Source Code Control Settings** window opens, enter the family attributes, component, work area, and release.

Starting Visual Basic: In order to use TeamConnection source code control with Visual Basic, a subset of Microsoft Source Safe DLLs are required. After you have Microsoft Source Safe installed, you must make the following changes:

1. Add or update the following lines in the vbaddin.ini file :
vbsc=0
SccAddIn.SourceCodeControlAddIn=1
2. When you start Visual Basic after making these changes, you will see another check box in the **Add-Ins->Add-In Manager** panel called **Source Code Control Add-In**. You should check the new box, and uncheck the **Source Code Control** checkbox.

These steps should enable you to use the TeamConnection source code control with Visual Basic

To create a new project in Visual Basic, do the following:

1. Start Visual Basic.
2. Create and save a new project.
3. Select **Add Project to TeamConnection** from the **TeamConnection** option in the **Add-Ins** menu. The **TeamConnection Settings** window will display. Fill in the family attributes, release, component, and work area, then select **OK**.
4. The **Add Project to TeamConnection** window opens. Select the files you want to add. Type a comment in the comment field. (Visual Basic requires that a comment be entered.) Select **OK**.

Starting Visual C++: To create a new project under Visual C++, do the following:

1. Start the Visual Developers Studio as normal.
2. From the **File** menu, select **New**. The **New Project Workspace** window will open.
3. On the **New Project Workspace** window, do the following:
 - a. Select the type of project
 - b. Type a name
 - c. Select **Create**.
4. Select **Project->Source Control->Add to Source Control** to open the **TeamConnection Settings** window.
5. On the **TeamConnection Settings** window, enter the family attributes, release, component, and work area. Then, select **OK**.
6. Files can now be added to the project using the **Insert** menu.

7. To place the files under source code control, select the **Add to Source Code Control** option of the **Tools** menu.

Starting PowerBuilder: To create a new project under PowerBuilder, do the following:

1. Start PowerBuilder.
2. Select the **Library** icon.

Note: If TeamConnection Source Code Control is already installed (the **TeamConnection Settings** window will open), you can proceed with Step 5

3. From the **Source** menu, select **Connect**. The **Connect** dialog box will display.
4. Select SCC API from the **Connect** list box, and then select **OK**.
5. The **TeamConnection Settings** window will open. On the **TeamConnection Settings** window, enter the family attributes, release, component, and work area. Then select **OK**.

Note: The **TeamConnection Settings** window may redisplay. If this happens, select **OK** again.

6. Add parts to the TeamConnection. Select all parts that you want to add to TeamConnection, including the project folder in the **Library** window. You can select all of these objects simultaneously by dragging the mouse pointer while depressing the Control (**CTRL**) key.
7. From the **Source** menu, select **Register**. This may take some time. These registered parts should be added to the TeamConnection work area previously specified in the **TeamConnection Settings** window

Appendix H. Supported expandable keywords

TeamConnection supports expandable keywords in text files. When a file containing expandable keywords is extracted from TeamConnection, the current value of each keyword is added to the file. This information can help you identify what version of source code is used for your deliverables.

TeamConnection supports the following keywords.

Keyword	Description
\$ChkD;	The time and date stamp applied during check in.
\$FN;	The file name complete with its path.
\$KW;	The start of keyword expansion. It expands to @(#).
\$EKW;	Keyword expansion is ended until the next \$KW=@(#); keyword.
\$Own;	The user ID of the owner of the component that manages the part.
\$Ver;	The version of the part in TeamConnection.

The following examples show lines of code that change in a text file as a user extracts a part. The text file used in this example is filex.hdr.

```
#ifndef_filex_hdr_
#define_filex_hdr_
    static char _filex_hdr[]="$KW; $FN; $Ver; $ChkD; $EKW;";
#endif
```

TeamConnection ignores keywords until it finds a \$KW; keyword. It then expands all keywords until a \$EKW; keyword is found. If the semicolon (;) following a keyword is omitted, the keyword is not expanded.

No change occurs when the part is checked in to TeamConnection. However, when the part is extracted, the keyword variables are updated. The following example shows how the keywords are expanded.

```
#ifndef_filex_hdr_
#define_filex_hdr_
    static char _filex_hdr[]="$KW=@(#) $FN=bin/filex.hdr; $Ver=1:1;
$ChkD=1998/03/20 18:13:19; $EKW;";
#endif
```

In the previous example, each keyword and its value appears in the output. The value of the keyword is replaced each time the part is extracted. If you do not want the keyword to appear in the output, add a minus sign (-) after the dollar sign (\$). For example, if the statement is prepared as follows:

```
static char _filex_hdr[]="$-KW; $-FN; $-Ver; $-ChkD; $-EKW;";
```

Then the expanded keywords will look like:

```
static char _filex_hdr[]="$@(#) bin/filex.hdr 1:1
1998/03/20 18:13:19 $-EKW;";
```

Be aware that if a file is extracted, then locked and checked in, the version information can no longer be updated because the keyword does not appear in the output.

Appendix I. Authority and notification for TeamConnection actions

TeamConnection ships with IBM-supplied authority groups, interest groups, component processes, and release processes. Your family administrator can modify these preconfigured authority groups, interest groups, and processes to fit the needs of your organization.

Each authority group consists of actions normally performed by a particular type of user. Your family administrator can modify these groups or create new ones to reflect the needs of your organization.

Authority groups provide explicit authority to perform the actions included in each group. You might also have implicit authority to perform certain actions according to the objects that you own. Authority groups are defined in a file called `authorit.Id`.

To determine your authority groups, from the Actions pull-down menu, select Lists → Access lists → Show authority actions. On the Show authority actions window select an action.

Each notification group consists of actions normally of interest to a particular type of user. Your family administrator can modify these groups or create new ones to reflect the needs of your organization. Interest groups are defined in a file called `interest.Id`.

To determine your interest notification groups, from the Actions pull-down menu, select Lists → Notification lists → Show interest actions. On the Show authority actions window select an action.

The following table lists all of the TeamConnection actions, the required level of implicit and explicit authority to perform the action, and the users who are notified when an action is performed. To explicitly assign authority to a user, add the user's ID to a component's access list.

Note: The user who performs the action is excluded from the notification that is sent out after the action is successfully completed.

For this action	These users have authority	These users are notified
AccessCreate	<ul style="list-style-type: none">Component ownerExplicitly defined for the component where access is being added	User being given new access, subscribers
AccessDelete	<ul style="list-style-type: none">Component ownerExplicitly defined for the component where access is being altered	User whose access was deleted, subscribers
AccessRestrict	<ul style="list-style-type: none">Component ownerExplicitly defined for the component where access is being restricted	User whose access is being restricted, subscribers

For this action	These users have authority	These users are notified
ApprovalAbstain	<ul style="list-style-type: none"> Approval record owner Explicitly defined for the component that manages the associated release 	Approval record owner, subscribers
ApprovalAccept	<ul style="list-style-type: none"> Approval record owner that manages the associated release 	Approval record owner, subscribers
ApprovalAssign	<ul style="list-style-type: none"> Approval record owner Explicitly defined for the component that manages the associated release 	New and original approval record owners, subscribers
ApprovalCreate	<ul style="list-style-type: none"> Work area owner Explicitly defined for the component that manages the associated release 	New approval record owner, subscribers
ApprovalDelete	<ul style="list-style-type: none"> Explicitly defined for the component that manages the associated release 	Approval record owner, subscribers
ApprovalReject	<ul style="list-style-type: none"> Approval record owner Explicitly defined for the component that manages the associated release 	Approval record owner, subscribers
ApproverCreate	<ul style="list-style-type: none"> Release owner Explicitly defined for the component that manages the associated release 	New approver, subscribers
ApproverDelete	<ul style="list-style-type: none"> Release owner Explicitly defined for the component that manages the associated release 	Deleted approver, subscribers
BuilderCreate	<ul style="list-style-type: none"> Explicitly defined for the component that manages the associated release 	Subscribers
BuilderDelete	<ul style="list-style-type: none"> Explicitly defined for the component that manages the associated release 	Subscribers
BuilderExtract	<ul style="list-style-type: none"> Explicitly defined for the component that manages the associated release 	Not applicable
BuilderModify	<ul style="list-style-type: none"> Explicitly defined for the component that manages the associated release 	Subscribers
BuilderView	<ul style="list-style-type: none"> Explicitly defined for the component that manages the associated release 	Not applicable
CollisionAccept	<ul style="list-style-type: none"> Component owner Explicitly defined for the component that manages the associated release 	Release owner, subscribers
CollisionReconc	<ul style="list-style-type: none"> Component owner Explicitly defined for the component that manages the associated release 	Release owner, subscribers

For this action	These users have authority	These users are notified
CollisionReject	<ul style="list-style-type: none"> Component owner Explicitly defined for the component that manages the associated release 	Release owner, subscribers
CompCreate	<ul style="list-style-type: none"> Parent component owner Explicitly defined for the parent component 	New component owner
CompDelete	<ul style="list-style-type: none"> Component owner Explicitly defined for the component being removed 	Component owner, subscribers
CompLink	<ul style="list-style-type: none"> Component owner of the component being linked Explicitly defined for the component being linked 	Owners of both components, subscribers
CompModify	<ul style="list-style-type: none"> Component owner Explicitly defined for the component being modified 	New component owner if applicable, subscribers
CompRecreate	<ul style="list-style-type: none"> Parent component owner Explicitly defined for the parent component 	Owners of both components, subscribers
CompUnlink	<ul style="list-style-type: none"> Component owner of the component being unlinked Explicitly defined for the component being unlinked 	Owners of both components, subscribers
CompView	<ul style="list-style-type: none"> Component owner Explicitly defined for the component being viewed 	Not applicable
CoreqCreate	<ul style="list-style-type: none"> Work area owner of all specified work areas Explicitly defined for the component managing the associated work area and release 	Not applicable
CoreqDelete	<ul style="list-style-type: none"> Work area owner of all specified work areas Explicitly defined for the component associated with the release 	Not applicable
DefectAccept	<ul style="list-style-type: none"> Defect owner for the component associated with the defect 	Defect owner, defect originator, duplicate defect originators, subscribers
DefectAssign	<ul style="list-style-type: none"> Defect owner, defect originator Explicitly defined for the component associated with the defect <p>Note: Originators who do not have DefectAssign authority can reassign the defect only when it is in the open state.</p>	New owner, defect originator, duplicate defect originators, subscribers
DefectCancel	<ul style="list-style-type: none"> Defect originator Explicitly defined for the component associated with the defect 	Defect owner, defect originator, duplicate defect originators, subscribers
DefectClose	Automatic action; no authority is required	Defect owner, defect originator, duplicate defect originators, subscribers

