StarTeam 16.2

Server Administration Help

Micro Focus
The Lawn
22-30 Old Bath Road
Newbury, Berkshire RG14 1QN
UK
http://www.microfocus.com

2017-10-13

Contents

StarTeam Server Administration

Welcome to StarTeam Server Administration

Welcome to StarTeam
Server Administration Overview
Licensing the Server
Guidelines for Deploying StarTeam
Working with Server Configurations
Server Administration Tool

Online resources

Micro Focus Infocenter
Micro Focus SupportLine
Micro Focus Product Updates
Micro Focus Knowledge Base
Micro Focus Community Forums
Micro Focus Training

Provide feedback

Contacting Support
Email us feedback regarding this Help

Welcome to StarTeam

StarTeam is a software change and configuration management solution designed to meet the needs of local and distributed teams regardless of size and work style. Team members can work whenever and wherever they like and benefit from integrated change management, defect tracking, file versioning, requirements management, and project and task management capabilities for flexible project control. StarTeam is a robust platform for coordinating and managing the entire software project throughout the software development life-cycle.

Introduction

Installation and Licensing for StarTeam

Installation

Installation instructions for installing StarTeam products can be found in the *StarTeam Installation Guide*.

Licensing

StarTeam is available in three licensing packages:

Enterprise StarTeam Enterprise provides a basic feature set, including the StarTeam Server, StarTeamMPX (MPX Event Transmitter and MPX Message Broker), the Cross-Platform Client, StarTeam Web Client, LDAP QuickStart Manager, and the SDK. The requirements component is not available with this license, however, it does provide access to custom fields.

Enterprise Advantage StarTeam Enterprise Advantage has all the StarTeam Enterprise features plus the Requirement component, StarTeamMPX (MPX Cache Agent and MPX File Transmitter), and StarTeam Workflow Extensions which include alternate property editors (APEs) that enable you to create custom forms and design workflow rules to control how all the items in a component move from state to state. StarTeam Datamart is available for purchase.

Products Included with StarTeam Enterprise Licenses

The following provides a summary of StarTeam products that come with the StarTeam Enterprise license. The installation instructions for some products are not in this consolidated installation guide, but are located in the respective guide of that product and are noted.

StarTeam Server A StarTeam Server stores artifacts (files, change requests/defects, tasks, and topics) for StarTeam clients. A server can support one or more server configurations on the same computer. Install StarTeam Server on a computer that is accessible to all users.

MPX Message Broker Pushes information from the StarTeam Server to its clients. Usually an administrator sets up a cloud of Message Brokers to improve server performance for users in diverse geographic locations. One (sometimes two) root Message Brokers are set up for the server, usually on the same computer or in a network-near location.

StarTeam Cross-Platform Client The StarTeam Cross-Platform Client is the most used client and provides users with access to all of the artifacts on the server. The Cross-Platform Client is a pure Java client that provides support of operating systems where a compatible JRE or JDK are available. As such, Cross-Platform Client is available for the Microsoft Windows, Solaris, and Linux operating systems.

StarTeam Visual Studio Plugin The StarTeam Visual Studio Plugin provides the StarTeam software configuration management capabilities tightly integrated with the Microsoft Visual Studio development environment. Using this integration makes it possible for you to develop applications in the Microsoft Visual Studio environment while simultaneously using the version control, change request, topic, task, and requirement component assets of StarTeam. The integration brings StarTeam main menu commands, context menu commands, and an embedded StarTeam client (providing much of the same look-and-feel as the full-featured Cross-Platform Client) to the Microsoft Visual Studio development environment.

StarTeam Eclipse Plugin StarTeam Edition for Eclipse allows you to share projects on StarTeam Server and projects in the Eclipse workspace, but it is much more than just a version control plug-in. This integration offers project teams a customizable solution providing requirements, task, and change management, defect tracking and threaded discussions tightly integrated within the Eclipse platform.

StarTeam Web Server and StarTeam Web Client The StarTeam Web Server makes it possible for users to access the server from their browsers using the StarTeam Web Client. The StarTeam Web Client is an intuitive web-based interface that many simultaneous users can use to connect to one or more StarTeam Servers to access projects and manage items. This product contains a core feature set designed to meet the needs of users responsible for viewing, creating, and editing StarTeam change requests, requirements, tasks, and topics. Browsing files and a limited set of file operations are also available.

Note: You must have a StarTeam user license to use the Web Client.

LDAP Quickstart Manager

The StarTeam Server can provide password authentication via a directory service, such as LDAP Quickstart Manager (QSM) to add users to the server, along with their distinguished names (DN) (needed for authentication) and other user information.

Layout Designer

Use Layout Designer to create forms for artifacts, such as change requests. This allows you to put the most important properties on the first tab, etc. With the web client and an Enterprise Advantage server, a Layout Designer form works with workflow. This is not true of the StarTeam Cross-Platform Client where Layout Designer's use is only for form building.

This product is translated into English, French, German, and Japanese.

StarTeam SDK

The StarTeam SDK is cross-compiled so that it can be offered both as Java and .NET applications. The full SDK is used by developers to create additional applications that use the StarTeam Server.

Usually, the StarTeam SDK runtime is installed with clients automatically so it can be used by them to access the StarTeam Server. Occasionally, you may need to install the runtime.

StarTeam SCC Integration

The StarTeam SCC Integration works with any application that uses the Microsoft Source Code Control (SCC) Application Programming Interface (API). This API, originally designed by Microsoft to allow applications to work with Microsoft Visual SourceSafe, enables you to perform version control operations, such as checking files in and out, using StarTeam as the SCC provider.

StarTeam Quality Center Synchronizer

This product is available with all licenses.

StarTeam Quality Center Synchronizer can ensure that the same data appears in Quality Center and a database used by StarTeam Server. The goal of the synchronization is to provide access to the latest information about defects, whether the defects are being processed from Quality Center or from StarTeam. You can use Quality Center to add defects, and you can use StarTeam to indicate that those defects have been fixed and vice versa. Team members do not need to be aware of where the defect was last processed. The latest data is available at all times, as long as the databases are synchronized frequently.

StarTeam Microsoft Project Integration

Available with all licenses.

The interaction of the StarTeam Microsoft Project Integration and Microsoft Project make the jobs of both project planners and team members easier. Project planners use Microsoft Project to list the tasks that workers must perform. After exporting the tasks to StarTeam, they can gather information about the work accomplished by each team member in StarTeam, rather than communicating individually with each team member.

StarTeam Toolbar Utility

The StarTeam **Toolbar Utility** (Toolbar) is a component of the StarTeam designed to make it easier for you to log on to multiple servers and to launch different programs. It automatically caches the user name and password used to log on to each StarTeam or Caliber server, reducing the number of times that you must enter your logon information. The Toolbar is initially populated with shortcuts for the tools of the StarTeam and Caliber products that are installed on your workstation. Because the Toolbar uses the standard Microsoft Windows program shortcut feature, you can easily add any other program as a tool.

File Compare/ Merge File Compare/Merge is a graphical compare/merge tool delivered with the StarTeam Cross-Platform Client. It enables you to compare a file dynamically with the file in the repository, and manually or automatically merge the content of the two files. File differences are highlighted in the File Compare/Merge panes using a configurable color scheme, and action buttons display in the highlighted areas to simplify the merging process.

View Compare/ Merge View Compare/Merge is a comprehensive tool for comparing and merging views available with the StarTeam Cross-Platform Client. There are two versions of View Compare/Merge:

- **Graphical** Provides interactive comparison and merging with per-item and per-folder interaction, allowing you to carefully control which items are compared and how each difference is resolved
- **Command-line** Enables batch/shell-directed sessions.

Products Included with StarTeam Enterprise Advantage Licenses

In addition to the products included with StarTeam Enterprise licenses, StarTeam Enterprise Advantage licenses also include the products listed below. The installation instructions for some products are not in this consolidated installation guide, but are located in the respective guide of that product.

MPX Cache Agent A root MPX Cache Agent monitors the server's repository for file content and object properties. Via Message Broker, the data is pushed to remove MPX Cache Agents that are network-near to members of dispersed teams, improving the speed with which users access the data they need.

StarTeam Extensions StarTeam Extensions enables clients to take advantage of workflow and custom toolbar applications. The StarTeam Extensions files must be checked into the StarFlow Extensions project on each server configuration. If there is no StarFlow Extensions project, you need to create one.

StarTeam Extensions also provides API documentation and samples.

StarTeam Workflow Designer Use the StarTeam Workflow Designer to create workflows for specific artifact types (such as change requests/defects) per project or even per view.

StarTeam Notification Agent The StarTeam Notification Agent runs on the same computer as the StarTeam Server (or on a network-near computer) so that it can monitor the server and send notifications set up in your workflow.

Search Search allows users to perform ad hoc queries across servers and projects. The query results reflect the access rights of the user logged on to Search so information is shared across the organization without compromising security.

Datamart* StarTeam Datamart is a complementary product to the StarTeam Server. StarTeam Datamart uses the StarTeam SDK to communicate with the StarTeam Server to create a reporting database that you can use with popular third party reporting applications such as Crystal Reports and Business Objects (reporting applications are not included with StarTeam Datamart). StarTeam Datamart extracts data from a StarTeam Server and places the data into a relational database, where reporting tools can access it. StarTeam Datamart can extract information from every project, every view in each project, every folder in each view, and every item in each folder, and labels, links, and history for each item. You can restrict extraction of data to a particular project and view or only extract certain tables.

Datamart stores the data in any StarTeam supported database..

The product comes with both an Extractor (for an initial retrieval) and with a Synchronizer to update existing data sets.

TeamInspector*

TeamInspector is a continuous integration build server and build inspection tool. It works with StarTeam, Subversion, Perforce, and ClearCase. It requires the use of a database: Microsoft SQL Server 2005 SP3, Oracle 10g Release 2 version 10.2.0.4, or Apache Derby 10.4.2.0 or later.

Rhythm*

Rhythm is an Agile project tracking tool designed to allow you to:

- Organize, prioritize, and manage your Agile teams' backlogs.
- Plan your sprints, task out the work, and then track progress throughout the sprint.
- Get comprehensive visibility of all your Agile assets.

* Can be purchased separately and added to the Enterprise package.

Contacting Support

Micro Focus is committed to providing world-class services in the areas of consulting and technical support. Qualified technical support engineers are prepared to handle your support needs on a case-by-case basis or in an ongoing partnership. Micro Focus provides worldwide support, delivering timely, reliable service to ensure every customer's business success.

For more information about support services, visit the Micro Focus SupportLine web site at *http://supportline.microfocus.com* where registered users can find product upgrades as well as previous versions of a product. Additionally, users can find the Knowledge Base, Product Documentation, Community Forums, and support resources.

When contacting support, be prepared to provide complete information about your environment, the product version, and a detailed description of the problem, including steps to reproduce the problem.