For this action	These users have authority	These users are notified
DefectConfiginfo	Not applicable; this is a base authority that can be performed by all users in the family	Defect owner, defect originator, duplicate defect originators, subscribers
DefectDesign	<ul style="list-style-type: none"> • Defect owner • Explicitly defined for the component associated with the defect 	Defect owner, defect originator, duplicate defect originators, subscribers
DefectModify	<ul style="list-style-type: none"> • Defect owner can modify: <ul style="list-style-type: none"> – answer, abstract, environment, driver, prefix, reference, release, and all configurable fields • Defect originator can modify: <ul style="list-style-type: none"> – originator, severity, name, abstract, environment, driver, prefix, reference, release, and all configurable fields • Explicitly defined for the component associated with the defect, these users can modify: <ul style="list-style-type: none"> – abstract, answer, name, environment, driver, originator, prefix, reference, release, severity, phaseFound*, phaseInject*, priority*, symptom*, and target* <p>*If these fields have been configured by the family administrator, the field names might differ from those shown.</p>	Defect owner, defect originator, duplicate defect originators, subscribers
DefectOpen	Not applicable; this is a base authority that can be performed by all users in the family	Component owner, subscribers
DefectReopen	<ul style="list-style-type: none"> • Defect originator • Explicitly defined for the component associated with the defect 	Defect owner, defect originator, duplicate defect originators, subscribers
DefectReturn	<ul style="list-style-type: none"> • Defect owner • Explicitly defined for the component associated with the defect 	Defect originator, duplicate defect originators, subscribers
DefectReview	<ul style="list-style-type: none"> • Defect owner • Explicitly defined for the component associated with the defect 	Defect owner, defect originator, duplicate defect originators, subscribers
DefectSize	<ul style="list-style-type: none"> • Defect owner • Explicitly defined for the component associated with the defect 	Defect owner, defect originator, duplicate defect originators, subscribers
DefectVerify	<ul style="list-style-type: none"> • Defect owner • Explicitly defined for the component associated with the defect 	Defect owner, defect originator, duplicate defect originators, subscribers

For this action	These users have authority	These users are notified
DefectView	<ul style="list-style-type: none"> • Defect owner, defect originator • Explicitly defined for the component associated with the defect 	Not applicable
DriverAssign	<ul style="list-style-type: none"> • Driver owner • Explicitly defined for the component associated with the release 	New owner, subscribers
DriverCheck	<ul style="list-style-type: none"> • Driver owner • Explicitly defined for the component associated with the release 	Not applicable
DriverCommit	<ul style="list-style-type: none"> • Explicitly defined for the component associated with the release 	Subscribers
DriverComplete	<ul style="list-style-type: none"> • Explicitly defined for the component associated with the release 	Subscribers
DriverCreate	<ul style="list-style-type: none"> • Release owner • Explicitly defined for the component associated with the release 	Subscribers
DriverDelete	<ul style="list-style-type: none"> • Driver owner • Explicitly defined for the component associated with the release 	Subscribers
DriverExtract	<ul style="list-style-type: none"> • Driver owner • Explicitly defined for the component associated with the release 	Not applicable
DriverFreeze	<ul style="list-style-type: none"> • Driver owner • Explicitly defined for the component associated with the release 	Driver owner, subscribers
DriverModify	<ul style="list-style-type: none"> • Driver owner • Explicitly defined for the component associated with the release 	Driver owner, subscribers
DriverRefresh	<ul style="list-style-type: none"> • Explicitly defined for the component associated with the release 	Component owner, subscribers
DriverRestrict	<ul style="list-style-type: none"> • Driver owner • Explicitly defined for the component associated with the release 	Driver owner, subscribers
DriverView	<ul style="list-style-type: none"> • Driver owner • Explicitly defined for the component associated with the release 	Not applicable
EnvCreate	<ul style="list-style-type: none"> • Release owner • Explicitly defined for the component associated with the release 	Tester, subscribers

For this action	These users have authority	These users are notified
EnvDelete	<ul style="list-style-type: none"> Release owner Explicitly defined for the component associated with the release 	Subscribers
EnvModify	<ul style="list-style-type: none"> Release owner Explicitly defined for the component associated with the release 	Tester, subscribers
FeatureAccept	<ul style="list-style-type: none"> Feature owner Explicitly defined for the component associated with the feature 	Feature owner, feature originator, duplicate feature originators, subscribers
FeatureAssign	<ul style="list-style-type: none"> Feature owner Explicitly defined for the component associated with the feature 	New owner, feature originator, duplicate feature originators, subscribers
FeatureCancel	<ul style="list-style-type: none"> Feature originator Explicitly defined for the component associated with the feature 	Feature owner, feature originator, duplicate feature originators, subscribers
FeatureClose	Occurs automatically; no authority is required	Feature owner, feature originator, duplicate feature originators, subscribers
FeatureComment	Not applicable; this is a base authority that can be performed by all users in the family	Feature owner, feature originator, duplicate feature originators, subscribers
FeatureDesign	<ul style="list-style-type: none"> Feature owner Explicitly defined for the component associated with the feature 	Feature owner, feature originator, duplicate feature originators, subscribers
FeatureModify	<ul style="list-style-type: none"> Feature owner can modify: <ul style="list-style-type: none"> abstract, prefix, reference, and all configurable fields Feature originator can modify: <ul style="list-style-type: none"> abstract, name, prefix, reference, and all configurable fields Explicitly defined for the component associated with the feature, these users can modify: <ul style="list-style-type: none"> abstract, name, originator, prefix, reference, priority*, and target* <p>*If these fields have been configured by the family administrator, the field names might differ from those shown.</p>	Feature owner, feature originator, duplicate feature originators, subscribers

For this action	These users have authority	These users are notified
FeatureOpen	Not applicable; this is a base authority that can be performed by all users in the family	Component owner, subscribers
FeatureReopen	<ul style="list-style-type: none"> • Feature originator • Explicitly defined for the component associated with the feature 	Feature owner, feature originator, duplicate feature originators, subscribers
FeatureReturn	<ul style="list-style-type: none"> • Feature owner • Explicitly defined for the component associated with the feature 	Feature owner, feature originator, duplicate feature originators, subscribers
FeatureReview	<ul style="list-style-type: none"> • Feature owner • Explicitly defined for the component associated with the feature 	Feature owner, feature originator, duplicate feature originators, subscribers
FeatureSize	<ul style="list-style-type: none"> • Feature owner • Explicitly defined for the component associated with the feature 	Feature owner, feature originator, duplicate feature originators, subscribers
FeatureVerify	<ul style="list-style-type: none"> • Feature owner • Explicitly defined for the component associated with the feature 	Feature owner, feature originator, duplicate feature originators, subscribers
FeatureView	<ul style="list-style-type: none"> • Feature owner • Explicitly defined for the component associated with the feature 	Not applicable
FixActive	<ul style="list-style-type: none"> • Fix record owner, component owner, work area owner • Explicitly defined for the component associated with the fix record 	Subscribers
FixAssign	<ul style="list-style-type: none"> • Fix record owner, component owner, work area owner • Explicitly defined for the component associated with the fix record 	New fix record owner, subscribers
FixComplete	<ul style="list-style-type: none"> • Fix record owner, component owner, work area owner • Explicitly defined for the component associated with the fix record 	Subscribers
FixCreate	<ul style="list-style-type: none"> • Defect or feature owner, work area owner • Explicitly defined for the component associated with the defect or feature 	Subscribers
FixDelete	<ul style="list-style-type: none"> • Defect or feature owner, work area owner • Explicitly defined for the component associated with the defect or feature 	Subscribers

For this action	These users have authority	These users are notified
HostCreate	<ul style="list-style-type: none"> Owner of the user ID for which a host list entry is being created or deleted Superuser 	Not applicable
HostDelete	<ul style="list-style-type: none"> Owner of the user ID for which a host list entry is being deleted Superuser 	Not applicable
MemberCreate	<ul style="list-style-type: none"> Driver owner Explicitly defined for the component associated with the release 	Driver owner, subscribers
MemberCreateR	<ul style="list-style-type: none"> Driver owner Explicitly defined for the component associated with the release 	Driver owner, subscribers
MemberDelete	<ul style="list-style-type: none"> Driver owner Explicitly defined for the component associated with the release 	Driver owner, subscribers
MemberDeleteR	<ul style="list-style-type: none"> Driver owner Explicitly defined for the component associated with the release 	Driver owner, subscribers
NotifyCreate	<ul style="list-style-type: none"> Component owner Explicitly defined for the component associated with the notification list 	Not applicable
NotifyDelete	<ul style="list-style-type: none"> Component owner Owner of user ID Explicitly defined for the component associated with the notification list <p>Note: Users can delete themselves from a notification list without requiring any authority</p>	Not applicable
ParserCreate	<ul style="list-style-type: none"> Explicitly defined for the component associated with the release 	Subscribers
ParserDelete	<ul style="list-style-type: none"> Explicitly defined for the component associated with the release 	Subscribers
ParserModify	<ul style="list-style-type: none"> Explicitly defined for the component associated with the release 	Subscribers
ParserView	<ul style="list-style-type: none"> Explicitly defined for the component associated with the release 	Not applicable
PartAdd	<ul style="list-style-type: none"> Component owner Explicitly defined for the component associated with the part 	Subscribers

For this action	These users have authority	These users are notified
PartBuild	<ul style="list-style-type: none"> Component owner Explicitly defined for the component associated with the part 	Subscribers
PartCheckIn	<ul style="list-style-type: none"> User who checked out or locked the part originally, component owner Explicitly defined for the component associated with the part <p>Note: The user who is explicitly given this authority can check in a part that is checked out by someone else.</p>	Subscribers
PartCheckOut	<ul style="list-style-type: none"> Component owner Explicitly defined for the component associated with the part 	Subscribers
PartChildInfo	<ul style="list-style-type: none"> Component owner Explicitly defined for the component associated with the part 	Not applicable
PartConnect	<ul style="list-style-type: none"> Component owner Explicitly defined for the component associated with the part 	Subscribers
PartDelete	<ul style="list-style-type: none"> Component owner Explicitly defined for the component associated with the part 	Subscribers
PartDeleteForce	<ul style="list-style-type: none"> Component owner Explicitly defined for the component associated with the part 	Subscribers
PartExtract	<ul style="list-style-type: none"> Component owner Explicitly defined for the component associated with the part 	Not applicable
PartForceIn	<ul style="list-style-type: none"> Component owner Explicitly defined for the component associated with the part 	Subscribers
PartForceOut	<ul style="list-style-type: none"> Component owner Explicitly defined for the component associated with the part 	Subscribers
PartLink	<ul style="list-style-type: none"> Component owner Explicitly defined for the component associated with the part 	Subscribers

For this action	These users have authority	These users are notified
PartLock	<ul style="list-style-type: none"> • Component owner • Explicitly defined for the component associated with the part 	Subscribers
PartLockForce	<ul style="list-style-type: none"> • Component owner • Explicitly defined for the component associated with the part 	Subscribers
PartMark	<ul style="list-style-type: none"> • Component owner • Explicitly defined for the component associated with the part 	Subscribers
PartMerge	<ul style="list-style-type: none"> • Component owner • Explicitly defined for the component associated with the part 	Subscribers
PartModify	<ul style="list-style-type: none"> • Component owner • Explicitly defined for the component associated with the part 	Subscribers
PartOverrideR	Explicitly defined for the component associated with the release	Subscribers, user granted the override (if a user specified)
PartReconcile	<ul style="list-style-type: none"> • Component owner • Explicitly defined for the component that manages the associated release 	Subscribers
PartRecreateForce	<ul style="list-style-type: none"> • Component owner • Explicitly defined for the component associated with the part 	Subscribers
PartRecreate	<ul style="list-style-type: none"> • Component owner • Explicitly defined for the component associated with the part 	Subscribers
PartRefresh	<ul style="list-style-type: none"> • Component owner • Explicitly defined for the component associated with the part 	Subscribers
PartRename	<ul style="list-style-type: none"> • Component owner • Explicitly defined for the component associated with the part 	Subscribers
PartRenameForce	<ul style="list-style-type: none"> • Component owner • Explicitly defined for the component associated with the part 	Subscribers
PartResolve	Not applicable; this is a base authority that can be performed by all users in the family	Not applicable
PartRestrict	Explicitly defined for the component associated with the release	Subscribers