For support on third-party tools or documentation, contact the vendor of the tool.

Standard StarTeam Architecture Overview

The standard architecture represents the minimal components present in a StarTeam instance: a StarTeam Server process managing a vault and a database and one or more StarTeam clients. With just these components, all basic StarTeam functionality is available. The core components of the standard StarTeam architecture are depicted below.

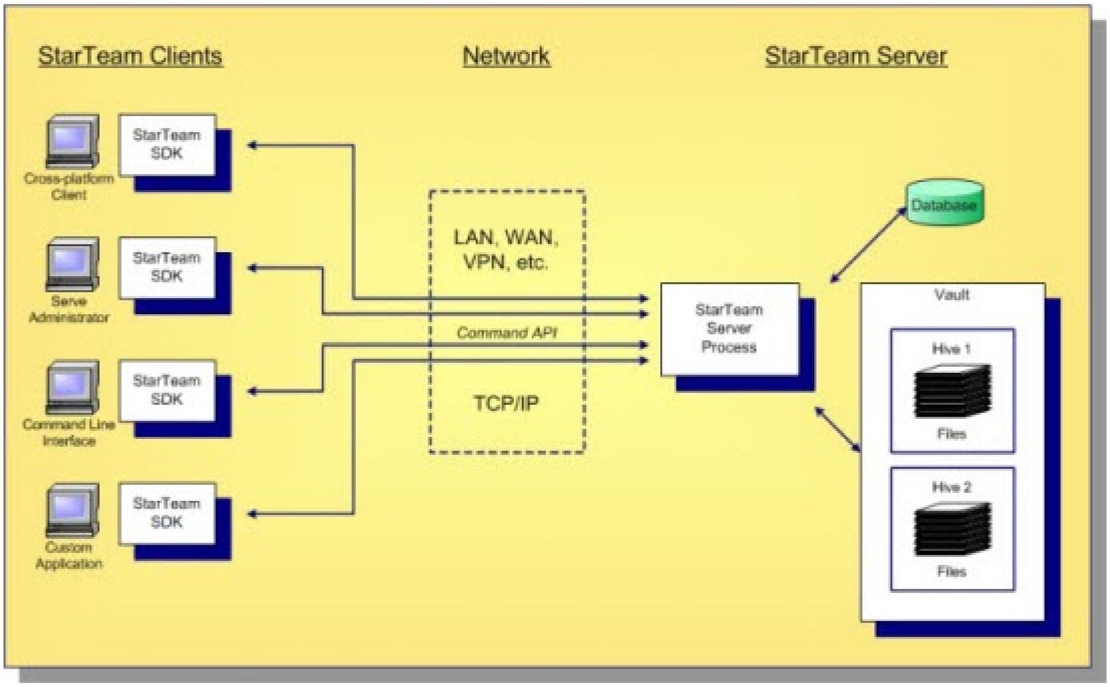

StarTeam employs a client/server architecture. The StarTeam Cross-Platform Client, **Server Administration** tool, and StarTeam Command Line Tools are examples of bundled StarTeam clients. StarTeam clients use the freely available StarTeam SDK, so you can write custom applications that have access to the same features as the bundled clients. The SDK is fully featured in Java, .NET, and COM, allowing you to write custom applications for any environment. A single StarTeam client can have multiple sessions to any number of StarTeam Servers.

All StarTeam clients connect to a StarTeam Server process using TCP/IP, so virtually any kind of network can be used: LAN, WAN, VPN, or the public Internet. StarTeam uses a proprietary protocol called the *command API*, which supports compression and multiple levels of encryption. The command API has been optimized to support high performance, automatic reconnect, delta check-out for slow connections, and other important features.

A single deployment instance of StarTeam is known as a *server configuration*, usually shortened to just *configuration*. The persistent data of a configuration consists of a database and a *vault* and is managed by a single StarTeam Server process. The database holds all metadata and non-file artifacts, whereas file contents are stored in the vault. The database can be any of the supported databases and it can reside on the same machine as the StarTeam Server process or a separate machine. The StarTeam database and vault can be backed-up dynamically, while the server is in use. This supports 24 x 7 operations that want to minimize down time.

StarTeam's vault is a critical component that affects performance and scalability. In contrast to the traditional delta storage technique, StarTeam's vault uses an innovative architecture designed for scalability, performance, high availability, and dynamic expandability. Today, customers are storing up to a terabyte of data in a single StarTeam vault, but it was designed to store content up to a petabyte and beyond.

Within the vault, files are stored in containers known as hives. A hive is a folder tree containing archive and cache files on a single disk volume. Hives can be dynamically added on existing or new disk volumes, thereby allowing virtually unlimited capacity. StarTeam stores each file revision in a separate archive file in a manner that minimizes space usage as well as duplicate content. StarTeam's vault uses less space than delta-based storage. In certain cases where it is more economical to send file deltas to clients instead of full versions, StarTeam generates and caches delta files. However, in most cases sending full versions is more economical.

MPX Components

Like all client/server architectures, as the number of clients grows, the server could potentially become a bottleneck. In fact, the scalability of many client/server systems is entirely limited by this bottleneck. Other client/server systems address scalability by deploying multiple instances and replicating information between them to attain synchronization.

is a unique solution to client/server scalability. is a publish/subscribe messaging framework that pushes update events that contain metadata and data to clients. It is optional because it is not required for basic StarTeam functionality. However, when is activated, it improves StarTeam Server scalability and improves StarTeam client responsiveness.

Message Broker

Basic requires the addition of a single extra component, known as the MPX Message Broker. The MPX Message Broker's role is illustrated below.

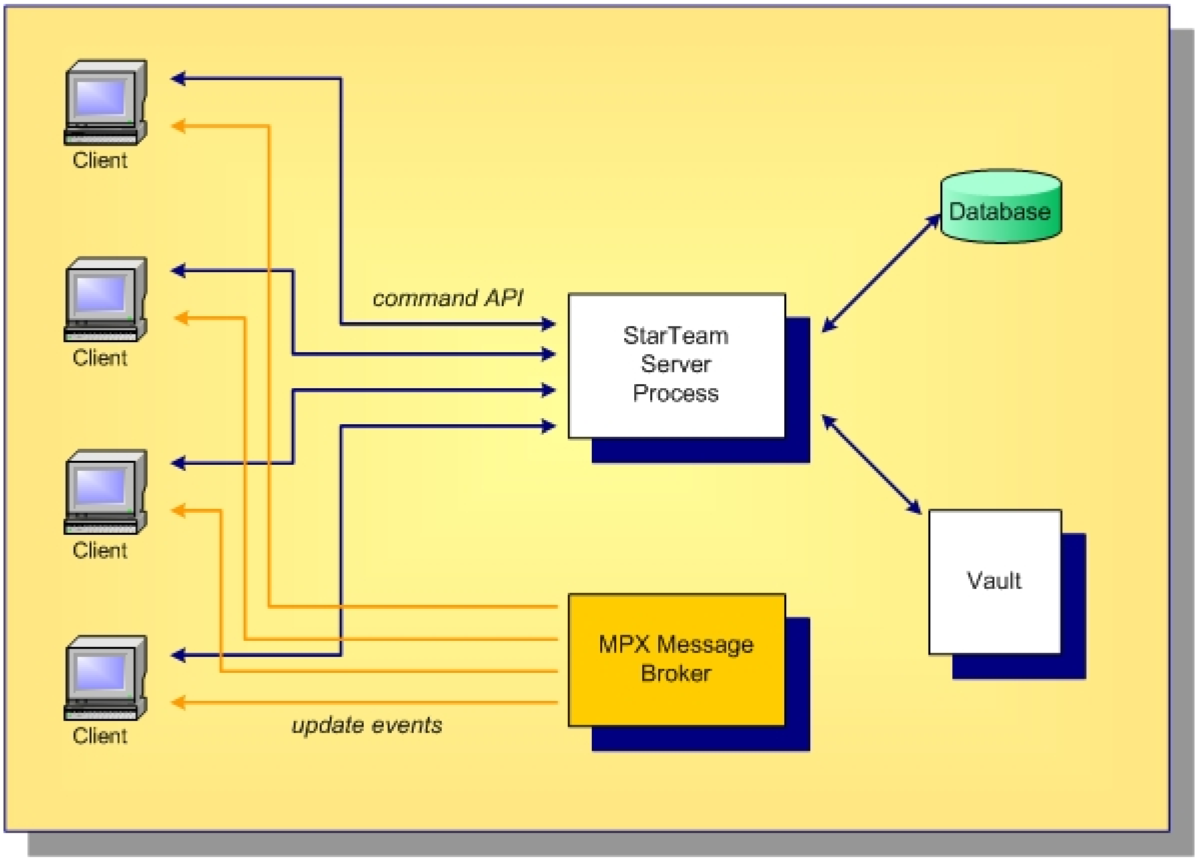

The MPX Message Broker is a messaging process that uses an event API to receive updates from the StarTeam Server process. The MPX Message Broker broadcasts encrypted messages containing updated artifacts. StarTeam clients subscribe to *subjects* and receive only messages relevant to them. By receiving updates as soon as they occur, StarTeam clients do not need to poll for updates or refresh information they have cached, significantly reducing the demand-per-client on the StarTeam Server. This improves server scalability, but it also improves client responsiveness since updates are received within seconds after they occur.

MPX Cache Agents

Messages broadcast by a MPX Message Broker benefit clients with active sessions. However, for files MPX offers an optional MPX Cache Agent process that manages its own persistent cache. MPX Cache Agents can be deployed at geographic locations, allowing clients to fetch file contents from the nearest

MPX Cache Agent, preventing the need to fetch this content across a longer (and potentially slower) network connection. MPX MPX Cache Agents are illustrated below.