For this action	These users have authority	These users are notified
PartTouch	<ul style="list-style-type: none"> • Component owner • Explicitly defined for the component associated with the part 	Subscribers
PartUndo	<ul style="list-style-type: none"> • Component owner • Explicitly defined for the component associated with the part 	Subscribers
PartUndoForce	<ul style="list-style-type: none"> • Component owner • Explicitly defined for the component associated with the part 	Subscribers
PartUnlock	<ul style="list-style-type: none"> • User who checked out or locked the part originally, component owner • Explicitly defined for the component associated with the part 	Subscribers
PartView	<ul style="list-style-type: none"> • Component owner • Explicitly defined for the component associated with the part 	Not applicable
PartViewMsg	<ul style="list-style-type: none"> • Component owner • Explicitly defined for the component associated with the part 	Not applicable
PrereqCreate	<ul style="list-style-type: none"> • Work area owner of all specified work areas • Explicitly defined for the component managing the associated work area and release 	Not applicable
PrereqDelete	<ul style="list-style-type: none"> • Work area owner of all specified work areas • Explicitly defined for the component managing the associated work area and release 	Not applicable
ReleaseCreate	<ul style="list-style-type: none"> • Explicitly defined for the component associated with the new release 	New release owner, component owner, subscribers
ReleaseDelete	<ul style="list-style-type: none"> • Release owner • Explicitly defined for the component associated with the release 	Release owner, component owner, subscribers
ReleaseExtract	<ul style="list-style-type: none"> • Release owner • Explicitly defined for the component associated with the release 	Not applicable
ReleaseLink	<ul style="list-style-type: none"> • Release owner • Explicitly defined for the component associated with the release 	Release owner, subscribers
ReleaseMerge	<ul style="list-style-type: none"> • Release owner • Explicitly defined for the component associated with the release 	Release owner, subscribers

For this action	These users have authority	These users are notified
ReleaseModify	<ul style="list-style-type: none"> Release owner Explicitly defined for the component associated with the release <p>Note: To identify a new component to manage the release, you must have ReleaseCreate in an authority group in the component that you are modifying</p>	Release owner, subscribers, new owner (if applicable)
ReleasePrune	<ul style="list-style-type: none"> Release owner Explicitly defined for the component associated with the release 	Subscribers
ReleaseRecreate	<ul style="list-style-type: none"> Release owner Explicitly defined for the component associated with the release 	Release owner, component owner, subscribers
ReleaseView	<ul style="list-style-type: none"> Release owner Explicitly defined for the component associated with the release 	Not applicable
Report	Not applicable; this is a base authority that can be performed by all users in the family	Not applicable
ShadowCreate	Explicitly defined for the component associated with the release	Not applicable
ShadowDefine	Superuser	Not applicable
ShadowDelete	Explicitly defined for the component associated with the release	Not applicable
ShadowDisable	Explicitly defined for the component associated with the release	Not applicable
ShadowEnable	Explicitly defined for the component associated with the release	Not applicable
ShadowModify	Explicitly defined for the component associated with the release	Not applicable
ShadowRedefine	Superuser	Not applicable
ShadowSync	Explicitly defined for the component associated with the release	Not applicable
ShadowUndefine	Superuser	Not applicable
ShadowVerify	Explicitly defined for the component associated with the release	Not applicable
ShadowView	Explicitly defined for the component associated with the release	Not applicable
SizeAccept	<ul style="list-style-type: none"> Sizing record owner Explicitly defined for the component associated with the sizing record 	Subscribers

For this action	These users have authority	These users are notified
SizeAssign	<ul style="list-style-type: none"> • Sizing record owner • Explicitly defined for the component associated with the sizing record 	New sizing record owner, defect/feature owner, subscribers
SizeCreate	<ul style="list-style-type: none"> • Defect/feature owner • Explicitly defined for the component associated with the defect/feature 	Component owner, defect/feature owner, subscribers
SizeDelete	<ul style="list-style-type: none"> • Defect/feature owner • Explicitly defined for the component associated with the defect/feature 	Subscribers, sizing record owner, defect/feature owner
SizeReject	<ul style="list-style-type: none"> • Sizing record owner • Explicitly defined for the component associated with the sizing record 	Subscribers
TestAbstain	<ul style="list-style-type: none"> • Test record owner • Explicitly defined for the component associated with the test record's release 	Subscribers
TestAccept	<ul style="list-style-type: none"> • Test record owner • Explicitly defined for the component associated with the test record's release 	Subscribers
TestAssign	<ul style="list-style-type: none"> • Test record owner • Explicitly defined for the component associated with the test record's release 	New test record owner, subscribers
TestReady	<ul style="list-style-type: none"> • Test record owner • Explicitly defined for the component associated with the test record's release 	Subscribers
TestReject	<ul style="list-style-type: none"> • Test record owner • Explicitly defined for the component associated with the test record's release 	Subscribers
UserCreate	Superuser	New user
UserDelete	Superuser	Not applicable
UserModify	<ul style="list-style-type: none"> • Owner of the user object can modify all characteristics except the superuser privilege • Must be a superuser to grant the superuser privilege 	Not applicable
UserRecreate	Superuser	Not applicable
UserView	Not applicable; this is a base authority that can be performed by all users in the family	Not applicable
VerifyAbstain	<ul style="list-style-type: none"> • Verification record owner • Explicitly defined for the component associated with the verification record's defect or feature 	Subscribers

For this action	These users have authority	These users are notified
VerifyAccept	<ul style="list-style-type: none"> • Verification record owner • Explicitly defined for the component associated with the verification record's defect or feature 	Subscribers
VerifyAssign	<ul style="list-style-type: none"> • Verification record owner • Explicitly defined for the component associated with the verification record's defect or feature 	New verification record owner, subscribers
VerifyReady	Takes place automatically; no authority is required	Verification record owners
VerifyReject	<ul style="list-style-type: none"> • Verification record owner • Explicitly defined for the component associated with the verification record's defect or feature 	Subscribers
WorkAreaAssign	<ul style="list-style-type: none"> • Work area owner • Explicitly defined for the component associated with the release 	New work area owner, subscribers
WorkAreaCancel	<ul style="list-style-type: none"> • Defect or feature owner • Explicitly defined for the component associated with the defect or feature 	Subscribers, owners of approval records for work area being canceled
WorkAreaCheck	<ul style="list-style-type: none"> • Work area owner • Explicitly defined for the component associated with the release 	Not applicable
WorkAreaCommit	<ul style="list-style-type: none"> • Work area owner • Explicitly defined for the component associated with the release 	Subscribers
WorkAreaCompleat	<ul style="list-style-type: none"> • Work area owner • Explicitly defined for the component associated with the release 	Subscribers
WorkAreaCreate	<ul style="list-style-type: none"> • Defect or feature owner • Explicitly defined for the component associated with the defect or feature 	Work area owner, subscribers
WorkAreaFix	<ul style="list-style-type: none"> • Work area owner • Explicitly defined for the component associated with the release 	Subscribers
WorkAreaFreeze	<ul style="list-style-type: none"> • Work area owner • Explicitly defined for the component associated with the release 	Subscribers
WorkAreaIntegra	<ul style="list-style-type: none"> • Work area owner • Explicitly defined for the component associated with the release 	Subscribers

For this action	These users have authority	These users are notified
WorkAreaModify	<ul style="list-style-type: none"> • Work area owner • Explicitly defined for the component associated with the release 	Subscribers
WorkAreaReconcile	<ul style="list-style-type: none"> • Work area owner • Explicitly defined for the component associated with the release 	Work area owner, subscribers
WorkAreaRefresh	<ul style="list-style-type: none"> • Work area owner • Explicitly defined for the component associated with the release 	Work area owner, subscribers
WorkAreaTest	<ul style="list-style-type: none"> • Work area owner • Explicitly defined for the component associated with the release 	Subscribers
WorkAreaUndo	<ul style="list-style-type: none"> • Work area owner • Explicitly defined for the component associated with the release 	
WorkAreaView	<ul style="list-style-type: none"> • Work area owner • Explicitly defined for the component associated with the release 	Not applicable

Appendix J. Sample REXX execs, build scripts, and parsers

This appendix is composed of the IBM-supplied REXX execs, build scripts, and parsers. Your family administrator can modify these samples to fit the needs of your organization.

The samples in this appendix may not be available on all platforms. Refer to the readme file for a complete list of samples available with TeamConnection. All samples are provided as-is and any use of or modifications to the samples are the sole responsibility of the customer.

Sample REXX execs

This section lists the sample REXX execs that are shipped with TeamConnection. The client.smp file contains this same listing. It is located in the bin subdirectory of the directory where the TeamConnection client is installed.

Users running these execs must have user and host access to your TeamConnection family.

Most of the execs require input parameters, and some require that the TC_FAMILY or TC_RELEASE environment variables be set. If the user who is running the script is acting for another user, the TC_BECOME environment variable must also be set. These variables can be set from a command line prompt.

The following convention is used to show the required, optional, and selective input parameters:

- Brackets ([]) indicate that the input or variable is optional.
- Braces ({ }) indicate that one of the inputs is required.
- An input or variable that is not surrounded by brackets or braces is required.

Script name	Function	Inputs	Environment variables
accComp	Lists the explicit access of users in a specified component.	componentName	TC_FAMILY [TC_BECOME]
compChld	Lists the direct children of a specified component. Also lists the description and owner of each child component.	componentName	TC_FAMILY [TC_BECOME]
compOwnr	Displays a list of component owners' addresses in a specified family.	familyName	[TC_BECOME]
compPrnt	Lists the parent component of a specified component.	componentName	TC_FAMILY [TC_BECOME]
compWalk	Displays the children and grandchildren of a specified component.	componentName	TC_FAMILY [TC_BECOME]
defClone	Creates a new defect based on values contained in a specified defect.	defectNumber	TC_FAMILY [TC_BECOME]

Script name	Function	Inputs	Environment variables
defDvr	Lists all defects for a specified driver.	driverName	TC_FAMILY [TC_BECOME]
defFfea	Creates a new defect based on values contained in a specified feature.	featureNumber	TC_FAMILY [TC_BECOME]
defNew	Displays the number of the most recent defect that was entered in the system.		TC_FAMILY [TC_BECOME]
defReopn	Reopens a previously canceled or returned defect.	defectNumber	TC_FAMILY [TC_BECOME]
defRept	Generates a global defect report showing work areas, test records, approval records, and fix records.		TC_FAMILY [TC_BECOME]
defState	Lists the defect number of all defects that are in a specified state.	stateName	TC_FAMILY [TC_BECOME]
defStats	Displays total active defect statistics on a defect owner area basis.	ownerArea	TC_FAMILY [TC_BECOME]
defWRef	Displays the full details of defects that contain the specified reference field value.	reference	TC_FAMILY [TC_BECOME]
dfDesc	Displays the full remarks that were entered when the specified defect or feature was created.	{defectNumber featureNumber}	TC_FAMILY [TC_BECOME]
feaClone	Creates a new feature based on values contained in a specified feature.	featureNumber	TC_FAMILY [TC_BECOME]
feaDvr	Lists all features contained in a specified driver.	driverName	TC_FAMILY [TC_BECOME]
feaFdef	Creates a new feature based on the values contained in the specified defect.	defectNumber	TC_FAMILY [TC_BECOME]
feaNew	Displays the number of the most recent feature that was entered in the system.		TC_FAMILY [TC_BECOME]
feaReopn	Reopens a previously canceled or returned feature.	featureNumber	TC_FAMILY [TC_BECOME]
feaRept	Generates a global feature report showing work areas, test records, size records, and fix records.		TC_FAMILY [TC_BECOME]
feaState	Lists the feature number of all features that are in a specified state.	stateName	TC_FAMILY [TC_BECOME]
feaStats	Displays total active feature statistics on a feature owner area basis.	ownerArea	TC_FAMILY [TC_BECOME]
drvByDF	Lists the name of the drivers that contain a specified defect or feature.	{defectNumber featureNumber}	TC_FAMILY [TC_BECOME]
drvMem	Lists the defect and feature members of a specified driver for a specified release.	driverName [releaseName]	TC_FAMILY [TC_RELEASE] [TC_BECOME]
mailTo	Sends a message to the addresses read through stdin.	messagefile subject	
ownerChg	Re-assigns all current work and objects owned by userLogin1 to userLogin2.	userLogin1 userLogin2	TC_FAMILY [TC_BECOME]