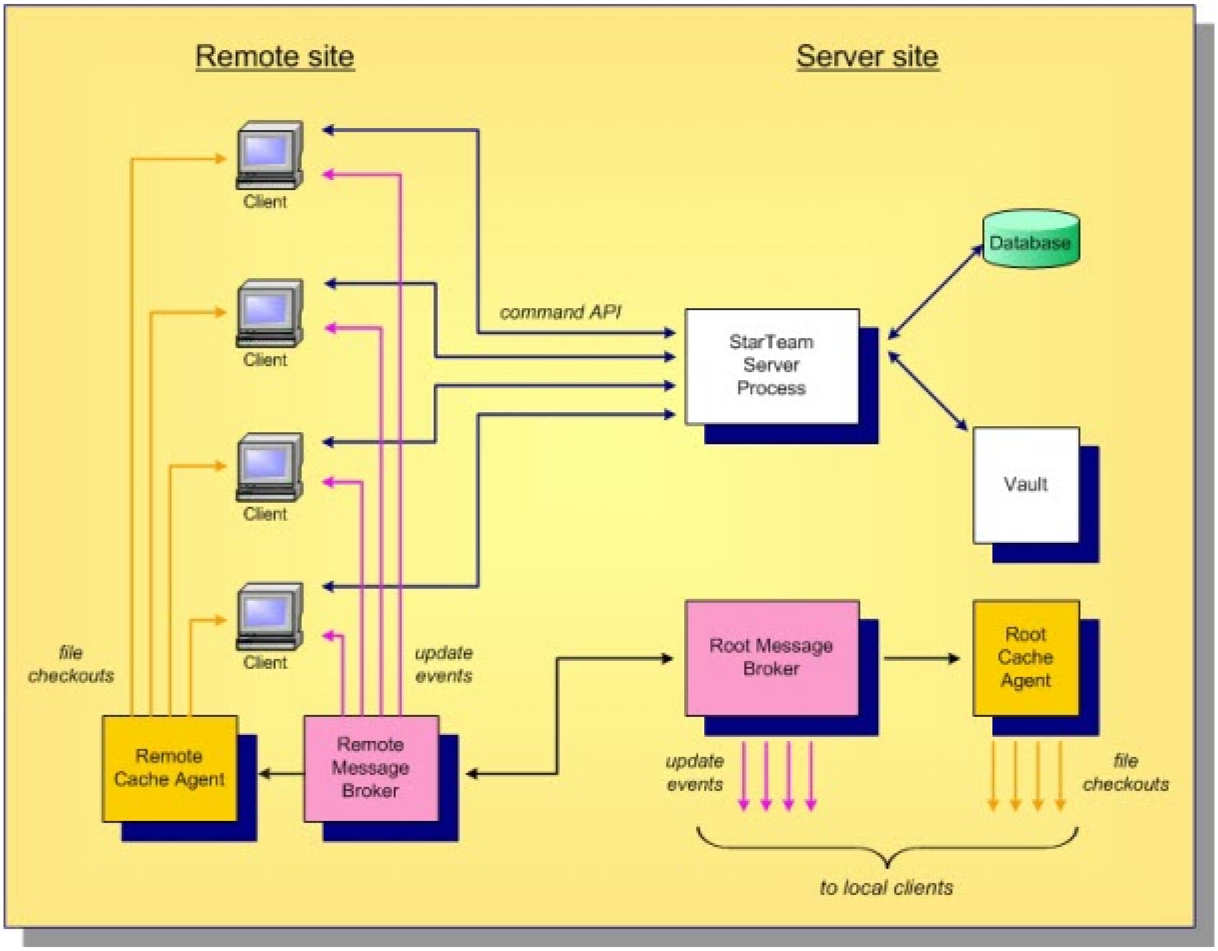

In this example, a Root MPX Cache Agent is deployed network-near to the StarTeam Server process. A Root MPX Cache Agent directly accesses the StarTeam vault, providing local clients with an alternate path to the vault for checking-out files. This reduces demand on the StarTeam Server, enhancing its scalability.

This example also shows a Remote Message Broker and a Remote MPX Cache Agent deployed at a remote site. Using *broker-to-broker forwarding*, each update event is forwarded once to the Remote Message Broker, which then broadcasts it to local clients. Files are streamed to the Remote MPX Cache Agent, which stores them in an encrypted private cache. StarTeam clients network-near to the Remote MPX Cache Agent can check out files at any time, leveraging the local high-speed network instead of pulling content across the WAN. This further reduces demand from the StarTeam Server while improving remote client responsiveness.

Other Options for Distributed Organizations

provides a unique solution for distributed teams. It leverages the benefits of a centralized server—lower total cost of ownership, better security, and simplified administration, while solving the traditional performance and scalability issues of client/server architectures. offers many advantages to distributed organizations:

- Any number of Message Brokers can be “chained” together (typically in a hub-and-spoke configuration) to form a “messaging cloud” that scales to any size organization. Message Broker limits can be configured to arbitrary values based on available resources such as file handles.
- Any number of MPX Cache Agents can be distributed globally. Clients can be configured to automatically locate and use the network-nearest MPX Cache Agent, or they can choose a specific MPX Cache Agent.

- MPX Cache Agents use *push caching* in which content is broadcast and stored by MPX Cache Agents as soon as it is created. This makes caches more effective than traditional “pull through” caching, in which every initial request results in a “cache miss”.
- MPX Cache Agents use advanced synchronization techniques that improve their effectiveness such as *pre-charging*, *tiering*, *request forwarding*, and *automatic catch-up*.

About Source Control

Source Control Basics

Each source control system consists of one or more centralized repositories and a number of clients. A repository is a database that contains not only the actual data files, but also the structure of each project you define.

Most source control systems adhere to a concept of a logical project, within which files are stored, usually in one or more tree directory structures. A source control system project might contain one or many IDE-based projects in addition to other documents and artifacts. The system also enforces its own user authentication or, very often, takes advantage of the authentication provided by the underlying operating system. Doing so allows the source control system to maintain an audit trail or snapshot of updates to each file. By storing only the differences, the source control system can keep track of all changes with minimal storage requirements. When you want to see a complete copy of your file, the system performs a merge of the differences and presents you with a unified view. At the physical level, these differences are kept in separate files until you are ready to permanently merge your updates, at which time you can perform a commit action.

This approach allows you and other team members to work in parallel, simultaneously writing code for multiple shared projects, without the danger of an individual team member's code changes overwriting another's. Source control systems, in their most basic form, protect you from code conflicts and loss of early sources. Most source control systems give you the tools to manage code files with check-in and check-out capabilities, conflict reconciliation, and reporting capabilities. Most systems do not include logic conflict reconciliation or build management capabilities.