Script name	Function	Inputs	Environment variables
prtChckin	Checks parts into the TeamConnection family. When common parts are encountered, the script requests the releases for which the part should remain common.	partPathName [releaseName]	TC_FAMILY [TC_RELEASE] [TC_BECOME]
prtChgDf	Lists all the parts that were changed for a specified defect or feature.	{defectNumber featureNumber}	TC_FAMILY [TC_BECOME]
prtChgDr	Lists the parts that were changed for a specified driver.	driverName	TC_FAMILY [TC_BECOME]
prtComGt	Extracts all the parts associated with a specific component. The parts are placed in a directory that represents the release name to which the version of the part is associated. This directory is created relative to the relativePathName parameter.	componentName relativePathName [committed]	TC_FAMILY [TC_BECOME]
prtComp	Lists all parts related to a specified component.	componentName	TC_FAMILY [TC_BECOME]
prtHist	Lists all defect and feature numbers and abstracts that caused a change to a specified part in a specified release.	partName [releaseName]	TC_FAMILY [TC_RELEASE] [TC_BECOME]
prtInfo	Displays information for a specified part.	partName	TC_FAMILY [TC_RELEASE] [TC_BECOME]
prtLock	Lists all parts that a specified user has locked.	userLogin	TC_FAMILY [TC_BECOME]
prtLokBy	Lists who has a specified part checked out.	partName	TC_FAMILY TC_RELEASE [TC_BECOME]
PrtPath	Finds and lists all parts that match a partial part path name.	partPathName	TC_FAMILY [TC_BECOME]
prtRel	Lists all parts related to a specified release.	releaseName	TC_FAMILY [TC_BECOME]
prtWaGt	Extracts all the parts associated with a specific work area and places them in the path specified by the relativePathName parameter.	releaseName workareaName relativePathName	TC_FAMILY [TC_BECOME]
rByArea	Generates a manager's report based on the specified areas or departments of interest.	areaName ...	TC_FAMILY [TC_BECOME]
relOwner	Displays a list of addresses of all release owners in a specified TeamConnection family.	familyName	[TC_BECOME]
showConf	Lists the valid values pertaining to a specified configurable type.	configType	TC_FAMILY [TC_BECOME]
userAuth	Lists the users who have the authority to give other users access to a specified component.	componentName	TC_FAMILY [TC_BECOME]
userInfo	Finds user information based on part of the user's name. A fuzzy search is performed.	userName	TC_FAMILY [TC_BECOME]
usersAll	Lists the addresses of all users in a specified TeamConnection family.	familyName	[TC_BECOME]

Script name	Function	Inputs	Environment variables
usrAcc	Lists the explicit access of a specified user for the specified component and its descendant components.	userLogin componentName	TC_FAMILY [TC_BECOME]
usrRept	Generates a user's report based on the specified user login.	userLogin	TC_FAMILY [TC_BECOME]
verByPrt	Lists the version numbers, release names, and path names for the specified part.	partName releaseName	TC_FAMILY [TC_BECOME]
waComit	Lists the work areas that are in the commit state for a specified release.	releaseName	TC_FAMILY [TC_BECOME]
waFix	Lists all the work areas that are in the fix state for a given release.	releaseName	TC_FAMILY [TC_BECOME]
walnLvl	Lists the work areas that are in the integrate state and are associated with at least one development driver for the specified release.	releaseName	TC_FAMILY [TC_BECOME]
walnt	Lists the work areas that are in the integrate state for a specified release.	releaseName	TC_FAMILY [TC_BECOME]
waPrdLv	Lists the work areas that are included in a production driver and are in the integrate state for a specified release.	releaseName	TC_FAMILY [TC_BECOME]
waStat	Generates a work area activity statistics report on a user area basis.	userArea	TC_FAMILY [TC_BECOME]
waTest	Lists the work areas that are in the test state for a specified release.	releaseName	TC_FAMILY [TC_BECOME]

Sample build scripts

fhbcob2.cmd

Calls the COBOL Visual Set for OS/2 compiler.

fhbcob2l.cmd

Calls the COBOL Visual Set for OS/2 compiler and link editor.

fhbocomp.cmd

Calls the VisualAge for C++ icc compiler.

fhbolib.cmd

Calls the OS/2 implib utility.

fhbolin2.cmd

Calls the VisualAge for C++ icc link editor.

fhbolink.cmd

Calls the link386 link editor.

fhborc.cmd

Calls the OS/2 resource compiler.

fhbplbld.cmd

Calls the OS/2 PL/1 compiler.

fhbplink.cmd
Calls the OS/2 PL/1 link editor.

edcc.jcl
Calls the C/370 JCL procedure.

fhbcobm.jcl
Calls the COBOL for MVS compiler.

fhbm370.jcl
Calls the C/370 compiler.

fhbmasm.jcl
Calls the MVS assembler.

fhbmc.jcl
Calls the C/370 compiler.

fhbmlink
Calls the MVS linkage editor.

fhbmpli.jcl
Calls the PL/1 MVS compiler.

fhbplked.jcl
Calls the C370 prelinker.

fhbtclnk
Calls the TeamConnection pseudo linker.

fhbwcomp.c
Calls the Microsoft Visual C++ compiler

fhbwlink.c
Calls the Microsoft linker

gather.cmd
Calls the Gather tool.

Sample parsers

fhbcbprs.cmd
A parser for COBOL applications.

fhbopars.cmd
A parser for C applications.

fhbcpp.c func.c fhbcpp.h
A C parser for COBOL applications.

mvsasmp.c mvsasmp.h
A parser for MVS assembly language applications.

fhbplprs.cmd
A parser for PL/1 applications.

Sample package files

gather.pkf

A package file for the Gather tool.

softdist.pkf

A package file for the Tivoli Software Distribution tool.

Service and Support

VisualAge TeamConnection Services!

VisualAge TeamConnection services offerings will provide customers with the tools to quickly establish a more productive and efficient development environment. These services offerings focus on the LAN library component of VisualAge TeamConnection and on the Repository function of VisualAge TeamConnection.

If you are interested in the VisualAge TeamConnection Services, select the **Support** item at:

<http://www.software.ibm.com/software/ad/teamcon>

VisualAge TeamConnection Support!

If you have a question or a problem, please take a moment to review the Customer Support section from any of the manuals for the VisualAge TeamConnection product. Your options for VisualAge TeamConnection support, as described in your License Information, include the following information (subject to availability).

IBM Lotus Passport Advantage Program

For more information on the IBM Lotus Passport Advantage volume licensing program that provides customers with a series of contract offerings under which they can acquire licenses, software subscriptions, and support, go to:

<http://www.lotus.com/passportadvantage>

DB2 Service Maintenance and Technical Library

To download the latest service maintenance for DB2, use the DB2 Service and Support on the World Wide Web at:

<http://www.software.ibm.com/data/db2/db2tech>

Note: Even though DB2 is bundled with VisualAge TeamConnection you should contact VisualAge TeamConnection Support to report DB2 problems. The licensing for VisualAge TeamConnection does not entitle you to contact DB2 Support directly. For a complete and up-to-date source of DB2 information, use the DB2 Product and Service Technical Library, in English only, on the World Wide Web at:

<http://www.software.ibm.com/data/db2/library>

For North American Customers

Electronic Forums

So you may electronically access VisualAge TeamConnection technical information, exchange messages with other VisualAge TeamConnection users, and receive information regarding the availability of FixPaks.

IBM TalkLink

Use the TEAMC CFORUM. In the United States, call 1-800-547-1283. For TalkLink information available via the Internet, go to:

<http://www.ibmmlink.ibm.com/talklink>

CompuServe

From any ! prompt, type GO SOFSOL, then select TeamConnection. Refer to the enclosed CompuServe brochure for additional information, or call 1-800-848-8199. For CompuServe information available via the Internet, refer to:

<http://www.compuserve.com>

Internet

Go to the IBM homepage, <http://www.ibm.com>. Use the search function with keyword TeamConnection to go to the VisualAge TeamConnection area. Access the TeamConnection directory in our ftp site. Use ftp and login as anonymous to <ftp://software.ibm.com>. In the directory [ps/products/teamconnection](ftp://software.ibm.com/ps/products/teamconnection) you can find fixes and information related to VisualAge TeamConnection.

Telephone Support

Direct customer support is provided by the Personal Systems Support Line and by the AIX Support Line. These fee services enhance customers' productivity by providing voice and electronic access to the IBM support organization. They will help answer questions pertaining to usage, "how to," and suspected software defects for eligible products.

The following are phone numbers for software support in the US:

- Personal Systems Support Line: 1-800-237-5511
- AIX Support Line: 1-800-CALL-AIX (1-800-225-5249)

In Canada, call 1-800-IBM-SERV (1-800-426-7378).

Note: In the US, 1-800-237-5511 is the Software Support phone number for all IBM software (390, OS/400, AIX, Personal Systems, etc.). You may call this number and take the option for OS/2 - DOS support, which then transfers you to 1-800-992-4777 for the Personal Systems (workstation) products, or you may call 1-800-992-4777 directly.

Obtaining product information packages:

- United States: 1-800-IBM-CALL (1-800-426-2255)
- Canada: 1-800-IBM-CALL (1-800-426-2255)

Ordering TeamConnection products:

- United States: 1-800-IBM-CALL (1-800-426-2255)
- Canada: 1-800-IBM-CALL (1-800-426-2255)

TeamConnection education:

- United States: 1-800-IBM-TEACH (1-800-426-8322) Canada 1-800-IBM-TEACH (1-800-426-8322)

While we may not be able to respond to or resolve all problems and questions, your satisfaction with our products and support is important to us. If you cannot access these forums, contact your IBM representative. There are several other support offerings available after purchasing IBM VisualAge TeamConnection.

If you live within the U.S.A., call any of the following numbers:

- 1-800-237-5511 to learn about available service options
- 1-800-IBM-CALL (1-800-426-2255) to order products or get general information
- 1-800-879-2755 to order publications.

Support for Customers Outside North America

For information on how to contact IBM outside of the United States, see Appendix A of the *IBM Software Support Handbook*, which can be located by selecting the **Service Offering** item at:

<http://www.ibm.com/support>:

Note: In some countries, IBM-authorized dealers should contact their dealer support structure instead of the IBM Support Center.

Bibliography

IBM VisualAge TeamConnection Enterprise Server library

The following is a list of the TeamConnection publications. For a list of other publications about TeamConnection, including white papers, technical reports, a product fact sheet, and the product announcement letter, refer to the IBM VisualAge TeamConnection Enterprise Server Library home page. To access this home page, select **Library** from the IBM VisualAge TeamConnection Enterprise Server home page at URL <http://www.software.ibm.com/ad/teamcon>.