Commonly, source control systems only allow you to compare and merge revisions for text-based files, such as source code files, HTML documents, and XML documents. StarTeam stores binary files, such as images or compiled code, in the projects you place under control. You cannot, however, compare or merge revisions of binary files. If you need to do more than store and retrieve specific revisions of these types of files, you might consider creating a manual system to keep track of the changes made to such files.

Repository Basics

Source control systems store copies of source files and difference files in some form of database repository. In some systems, such as CVS or VSS, the repository is a logical structure that consists of a set of flat files and control files. In other systems, such as StarTeam, the repositories are instances of a particular database management system (DBMS).

Repositories are typically stored on a remote server, which allows multiple users to connect, check files in and out, and perform other management tasks simultaneously.

With StarTeam, you create a server configuration to identify a repository for StarTeam projects. Each server configuration acquires its own set of projects as they are created. The Server can run any number of server configurations. Because each server configuration must use a database, you need to make sure that you establish connectivity not only with the server, but also with the database instance.

Atomic Check-ins

All check-ins in StarTeam are atomic. Whenever more than one file is checked in as the result of a single transaction all of the files, and their associated process items, are updated in a single action. If for some reason, the check-in fails, none of the files are checked in, and the status of the associated process items is not updated.

For example, suppose `User A` selects to check in all modified files in a StarTeam folder, but one of the selected files is locked by `User B`. Because of the locked file, none of the files are checked in (and none of the process items are updated as fixed) and `User A` is notified that none of the files were checked in because one of the files was locked by `User B`.

Deployment Guidelines

This section discusses high-level options for hardware deployment with StarTeam. Because StarTeam can be used by small teams, enterprise-scale organizations, and everything in between, there are many options for deploying its components that impact performance, scalability, fail-over, and other factors such as minimum hardware requirements, high availability options, and options for distributed teams.

Performance and Scalability Factors

The good news is that StarTeam is a rich application that can be used in a variety of ways. The bad news is that this flexibility makes it difficult to predict exactly what hardware configuration is perfect for your organization. Here are the major factors that affect the performance and scalability of a StarTeam configuration:

Repository Size The number of views and items affect the StarTeam Server process's memory usage, database query traffic, and other resource factors more than any other type of data. Other kinds of data such as users, groups, queries, and filters have a lesser effect on resource demand. Simply put, as the repository gets bigger, more demand is placed on server caching and database queries.

Concurrent Users The number of concurrent users during peak periods has a significant affect on a server. Each concurrent user requires a session, which maintains state, generates commands that utilize worker threads, incurs locking, and so forth. The number of defined users is not nearly as important as the number concurrent users during peak periods. If you use a single metric for capacity planning, use concurrent users.

It boosts server scalability, so whether or not you deploy it and whether or not clients enable it will affect scalability. MPX Cache Agents not only significantly boost check-out performance for remote users, but they also remove significant traffic from the server. In short, deploying will bolster your configuration's scalability.

Bulk Applications On-line users that utilize a graphical client typically incur low demand on the server. In contrast, bulk applications such as “extractors” for Datamart or Search and “synchronizers” for integrations such as Caliber or StarTeam Quality Center Synchronizer tend to send continuous streams of commands for long durations. A single bulk application can generate demand comparable to 10-20 on-line users.

Application Complexity Due to its customizability, StarTeam allows you to build sophisticated custom forms, add lots of custom fields to artifact types, create custom reports, and so forth. The more sophisticated your usage becomes, the more commands will be generated and the bigger artifacts will get, both of which increase demand.

Consider these factors when deciding the size of your configuration. Because of the unique factors that define your environment, take these deployment suggestions as guidelines only.

Configuration Size

There are no hard rules about what makes a StarTeam configuration small, medium, or large. However, for our purposes, we'll use these definitions based on concurrent users:

Small configuration	< 50 concurrent users.

Medium configuration	< 200 concurrent users.
Large configuration	> 200 concurrent users or more.

The concurrent user count, rather than data volume or type of users, seems to be the best metric for judging configuration size for purposes of deployment planning. In our experience, the amount of data managed by a StarTeam configuration (particularly items) tends to grow proportionally with the number of projects and views, which grow in proportion to the team size. Moreover, the ratio of online users to bulk applications tends to be roughly the same across organization sizes.

So how big can a configuration get? To date, we've seen single StarTeam instances with over 500 concurrent users, over 10,000 total "defined" users, over 4,000 views, tens of millions of items, and up to a terabyte of vault data. With continuous hardware advances and software improvements, these limits get pushed every year.

Note: Not all of these limits have been reached by the same configuration. Although some customers have 4,000 views, not all are actively used. A customer with 10,000 total users typically sees 250-300 concurrent users during peak periods. Interestingly, however, the amount of data managed by the vault seems to have little effect on performance or scalability.

The factors to consider as a configuration size increases are:

Start-up Time
The StarTeam Server process performs certain maintenance tasks when it starts such as purging aged audit and security records in the database. As the amount of activity and time-between-restarts increases, these tasks increase the start-up time. Also, start-up time is affected by the number of unique "share trees" due to initial caches built at start-up time. With well-tuned options, even a large server can start in a few minutes, but it can also take up to 15 minutes or more.