- **License Information (GC34-4497):**
Contains license, service, and warranty information.
- **Installation Guide (GC34-4742):**
Lists the hardware and software that are required before you can install and use the IBM VisualAge TeamConnection Enterprise Server product, provides detailed instructions for installing the TeamConnection server and client.
- **Administrator's Guide (GC34-4551):**
Provides instructions for configuring the TeamConnection family server and administering a TeamConnection family.
- **Getting Started with the TeamConnection Clients (SC34-4552):**
Tells first-time users how to install the TeamConnection clients on their workstations, and familiarizes them with the command line and graphical user interfaces.
- **User's Guide (SC34-4499):**
A comprehensive guide for TeamConnection administrators and client users that helps them install and use TeamConnection.
- **Commands Reference (SC34-4501):**
Describes the TeamConnection commands, their syntax, and the authority required to issue each command. This book also provides examples of how to use the various commands.
- **Quick Commands Reference (GC34-4500):**
Lists the TeamConnection commands along with their syntax.
- **Staying on Track with TeamConnection Processes (83H9677):**
Poster showing how objects flow through the states defined for each TeamConnection process.
- The following publications can be ordered as a set (SBOF-8560):
 - Administrator's Guide**
 - Getting Started with the TeamConnection Clients**
 - User's Guide**
 - Commands Reference**
 - Quick Commands Reference**
 - Staying on Track with TeamConnection Processes**

TeamConnection technical reports

The following is a list of technical reports available for TeamConnection. Refer to the IBM VisualAge TeamConnection Enterprise Server Library home page for the most up-to-date list of technical reports. To access this home page, select **Library** from the IBM VisualAge TeamConnection Enterprise Server home page at URL <http://www.software.ibm.com/ad/teamcon>.

29.2147	SCLM Guide to TeamConnection Terminology
29.2196	Using REXX command files with TeamConnection MVS Build Scripts
29.2231	TeamConnection Interoperability with MVS and SCLM
29.2235	Using REXX command files with TeamConnection MVS Build Scripts for PL/I programs
29.2253	Comparison between CMVC 2.3 and TeamConnection 2
29.2254	Migrating from CMVC 2.3 to TeamConnection 2
29.2267	TeamConnection frequently asked questions: how to do routine operating system tasks

DB2

The following publications are part of the IBM DB2 Universal Database library of documents for DB2 administration. DB2 publications are available in HTML format from the DB2 Product and Service Technical Library at the following URL:

<http://www.software.ibm.com/data/db2/library/>

- *Administration Getting Started* (S10J-8154-00)
An introductory guide to basic administration tasks and the DB2 administration tools.
- *SQL Getting Started* (S10J-8156-00)
Discusses basic concepts of DB2 SQL.
- *Administration Guide* (S10J-8157-00)
A complete guide to administration tasks and the DB2 administration tools.
- *SQL Reference* (S10J-8165-00)
A reference to DB2 SQL for programmers and database administrators.
- *Troubleshooting Guide* (S10J-8169-00)
A guide to identifying and solving problems with DB2 servers and clients and to using the DB2 diagnostic tools.
- *Messages Reference* (S10J-8168-00)
Provides detailed information about DB2 messages.
- *Command Reference* (S10J-8166-00)
Provides information about DB2 system commands and the command line processor.
- *Replication Guide* (S10J-0999-00)
Describes how to plan, configure, administer, and operate IBM replication tools available with DB2.
- *System Monitor Guide and Reference* (S10J-8164-00)

Describes how to monitor DB2 database activity and analyze system performance.

- *Glossary*
A comprehensive glossary of DB2 terms.

Related publications

- Transmission Control Protocol/Internet Protocol (TCP/IP)
 - *TCP/IP 2.0 for OS/2: Installation and Administration* (SC31-6075)
 - *TCP/IP for MVS Planning and Customization* (SC31-6085)
- MVS
 - *MVS/XA JCL User's Guide* (GC28-1351)
 - *MVS/XA JCL Reference* (GC28-1352)
 - *MVS/ESA JCL User's Guide* (GC28-1830)
 - *MVS/ESA JCL Reference* (GC28-1829)
- NLS and DBCS
 - *AIX 4, General Programming Concepts: Writing and Debugging Programs*. (SC23-2533-02). See chapter 16 "National Language Support" for an updated contents of the AIX 3 material (see below).
 - *AIX 4, System Management Guide: Operating System and Devices* (SC23-2525-03). See chapter 10, "National Language Support" for system tasks.
 - *AIX Version 3.2 for RISC System/6000, National Language Support* (GG24-3850).
 - *Internationalization of AIX Software, A Programmer's Guide* (SC23-2431).
 - *National Language Design Guide Volume 1* (SE09-8001-02). This manual contains very good information on how to enable an application for NLS.
 - *National Language Design Guide Volume 2* (SE09-8002-02). This manual provides information on the IBM language codes (consult the "Language codes" chapter).

Glossary

This glossary includes terms and definitions from the *IBM Dictionary of Computing*, 10th edition (New York: McGraw-Hill, 1993). If you do not find the term you are looking for, refer to this document's index or to the *IBM Dictionary of Computing*.

This glossary uses the following cross-references:

Compare to

Indicates a term or terms that have a similar but not identical meaning.

Contrast with

Indicates a term or terms that have an opposed or substantially different meaning.

See also

Refers to a term whose meaning bears a relationship to the current term.

A

absolute path name. A directory or a part expressed as a sequence of directories followed by a part name beginning from the root directory.

access list. A set of objects that controls access to data. Each object consists of a component, a user, and the authority that the user is granted or is restricted from in that component. See also *authority*, *granted authority*, and *restricted authority*.

action. A task performed by the TeamConnection server and requested by a TeamConnection client. A TeamConnection action is the same as issuing one TeamConnection command.

agent. See *build agent*.

alternate version ID. In collision records, the database ID of the version of a driver, release, or work area where the conflicting version of a part is visible.

approval record. A status record on which an approver must give an opinion of the proposed part changes required to resolve a defect or implement a feature in a release.

approver. A user who has the authority to mark an approval record with accept, reject, or abstain within a specific release.

approver list. A list of user IDs attached to a release, representing the users who must review part changes that are required to resolve a defect or implement a feature in that release.

attribute. Information contained in a field that is accessible to the user. TeamConnection enables family administrators to customize defect, feature, user, and part tables by adding new attributes.

authority. The right to access development objects and perform TeamConnection commands. See also *access list*, *base authority*, *explicit authority*, *granted authority*, *implicit authority*, *restricted authority*, and *superuser privilege*.

B

base authority. The set of actions granted to a user when a user ID is created within a TeamConnection family. See also *authority*. Contrast with *implicit authority* and *explicit authority*.

base name. The name assigned to the part outside of the TeamConnection server environment, excluding any directory names. See also *path name*.

base part tree. The base set of parts associated with a release, to which changes are applied over time. Each committed driver or work area for a release updates the base part tree for that release.

build. The process used to create applications within TeamConnection.

build associate. A TeamConnection part that is not an input to or an output from a build. An example of such a part is a read.me file.

build cache. A directory that the build processor uses to enhance performance.

build dependent. A TeamConnection part that is needed for the compile operation to complete, but it will not be passed directly to the compiler. An example of this is an include file. See also *dependencies*.

builder. An object that can transform one set of TeamConnection parts into another by invoking tools such as compilers and linkers.

build event. An individual step in the build of an application, such as the compiling of hello.c into hello.obj.

build input. A TeamConnection part that will be used as input to the object being built.

build output. A TeamConnection part that will be generated output from a build, such as an .obj or .exe file.

build pool. A group of build servers that resides in an environment. The environment in which several build servers operate. Typically, several servers are set up for each environment that the enterprise develops applications for.

build scope. A collection of build events that implement a specific build request. See also *build event*.

build script. An executable or command file that specifies the steps that should occur during a build operation. This file can be a compiler, a linker, or the name of a .cmd file you have written.

build server. A program that invokes the tools, such as compilers and linkers, that construct an application.

build target. The name of the part at the top of the build tree which is the final output of a build.

TeamConnection uses the build target to determine the scope of the build. See also *build tree*.

build tree. A graphical representation of the dependencies that the parts in an application have on one another. If you change the relationship of one part to another, the build tree changes accordingly.

C

change control process. The process of limiting and auditing changes to parts through the mechanism of checking parts in and out of a central, controlled, storage location. Change control for individual releases can be integrated with problem tracking by specifying a process for the release that includes the tracking subprocess.

check in. The return of a TeamConnection part to version control.

check out. The retrieval of a version of a part under TeamConnection control. In non-concurrent releases, the check out operation does not allow a second user to check out a part until the first user has checked it back in.

child component. Any component in a TeamConnection family, except the root component, that is created in reference to an existing component. The existing component is the parent component, and the new component is the child component. A parent component can have more than one child component, and a child component can have more than one parent component. See also *component* and *parent component*.

child part. Any part in a build tree that has a parent defined. A child part can be input, output, or dependent. See also *part* and *parent part*.

client. A functional unit that receives shared services from a server. Contrast with *server*.

collision record. A status record associated with a work area or driver, a part, and one of the following:

- The work area or driver's release
- Another work area

TeamConnection generates a collision record when a user attempts to replace an older version of a part with a modified version, another user has already modified that part, and the first user's modification is not based on this latest version of the part.

command. A request to perform an operation or run a program from the command line interface. In TeamConnection, a command consists of the command name, one action flag, and zero or more attribute flags.

command line. (1) An area on the Tasks window or in the TeamConnection Commands window where a user can type TeamConnection commands. (2) An area on an operating system window where you can type TeamConnection commands.

committed version. The revision of a part that is visible from the release.

common part. A part that is shared by two or more releases, and the same version of the part is the current version for those releases.

comparison operator. An operator used in comparison expressions. Comparison operators used in TeamConnection are > (greater than), < (less than), >= (greater than or equal to), <= (less than or equal to), = (equal to), and <> (different from).

component. A TeamConnection object that organizes project data into structured groups, and controls configuration management properties. Component owners can control access to data and notification of TeamConnection actions. Components exist in a parent-child hierarchy, with descendant components inheriting access and notification information from ancestor components. See also *access list* and *notification list*.

concurrent development. Several users can work on the same part at the same time. TeamConnection requires these users to reconcile

their changes when they commit or integrate their work areas and drivers with the release. Contrast with *serial development*. See also *work area*.

configuration management. The process of identifying, managing, and controlling software modules as they change over time.

connecting parts. The process of linking parts so that they are included in a build.

context. The current work area or driver used for part operations.

corequisite work areas. Two or more work areas designated as corequisites by a user so that all work areas in the corequisite group must be included as members in the same driver, before that driver can be committed. If the driver process is not used in the release, then all corequisite work areas must be integrated by the same command. See also *prerequisite work areas*.

current version. The last visible modification of a part in a driver, release, or work area.

current working directory. (1) The directory that is the starting point for relative path names. (2) The directory in which you are working.

D

daemon. A program that runs unattended to perform a standard service. Some daemons are triggered automatically to perform their task; others operate periodically.

database. A collection of data that can be accessed and operated upon by a data processing system for a specific purpose.

default. A value that is used when an alternative is not specified by the user.

default query. A database search, defined for a specific TeamConnection window, that is issued each time that TeamConnection window is opened. See also *search*.

defect. A TeamConnection object used to formally report a problem. The user who opens a defect is the defect originator.

delete. If you delete a development object, such as a part or a user ID, any reference to that object is removed from TeamConnection. Certain objects can be deleted only if certain criteria are met. Most objects that are deleted can be re-created.

delta part tree. A directory structure representing only the parts that were changed in a specified place.

dependencies. In TeamConnection builds there are two types of dependencies:

- **automatic.** These are build dependencies that a parser identifies.
- **manual.** These are build dependencies that a user explicitly identifies in a build tree.

See also *build dependent*.

descendant. If you descendant a development object, such as, a part or a user ID, any reference to that object is removed from TeamConnection. Certain objects can be descendant only if certain criteria are met. Most objects that are descendants can be re-created.

disconnecting parts. The process of unlinking parts so that they are not included in a build.

driver. A collection of work areas that represent a set of changed parts within a release. Drivers are only associated with releases whose processes include the track and driver subprocesses.

driver member. A work area that is added to a driver.

E

end user. See *user*.

environment. (1) A user-defined testing domain for a particular release. (2) A defect field, in which

case it is the environment where the problem occurred. (3) The string that matches a build server with a build event.

environment list. A TeamConnection object used to specify environments in which a release should be tested. A list of environment-user ID pairs attached to a release, representing the user responsible for testing each environment. Only one tester can be identified for an environment.

explicit authority. The ability to perform an action against a TeamConnection object because you have been granted the authority to perform that action. Contrast with *base authority* and *implicit authority*.

extract. A TeamConnection action you can perform on a builder, part, driver or release builder. An extraction results in copying the specified builder, part, or parts contained in the driver or release to a client workstation.

F

family. A logical organization of related data. A single TeamConnection server can support multiple families. The data in one family cannot be accessed from another family.

family administrator. A user who is responsible for all nonsystem-related tasks for one or more TeamConnection families, such as planning, configuring, and maintaining the TeamConnection environment and managing user access to those families.

family server. A workstation running the TeamConnection server software.

FAT. See *file allocation table*.

feature. A TeamConnection object used to formally request and record information about a functional addition or enhancement. The user who opens a feature is the feature originator.

file. A collection of data that is stored by the TeamConnection server and retrieved by a path name. Any text or binary file used in a

development project can be created as a TeamConnection file. Examples include source code, executable programs, documentation, and test cases.

file allocation table (FAT). The DOS-, OS/2-, Windows 95-, and Windows NT-compatible file system that manages input, output, and storage of files on your system. File names can be up to 8 characters long, followed by a file extension that can be up to 3 characters.

fix record. A status record that is associated with a work area and that is used to monitor the phases of change within each component that is affected by a defect or feature for a specific release.

freeze. The freeze action saves changed parts to the work area. Thus, TeamConnection takes a snapshot of the work area, including all of the current versions of parts visible from that work area, and saves this image of the system. The user can always come back to this stage of development in the work area. Note, however, that a freeze action does not make the changes visible to the other people working in the release. Compare with *refresh*.

full part tree. A directory structure representing a complete set of active parts associated with the release.

G

Gather. A tool to organize files for distribution into a specified directory structure. This tool can be used as a prelude to further distribution, such as using CD-ROM or through electronic means like NetView DM/2. It can also be used by itself for distributing file copies to network-attached file systems.

GID. A number which uniquely identifies a file's group to a UNIX system.

granted authority. If an authority is granted on an access list, then it applies for all objects managed by this component and any of its descendants for which the authority is not

restricted. See also *access list*, *authority*, and *inheritance*. Contrast with *restricted authority*.

graphical user interface (GUI). A type of computer interface consisting of a visual metaphor of a real-world scene, often as a desktop. Within that scene are icons, representing actual objects, that the user can access and manipulate with a pointing device.

GUI. Graphical user interface.

H

high-performance file system (HPFS). In the OS/2 operating system, an installable file system that uses high-speed buffer storage, known as a cache, to provide fast access to large disk volumes. The file system also supports the existence of multiple, active file systems on a single personal computer, with the capacity of multiple and different storage devices. File names used with HPFS can have as many as 254 characters.

host. A host node, host computer, or host system.

host list. A list associated with each TeamConnection user ID that indicates the client machine that can access the family server and act on behalf of the user. The family server uses the list to authenticate the identity of a client machine when the family server receives a command. Each entry consists of a login, a host name, and a TeamConnection user ID.

host name. The identifier associated with the host computer.

HPFS. See *high-performance file system*.

I

implicit authority. The ability to perform an action on a TeamConnection object without being granted explicit authority. This authority is

automatically granted through inheritance or object ownership. Contrast with *base authority* and *explicit authority*.

import. To bring in data. In TeamConnection, to bring selected items into a field from a matching TeamConnection object window.

inheritance. The passing of configuration management properties from parent to child component. The configuration management properties that are inherited are access and notification. Inheritance within each TeamConnection family or component hierarchy is cumulative.

integrated problem tracking. The process of integrating problem tracking with change control to track all reported defects, all proposed features, and all subsequent changes to parts. See also *change control*.

interest group. The list of actions that trigger notification to the user IDs associated with those actions listed in the notification list.

J

job queue. A queue of build scopes. One job queue exists for each TeamConnection family.

L

local version ID. In collision records, the database ID of the version of the current work area.

lock. An action that prevents editing access to a part stored in the TeamConnection development environment so that only one user can change a part at a time.

login. The name that identifies a user on a multi-user system, such as AIX or HP-UX, Solaris, or Windows NT. In OS/2 and Windows 95, the login value is obtained from the TC_USER environment variable.

M

map. The process of reassigning the meaning of an object.

metadata. In databases, data that describe data objects.

N

name server. In TCP/IP, a server program that supplies name-to-address translation by mapping domain names to Internet addresses.

National Language Support (NLS). The modification or conversion of a United States English product to conform to the requirements of another language or country. This can include the enabling or retrofitting of a product and the translation of nomenclature, MRI, or documentation of a product.

Network File System (NFS). The Network File System is a program that enables you to share files with other computers in networks over a variety of machine types and operating systems.

notification list. An object that enables component owners to configure notification. A list attached to a component that pairs a list of user IDs and a list of interest groups. It designates the users and the corresponding notification interest that they are being granted for all objects managed by this component or any of its descendants.

notification server. A server that sends notification messages to the client.

NTFS. NT file system.

NVBridge. A tool for automatic electronic distribution of TeamConnection software deliverables within a NetView DM/2 network.

O

operator. A symbol that represents an operation to be done. See also *comparison operators*.

originator. The user who opens a defect or feature and is responsible for verifying the outcome of the defect or feature on a verification record. This responsibility can be reassigned.

owner. The user who is responsible for a TeamConnection object within a TeamConnection family, either because the user created the object or was assigned ownership of the object.

P

parent component. All components in each TeamConnection family, except the root component, are created in reference to an existing component. The existing component is the parent component. See also *child component* and *component*.

parent part. Any part in a build tree that has a child defined. See also *part* and *child part*.

parser. A tool that can read a source file and report back a list of dependencies of that source file. It frees a developer from knowing the dependencies one part has on other parts to ensure a complete build is performed.

part. A collection of data that is stored by the family server and retrieved by a path name. They include text objects, binary objects, and modeled objects. These parts can be stored by the user or the tool, or they can be generated from other parts, such as when a linker generates an executable file.

path name. The name of the part under TeamConnection control. A path name can be a directory structure and a base name or just a base name. It must be unique within each release. See also *base name*.

pool. See *build pool*.

pop-up menu. A menu that, when requested, appears next to the object it is associated with.

prerequisite work areas. If a part is changed to resolve more than one defect or feature, the work area referenced by the first change is a

prerequisite of the work area referenced by later changes. A work area is a prerequisite to another work area if:

- Part changes are checked in, but not committed, for the first work area.
- One or more of the same parts are checked out, changed, and checked in again for the second work area.

problem tracking. The process of tracking all reported defects through to resolution and all proposed features through to implementation.

process. A combination of TeamConnection subprocesses, configured by the family administrator, that controls the general movement of TeamConnection objects (defects, features, work areas, and drivers) from state to state within a component or release. See also *subprocess* and *state*.

Q

query. A request for information from a database, for example, a search for all defects that are in the open state. See also *default query* and *search*.

R

raw format. Information retrieved on the report command that has the vertical bar delimiter separating field information, and each line of output corresponds to one database record.

refresh. This TeamConnection action updates a work area with any changes from the release, and it also freezes the work area, if it is not already frozen.

relative path name. The name of a directory or a part expressed as a sequence of directories followed by a part name, beginning from the current directory.

release. A TeamConnection object defined by a user that contains all the parts that must be built, tested, and distributed as a single entity.

restricted authority. The limitation on a user's ability to perform certain actions at a specific component. Authority can be restricted by the superuser, the component owner, or a user with AccessRestrict authority. See also *authority*.

root component. The initial component that is created when a TeamConnection family is configured. All components in a TeamConnection family are descendants of the root component. Only the root component has no parent component. See also *component*, *child component*, and *parent component*.

S

search. To scan one or more data elements of a set in a database to find elements that have certain properties.

serial development. While a user has parts checked out from a work area, no one else on the team can check out the part. The user develops new material without interacting with other developers on the project. TeamConnection provides the opportunity to hold the part until the user is sure that it integrates with the rest of the application. Thus, the lock is not released until the work area as a whole is committed. Contrast with *concurrent development*. See also *work area*.

server. A workstation that performs a service for another workstation.

shadow. A collection of parts in a filesystem that reflects the contents of a TeamConnection workarea, driver, or release.

shared part. A part that is contained in two or more releases.

shell script. A series of commands combined in a file that carry out a function when the file is run.

SID. The name of a version of a driver, release, or work area.

sizing record. A status record created for each component-release pair affected by a proposed defect or feature. The sizing record owner must

indicate whether the defect or feature affects the specified component-release pair and the approximate amount of work needed to resolve the defect or implement the feature within the specified component-release pair.

stanza format. Data output generated by the Report command in which each database record is a stanza. Each stanza line consists of a field and its corresponding values.

state. Work areas, drivers, features, and defects move through various states during their life cycles. The state of an object determines the actions that can be performed on it. See also *process* and *subprocess*.

subprocess. TeamConnection subprocesses govern the state changes for TeamConnection objects. The design, size, review (DSR) and verify subprocesses are configured for component processes. The track, approve, fix, driver, and test subprocesses are configured for release processes. See also *process* and *state*.

superuser. This privilege lets a user perform any action available in the TeamConnection family.

system administrator. A user who is responsible for all system-related tasks involving the TeamConnection server, such as installing, maintaining, and backing up the TeamConnection server and the database it uses.

T

task list. The list of tasks displayed in the Tasks window. The user can customize this list to issue requests for information from the server. Tasks can be added, modified, or deleted from the lists.

TCP/IP. Transmission Control Protocol/Internet Protocol.

TeamConnection client. A workstation that connects to the TeamConnection server by a TCP/IP connection and that is running the TeamConnection client software.

TeamConnection part. A part that is stored by the TeamConnection server and retrieved by a path name, release, type, and work area. See also *part*, *common part*, and *type*.

TeamConnection superuser. See *superuser*.

tester. A user responsible for testing the resolution of a defect or the implementation of a feature for a specific driver of a release and recording the results on a test record.

test record. A status record used to record the outcome of an environment test performed for a resolved defect or an implemented feature in a specific driver of a release.

track subprocess. An attribute of a TeamConnection release process that specifies that the change control process for that release will be integrated with the problem tracking process.

Transmission Control Protocol/Internet Protocol (TCP/IP). A set of communications protocols that support peer-to-peer connectivity functions for both local and wide area networks.

type. All parts that are created through the TeamConnection GUI or on the command line will show up in reports with the type of TcPart as the part type. The TeamConnection GUI and command line can only check in, check out, and extract parts of the type TcPart.

U

user exit. A user exit allows TeamConnection to call a user-defined program during the processing of TeamConnection transactions. User exits provide a means by which users can specify additional actions that should be performed before completing or proceeding with a TeamConnection action.

user ID. The identifier assigned by the system administrator to each TeamConnection user.

V

verification record. A status record that the originator of a defect or a feature must mark before the defect or feature can move to the closed state. Originators use verification records to verify the resolution or implementation of the defect or feature they opened.

version. (1) A specific view of a driver, release, or work area. (2) A revision of a part.

version control. The storage of multiple versions of a single part along with information about each version.

view. An alternative and temporary representation of data from one or more tables.