Memory Usage
The StarTeam Server process's memory usage is affected by several factors such as the total number of items, the server caching option settings, the number of active sessions (concurrent users), the number of active views, and the number of command threads required. Caching options can be used to manage memory usage to a point, but sessions, active views, and other run-time factors dictate a certain amount of memory usage. On a 32-bit Microsoft Windows platform, the StarTeam Server process is limited to 2 GB of virtual memory. If you enable 4 GT RAM Tuning, which boosts the virtual memory limit of a single process on a 32-bit system, this limit can be pushed closer to 3 GB. Running 32-bit on 64-bit operating system allows the process to grow up to 4 GB. Running Native 64-bit removes memory restrictions and you are constrained by the physical memory available on the server.

Tip: It is highly recommended to use the 64-bit version of the StarTeam Server for better performance and scalability.

Command Size
Some client requests return a variable response size based on the number of items requested, the number of users or groups defined, the number of labels owned by a view, and so forth. Large server configurations can cause certain commands to return large responses, which take longer to transfer, especially on slower networks. Clients will see this as reduced performance for certain operations such as opening a project or a custom form.

Multiple Configurations on the Same Server

For small- to medium-sized server configurations, you can place all StarTeam Server components on a single machine. Furthermore, you can also deploy all components for multiple configurations on the same machine depending on the sum of concurrent users of all configurations. The diagram below shows both basic and components deployed.

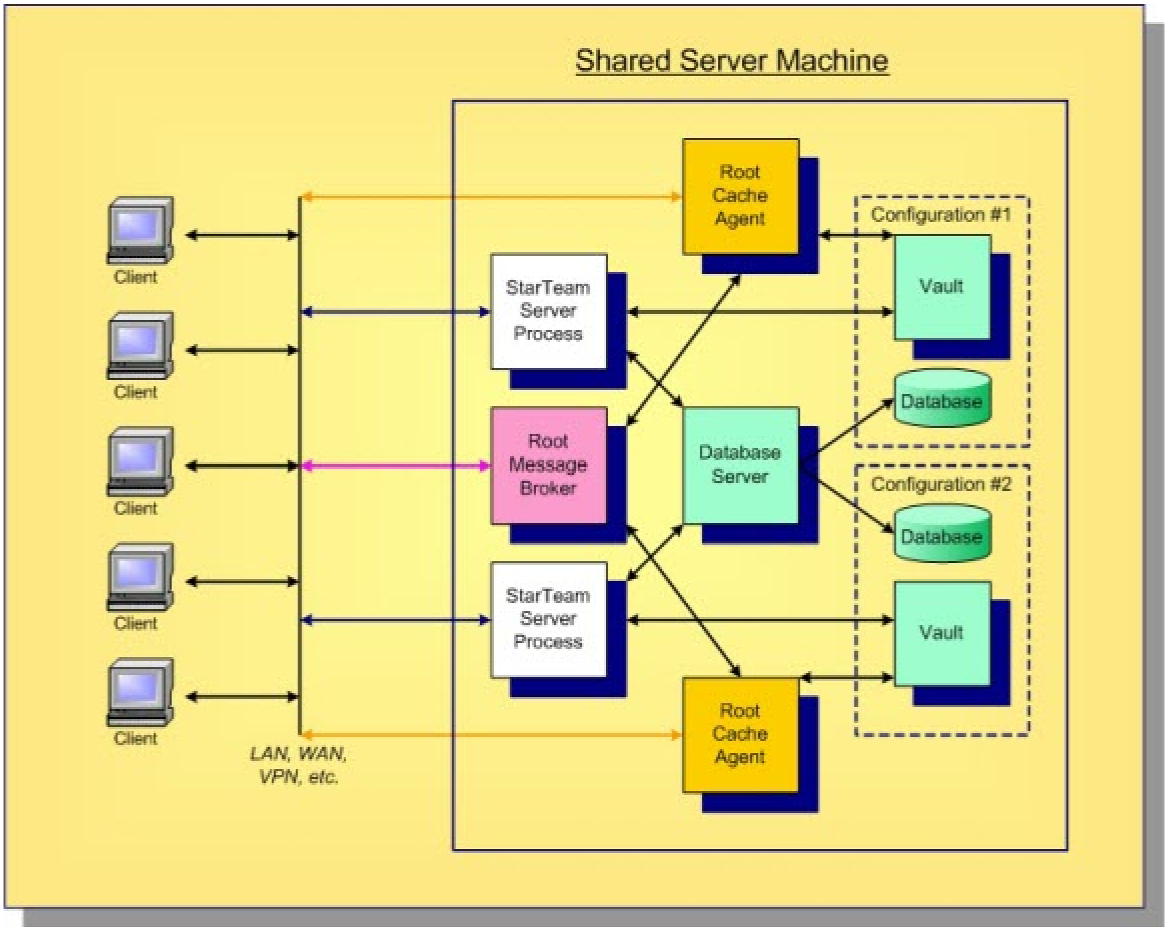

You should use a single machine for all StarTeam Server components only when the total number of concurrent users for all configurations does not exceed 100. Even though a single configuration can support more than 100 users, each configuration has a certain amount of overhead. Consequently, we recommend that when the total peak concurrent user count reaches 100, it's time to move at least one configuration to its own machine.

With a single machine, all StarTeam Server processes, the root Message Broker, root MPX Cache Agents, and the database server process execute on one machine. Here are some rules of thumb for this layout:

- Start with 2 cores and 2 GB of memory for the database server process.
- Add 2 cores and 2 GB of memory per StarTeam configuration.
- If you use locally-attached disk for each StarTeam configuration's vault and database partitions, use separate, fast drives to improve concurrency. Also, the disks should be mirrored to prevent a single point of failure.
- If you deploy , all StarTeam configurations can share a single root MPX Message Broker. Though not shown, one or more remote Message Brokers may be connected to the root Message Broker.
- If you deploy MPX Cache Agents, each configuration needs its own root MPX Cache Agent, which can share the root Message Broker. Though not shown, one or more remote MPX Cache Agents may be connected to each root MPX Cache Agent.
- Be sure to configure each StarTeam Server, Message Broker, and root MPX Cache Agent process to accept TCP/IP connections on a different port.