W

work area. An object in TeamConnection that you create and associate with a release. When the work area is created, you see the most current view of the release and all the parts that it contains. You can check out the parts in the work area, make modifications, and check them back into the work area. You can also test the modifications without integrating them. Other users are not aware of the changes that you make in the work area until you integrate the work area to the release. While you work on files in a work area, you do not see subsequent part changes in the release until you integrate or refresh your work area.

working part. The checked-out version of a TeamConnection part.

Y

year 2000 ready. IBM VisualAge TeamConnection Enterprise Server is Year 2000 ready. When used in accordance with its associated documentation, TeamConnection is capable of correctly processing, providing and/or receiving date data within and between the twentieth and twenty-first centuries, provided that all products (for example, hardware, software and

firmware) used with the product properly exchange accurate date data with it.

Index

Special Characters

/Ft(dir) builder parameter 157

A

Accept Defects window 52
Accept Test Records 96
Activate Fix Records 90
Add Driver Members 88
approval command
 approving a fix 82
Approval Records window 81
approve state 45
authority
 basic 26
 build 177
 for checking in parts 34
 for checking out parts 33
 for extracting parts 33

B

build action 7
build administrator 13
build agent
 installing on separate machines 137
 startup file, creating 144
 stopping 145
build environment 149
build event
 criteria used to determine success 150
 defining multiple outputs from one event 195
 definition of 130
 timeout setting 150
 with VisualAge C ++ templates 157
build function
 authority 177
 building a driver 92
 canceling a build 194
 collector part 196
 concepts of 129
 creating an MVS build server 137
 definition of 11
 diagram showing physical structure of 129
 features of 129
 installing 137
 installing an MVS build server 137
 monitoring build progress
 using Build Progress window 192
 using part -viewmsg command 192
 object model 132
 startup files, creating 144
 testing part updates 83

build mode 190
Build Parts window 60
build pool
 specifying when starting build 190
build scripts
 at work 191
 debug variable 153
 definition of 130
 for MVS
 compile example 168
 definition of 159
 file name conversions 164
 link example 172
 steps for working with 147, 159
 supported JCL syntax 165, 166
 writing 163
 for OS/2 150
 modifying contents of 155
 samples shipped 298
 testing 154
 timeout of 150
 writing 153
 writing an executable file for 153
build server
 creating for MVS 137
 installing 137
 installing on separate machines 137
 starting MVS 141
 startup file, creating 144
 stopping 145
 timeout setting 150
build target 190
build tree
 creating 184
 display of 188
 example of 134
 multiple outputs from single event 195
 setting up for packaging
 for the gather tool 202
 setting up for packaging 201
 versions of 134
 working with 133
builder
 command
 connecting builder to its parts 156
 creating a builder 148, 154
 extracting a builder 155
 modifying builder contents 155
 connecting to parts 155, 178
 creating a null builder 149, 196
 definition 130

- builder (*continued*)
 - for MVS
 - creating 159
 - environment supported 162
 - naming 160
 - passing parameters to a build script 164
 - versions of 160
 - for OS/2
 - creating 147
 - environment supported 149
 - naming 148
 - passing parameters to a build script 151
 - versions of 148
 - removing from a part 156
- building an application
 - an example
 - build scripts at work 191
 - creating a build tree 184
 - creating builders and parsers 183
 - extracting resulting executable 192
 - list of tasks 181
 - starting a build on the client 189
 - starting servers and agents 182
 - authority 177
 - building all parts ignoring times 193
 - canceling a build 194
 - monitoring progress of build 192
 - preparing your parts 30, 135
 - report of which parts will be built 194
 - running in spite of errors 193
 - testing part updates 60
 - with VisualAge C++ and templates 157
- buildView action 28

C

- canceled state 43
- change control 3
- check-in action 7
- Check In Parts window 58
- check-out action 7
- checking in parts
 - authority needed to 34
 - example of 58
 - explanation of 34
- checking out parts
 - an example of 54
 - authority needed to 33
 - explanation of 33
- Checkout Parts window 55
- client
 - definition 5
 - starting 18
 - stopping 19
- closed state 44

- collector part
 - example of 196
 - using a null builder 149
- collision command
 - reconciling differences 76
- collision record
 - example of 34
 - reconciling using command line 75
 - reconciling using GUI 74
 - when creating driver members 89
- command file
 - fhbopars.cmd sample shipped 178
 - specifying 176
 - writing 177
- command line interface
 - using 22
 - viewing syntax online 22
- commands
 - becoming familiar with 22
 - fhomigmk 237
 - for client
 - teamc.log file 22
 - viewing syntax online 22
 - tmerge 75
 - teamcgui 18
- Commit Drivers 94
- commit state
 - of drivers 48
 - of work areas 47
- committing
 - a driver 94
 - versus integrating 97
- common parts
 - between releases 32
 - between work areas 32
 - breaking link 32
 - definition of 32
 - locking 32
- Complete Drivers 95
- Complete Fix Records 87
- complete state
 - moving work area to 96
 - of drivers 48
 - of work areas 47
- components
 - definition of 7
 - displaying structure of 25
 - example of hierarchy 7
 - information stored about 7
- Components window 25
- concepts of
 - build function 129
 - TeamConnection 3

- concurrent development
 - definition of 51
 - example of 34, 70
 - how to work in 31
 - reconciling differences in no-track process 73
 - reconciling differences in tracking process 89
- configuration management 3
- configuring
 - processes 96
- connect function 31
- create action 6
- Create Builder 147
- Create Parser 175
- Create Parts window 184
- Create Work Areas window 54
- D**
- DEBUG 153, 154
- defect command
 - accepting 53
 - closing 70
 - modifying ownership 79
 - verifying 70
- defects
 - accepting 52, 80
 - approving the fix 81
 - closing 69
 - definition of 10
 - designing 78
 - reassigning ownership 79
 - reviewing 78
 - sizing 78
 - states of
 - canceled state 43
 - closed state 44
 - design state 43
 - open state 43
 - return state 43
 - review state 43
 - size state 43
 - verify state 43
 - working state 43
 - verifying 69
 - working with 38
- delta file tree 48
- dependencies on a build
 - defining through parsers 131
 - viewing through a build tree 133
- dependent part 31
- design, size, review process 78
- design changes 38
- design state 43
- development mode 32
- differences in parts
 - reconciling in no-track process 73

- differences in parts *(continued)*
 - reconciling in tracking process 89
- driver command
 - committing a driver 94
 - completing a driver 95
 - refreshing driver 92
 - restricting driver 94
- driver member 88
- driverMember command
 - adding driver members 88
- drivers
 - building 92
 - committing changes into release 94
 - completing 95
 - definition of 9
 - preparing for formal test 95
 - refreshing 91
 - states of
 - commit state 48
 - complete state 48
 - integrate state 47
 - restrict state 48, 93
 - working state 47
 - versioning 37
- E**
- edit action 7
- Edit Task List window 56
- editing parts 34
- environment variables 229
 - setting 235
 - setting before invoking MVS build script 165
 - setting before invoking OS/2 build script 151
 - setting before invoking Windows NT build script 151
 - setting for command line usage 22
 - setting for GUI usage 20
- examples of
 - build components on separate machines 137
 - build script for an MVS compile 168
 - build script for an MVS link 172
 - build tree 134
 - building an application 181
 - client/server network 4
 - component hierarchy 7
 - concurrent development 70
 - display of a build tree 188
 - executable file for a build script 153
 - following a no-track process 51
 - following a tracking process 77
 - listing work areas in a release 100
 - retrieving past part versions 97
 - rules file 240
 - serial development 51
 - showing part/release/component relationship 8

- examples of *(continued)*
 - starting make import 238
 - teamcpak command for Gather 208
 - teamcpak command for NVBridge 217
 - writing a build script 153
- expand keywords 20
- extract action 7
- Extract Parts window 62
- extracting parts
 - an example of 62
 - authority needed 33
 - previous versions of 100
 - resulting build executable 192
 - versus checking out 33

F

- family 6
- family administrator responsibilities 12
- features
 - definition of 10
 - states of
 - canceled state 43
 - closed state 44
 - design state 43
 - open state 43
 - return state 43
 - review state 43
 - size state 43
 - verify state 43
 - working state 43
 - working with 38
- fhbopars.cmd, sample command file 178
- files
 - teamc.log 22
- filter windows for parts 27
- finding objects 27
- fix command
 - accepting fix records 87
 - completing fix records 87
 - reactivate 90
- fix records
 - accepting 87
 - completing 87
 - moving back to active state 90
 - reactivating 90
 - when created 81
- fix state 45
- Fix Work Areas 89
- Freeze Work Areas window 65
- freezing work areas
 - examples of 65, 85
 - explanation of 28
- full part tree 48

G

- Gather tool
 - explanation of 205
 - packaging file
 - example of syntax 208
 - keywords for 209
 - specifying 206, 216
 - syntax rules for 208
 - writing 208
 - teamcpak command 205
- GUI
 - client
 - accepting a defect 52
 - accepting fix records 87
 - accepting test records 96
 - accessing online help 21
 - adding driver member 88
 - approving the fix 81
 - building a driver 92
 - building parts 60
 - checking in part 58
 - checking out a part 54
 - closing a defect 70
 - committing driver changes into release 94
 - completing a driver 95
 - connecting builder to its parts 155
 - connecting parser to parts 178
 - creating a work area 53
 - creating parsers 175
 - extracting parts 62
 - fast path 20
 - freezing a work area 65
 - integrating a work area, concurrent development 72
 - integrating a work area, serial development 68
 - moving fix records back to active state 90
 - reactivating fix records 90
 - reassigning defect ownership 79
 - reconciling collisions 74
 - refreshing a driver 91
 - refreshing a work area, concurrent development 71
 - refreshing a work area, serial development 66
 - removing builder connection from parts 156
 - removing parser connection from parts 179
 - restricting a driver 93
 - returning work area to fix state 89
 - searching for parts 55
 - Settings notebook 20
 - starting 18
 - stopping 19
 - Tasks window 18
 - using 17
 - verifying a defect 70

H

help

- diagram push button 21
- how do I 21
- how to access 21

hierarchy

- component example 7
- displaying component structure 25

how do I help 21

I

importing makefile information 237

input part 30

installation

- build components on separate machines 137
- creating an MVS build server 137

integrate state

- of drivers 47
- of work areas 46

Integrate Work Areas window 68

integrating

- commit versus integrate 97

interfaces

- becoming familiar with 17
- description of 5

K

keyword

- DATA 209, 218
- EXITPOST 212
- EXITPRIOR 212
- EXITREPLACE 212
- for Gather utility 208, 209
- in text files 277
- INSTALLS 216
- MAIL 215
- NVGLOBALS 218, 219, 220
- package file 208
- PACKAGEFORMAT 209, 218
- RULE 209
- SENDS 216
- SOURCE 210
- TARGET 211
- TARGETROOT 209

L

LANG 229

M

makefile

- creating rules file 239
- example of starting import 238
- importing information 237
- merging differences 75
- Modify Defect Owner window 79

Modify Part Properties 155

MVS

build script

- definition of 159
- file name conversions 164
- for a compile 168
- for a link 172

builder 159

- creating a build server 137
- installing a build server 137
- modifying RUNPGM JCL 141
- starting build server 141
- stopping build server 145
- supported JCL syntax 165, 166
- syntax for builds 165

N

naming

- builders 148
- parts 30
- work areas 29

network 4

NLSPATH 229

no-track process, scenario 51

O

online help

- diagram push button 21
- how to access 21

open state 43

output part 30

P

package file

- for Gather tool
 - keywords for 209
 - specifying 206, 216
 - syntax rules for 208
 - writing 208
- for Tivoli Software Distribution tool
 - example of syntax 217
 - keywords for 218
 - syntax rules for 217
 - writing 217

packaging

- definition of 12
- explanation of Gather tool 205
- sample files shipped 300
- setting up build tree
 - for the gather tool 202
- tasks involved in 201, 215

parameters

- passing to a build script 151
- where specified
 - attributes of a builder 151

- parameters *(continued)*
 - where specified *(continued)*
 - attributes of part in a build tree 152
 - parameters of part -build command 152
- parser command
 - connecting parser to parts 178
 - creating a parser 176
- parsers
 - command file
 - fhbopars.cmd, sample shipped 178
 - specifying 176
 - writing 177
 - creating 175
 - definition of 131
 - explanation of 175
 - removing from a part 179
 - samples shipped 299
- part command
 - building a driver 93
 - building parts 61
 - canceling a build 194
 - checking in parts 59
 - checking out parts 55, 58
 - extracting build executable 192
 - extracting parts 62
 - listing parts that will be built 194
 - monitoring progress of build 192
 - removing builder connection 157
 - removing parser connection 180
 - starting a build 190
 - viewing all version parts in a work area 98
- partFull action 28
- parts
 - authority needed to check in 34
 - authority needed to check out and extract 33
 - checking in 34
 - checking in, an example 58
 - checking in versus integrating 34
 - checking out, an example 54
 - checking out versus extracting 33
 - common
 - between releases 32
 - between work areas 32
 - breaking link 32
 - definition of 32
 - connecting builder 155
 - connecting parser 178
 - creating 29
 - definition of 6
 - dependent 31
 - editing parts 34
 - empty 29
 - extracting 33
 - extracting, an example 62

- parts *(continued)*
 - finding
 - using BuildView action 28
 - using Filter window 27
 - using PartFull action 28
 - using Parts action 27
 - using report command 27
 - input 30
 - linking 32
 - linking between releases 32
 - linking between work areas 32
 - locked 31
 - making changes visible to a release 94
 - making changes visible without driver subprocess 97
 - naming 30
 - preparing for build 30
 - removing builder from 156
 - removing parser from 179
 - retrieving past versions 97, 100
 - searching for 27
 - searching for, an example 55
 - testing updates to 60, 83
 - versioning 38
 - viewing versions in a work area 98
 - where placed on workstation 33
- parts action 27
- Parts Filter window 56
- PATH 229
- placeholder parts 29
- problem information, reporting 38
- processes
 - definition of 10
 - example of using a configured process 96
 - relating to defects and features 38

R

- reassigning ownership of a defect 79
- Reconcile Collision Record window 74
- reconciling differences 73, 89
- Refresh Drivers 91
- Refresh Work Areas window 66
- refreshing
 - a driver 91
 - a work area 66
- relative flag 33
- release management 3
- releases
 - committing driver changes to 94
 - common parts 32
 - definition of 8
 - example of relationship with other objects 8
 - linking parts between releases 32
 - parts common to more than one 32

- releases (*continued*)
 - using a configured process 96
 - versioning 35
- report command
 - listing work areas in a release 100
 - searching for parts 57
 - to find parts 27
 - used to view differences 75
- Restrict Drivers 93
- restrict state
 - of drivers 48, 93
 - of work areas 46
- retrieving past versions of objects 35
- return state 43
- review state 43
- REXX execs 295
- rules file
 - creating 239
 - example of 240
- RUNPGM JCL 141

S

- sample files shipped
 - build script for NVBridge tool 215
 - build scripts 298
 - for package function 300
 - parsers 299
 - REXX execs 295
- Save to Task List push button 56
- scenarios
 - concurrent development 70
 - no-track process
 - accepting a defect 52
 - checking in a part 58
 - checking out a part 54
 - closing a defect 69
 - creating a work area 53
 - extracting a part 62
 - freezing the work area 65, 66
 - integrating a work area 68
 - reconciling differences 73
 - searching for a part 55
 - testing part updates 60
 - verifying a defect 69
 - tracking process
 - accepting a defect 80
 - accepting fix records 87
 - accepting test records 95
 - adding driver member to a driver 88
 - approving the fix 81
 - building a driver 92
 - checking out a part 82
 - committing driver changes into release 94
 - completing a driver 95
 - completing fix records 87

- scenarios (*continued*)
 - tracking process (*continued*)
 - designing a defect 78
 - freezing the work area 85
 - reactivating fix record 90
 - reassigning defect ownership 79
 - refreshing a driver 91, 93
 - returning work area to fix state 89
 - reviewing a defect 78
 - sizing a defect 78
 - testing part updates 83
 - searching for objects 27
- serial development
 - definition of 51
 - example of 51
 - how to work in 31
- servers
 - definition 5
 - family server
 - definition 5
- Settings notebook
 - for client 20
 - list of variables 20
- size state 43
- sizing records 78
- starting
 - a build on the client 189
 - GUI client 18
 - MVS build server 141
- startup files 144
- states of objects
 - defects 41
 - drivers 47
 - features 41
 - work areas 45
- stopping
 - build agent 145
 - MVS build server 145
 - OS/2 build server 145
- superuser 6
- syntax
 - for MVS builds 165
 - how to view online 22
 - supported JCL syntax for build 165, 166
- system administrator, responsibilities 12

T

- tasks
 - authority to perform 26
 - preparing to build an application 135
 - understanding the basics 25
 - when following a no-track process 52, 77
- Tasks window 18
- TC_BECOME 229, 295
- TC_BUILDPOOL 229

- TC_CASESENSE 229
- TC_COMPONENT 22, 229, 240
- TC_DBPATH 229
- TC_FAMILY 22, 23, 51, 151, 182, 229, 235, 237, 238, 247, 295
- TC_INPUT 151, 153, 154, 247
- TC_INPUTTYPE 151, 247
- TC_LOCATION 151, 247
- TC_MAKEIMPORTRULES 229, 237, 240
- TC_MAKEIMPORTTOP 229, 238
- TC_MAKEIMPORTVERBOSE 229, 238
- TC_MIGRATERULES 229
- TC_NOTIFY_DAEMON 229
- TC_OUTPUT 153, 154
- TC_OUTPUTTYPE 151, 247
- TC_RELATIVE 83
- TC_RELEASE 151, 229, 237, 238, 247, 295
- TC_TOP 33, 229
- TC_TRACE 229
- TC_TRACEFILE 229
- TC_TRACESIZE 229
- TC_USER 229
- TC_WORKAREA 151, 229, 237, 247
- tlogin command 22
- tcmerge 75
- teamcgui command 18
- TeamConnection
 - concepts of 3
 - diagram showing physical structure of 129
 - introducing 3
 - the basics of using 25
- teamcpak command
 - for Gather tool
 - command line flags 206
 - examples of 208
 - starting 205
 - syntax of 205
 - for NVBridge tool
 - examples of 217
- templates 157
- test command
 - accepting test records 96
- test records
 - accepting 95
 - when created 95
- test state
 - moving work area to 95
 - of work areas 47
- Tivoli Software Distribution tool
 - explanation of 215
 - packaging file
 - example of syntax 217
 - keywords for 218
 - syntax rules for 217

- Tivoli Software Distribution tool *(continued)*
 - packaging file *(continued)*
 - writing 217
 - problem determination 220
 - sample build script 215
 - teamcpak command 215
 - Tivoli Software Distribution output files 215
 - used as a builder for packaging 215
- tracking
 - following a no-track process 51
 - following a tracking process 77

V

- verification record 43
- Verify Defects window 70
- verify state 43
- version control 3, 35
- versions 35
 - builders 148
 - of build trees 134
 - of drivers 37
 - of parts 38
 - of releases 35
 - of work areas 36
 - retrieving past part versions 97
 - viewing all part versions in a work area 98
- VisualAge C ++ 157

W

- window examples
 - Accept Defects 52
 - Accept Test Records 96
 - Activate Fix Records 90
 - Add Driver Members 88
 - Approval Records 81
 - Build Parts 60
 - BuildView 28
 - Check In Parts 58
 - Check Out Parts 55
 - Commit Drivers 94
 - Complete Drivers 95
 - Complete Fix Records 87
 - Components window 25
 - Create Builder 147
 - Create Parser 175
 - Create Parts 184
 - Create Work Areas 54
 - Edit Task List 56
 - Extract Parts 62
 - Fix Work Areas 89
 - Freeze Work Areas 65
 - Integrate Work Areas 68
 - Modify Defect Owner 79
 - Modify Part Properties 155
 - PartFull 28

- Parts Filter 27, 56
 - Reconcile Collision Record 74
 - Refresh Drivers 91
 - Refresh Work Areas 66
 - Restrict Driver 93
 - Tasks 18
 - Verify Defects 70
 - work area
 - automatic creation of 28
 - canceling 29
 - creating 53
 - definition of 8
 - freezing 28
 - freezing, an example 65, 85
 - integrating, concurrent development 72
 - integrating, serial development 68
 - linking parts between work areas 32
 - moving to complete state 96
 - moving to test state 95
 - naming 29
 - reactivating to fix state 90
 - refreshing, concurrent development 71
 - refreshing, serial development 66
 - returning to fix state 89
 - reusing 29
 - states of
 - approve state 45
 - commit state 47
 - complete state 47
 - fix state 45
 - integrate state 46
 - restrict state 46
 - test state 47
 - things you can do with 9
 - using 28
 - versioning 36
 - viewing all version parts 98
 - when parts become visible to others 28
 - when to create one 28
 - workarea command
 - creating 54
 - freezing work areas 65
 - integrating when driver subprocess is not enabled 97
 - integrating work areas, concurrent development 73
 - integrating work areas, serial development 69
 - refreshing work areas, concurrent development 71
 - refreshing work areas, serial development 67
 - return work area to fix state 89
 - working state
 - of defects and features 43
 - of drivers 47

Readers' Comments — We'd Like to Hear from You

**IBM VisualAge TeamConnection Enterprise Server
User's Guide
Version 3.0**

Publication No. SC34-4499-04

Overall, how satisfied are you with the information in this book?

	Very Satisfied	Satisfied	Neutral	Dissatisfied	Very Dissatisfied
Overall satisfaction	<input type="checkbox"/>				

How satisfied are you that the information in this book is:

	Very Satisfied	Satisfied	Neutral	Dissatisfied	Very Dissatisfied
Accurate	<input type="checkbox"/>				
Complete	<input type="checkbox"/>				
Easy to find	<input type="checkbox"/>				
Easy to understand	<input type="checkbox"/>				
Well organized	<input type="checkbox"/>				
Applicable to your tasks	<input type="checkbox"/>				

Please tell us how we can improve this book:

Thank you for your responses. May we contact you? Yes No

When you send comments to IBM, you grant IBM a nonexclusive right to use or distribute your comments in any way it believes appropriate without incurring any obligation to you.

Name

Address

Company or Organization

Phone No.



Cut or Fold
Along Line

Fold and Tape

Please do not staple

Fold and Tape



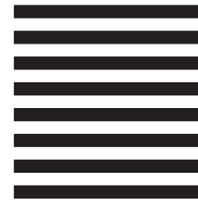
NO POSTAGE
NECESSARY
IF MAILED IN THE
UNITED STATES

BUSINESS REPLY MAIL

FIRST-CLASS MAIL PERMIT NO. 40 ARMONK, NEW YORK

POSTAGE WILL BE PAID BY ADDRESSEE

IBM Corporation
Information Development
Department G71A / Bldg 062
P.O. Box 12195
Research Triangle Park, NC
27709-2195



Fold and Tape

Please do not staple

Fold and Tape

Cut or Fold
Along Line



Part Number: CT66GNA



Printed in the United States of America
on recycled paper containing 10%
recovered post-consumer fiber.

SC34-4499-04



CT66GNA