Using these guidelines, you can deploy three to four small StarTeam configurations on one machine, only if the total number of concurrent users doesn't peak above 100 or so. Otherwise, the various processes could begin to compete for resources (CPU, memory, disk I/O, and/or network bandwidth), adversely affecting responsiveness. Also, if you start out with the single-server configuration, don't forget to plan on moving components to their own machines when demand grows over time.

⚠ **Caution:** The disadvantage of deploying multiple configurations on a single machine is that they are all affected when the machine must be upgraded, patches need to be installed, someone kicks the power plug, and so forth.

Medium Configurations

As your configuration size grows beyond what could be called a small configuration, the first thing to move to its own machine is the database process. When you move the database process to its own machine, install a high-speed dedicated link between the StarTeam Server and database machines. Trace route between the StarTeam Server and the database machine should ideally be one hop.

Separate Database Machine

Using a separate machine for the database server, multiple StarTeam Server processes and MPX components can still be deployed on the same shared server machine. Because the database processing is offloaded to another machine, the total number of current users can be higher, up to 200-300 or so. A shared database server is shown below.

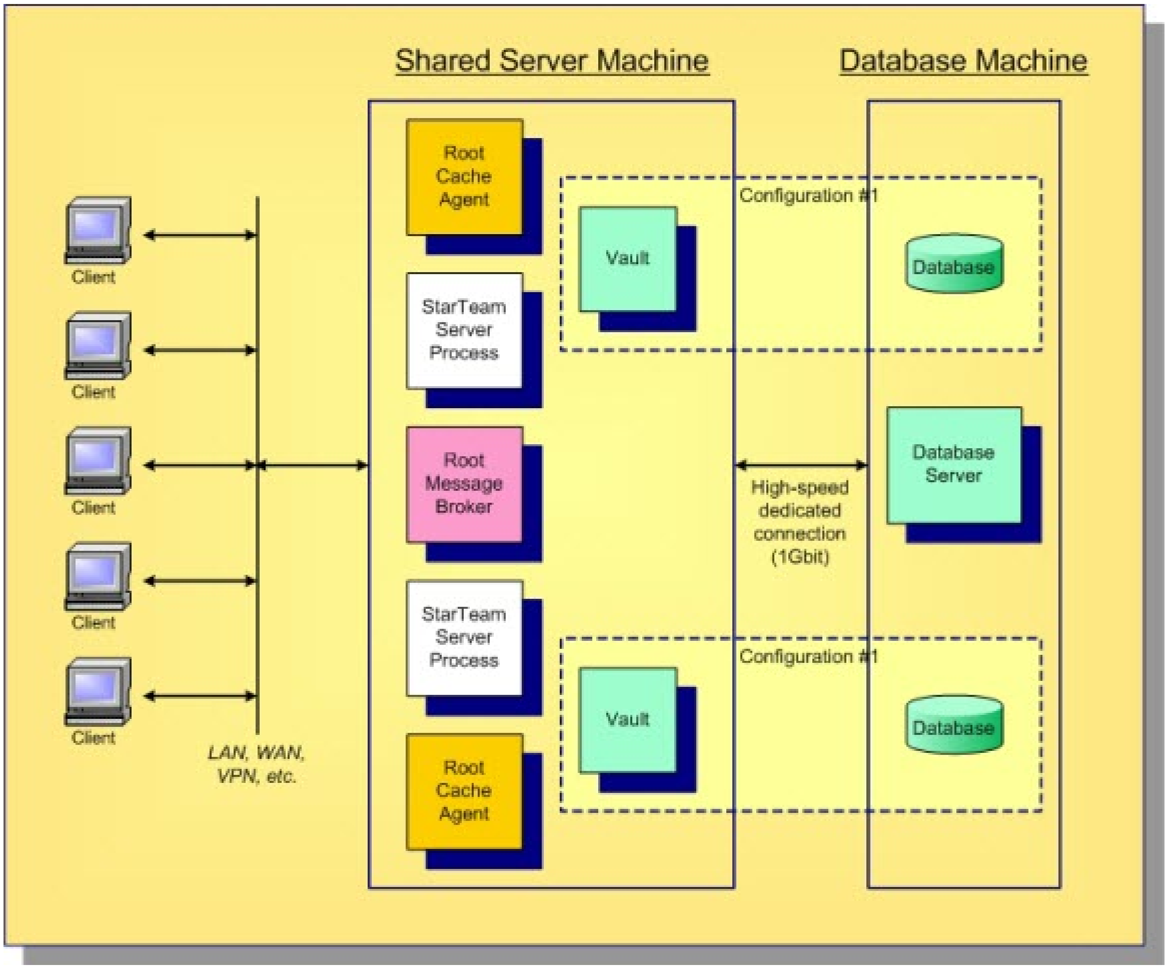

In this diagram, a locally-attached disk is assumed for the server and database machines.

Storage Server

With multiple configurations, you have multiple vaults and databases, possibly on separate disks. As you consider backup procedures, mirroring for high availability, and other administrative factors, you may find it more cost-effective to place all persistent data on a shared disk server (SAN or NFS), as shown below.

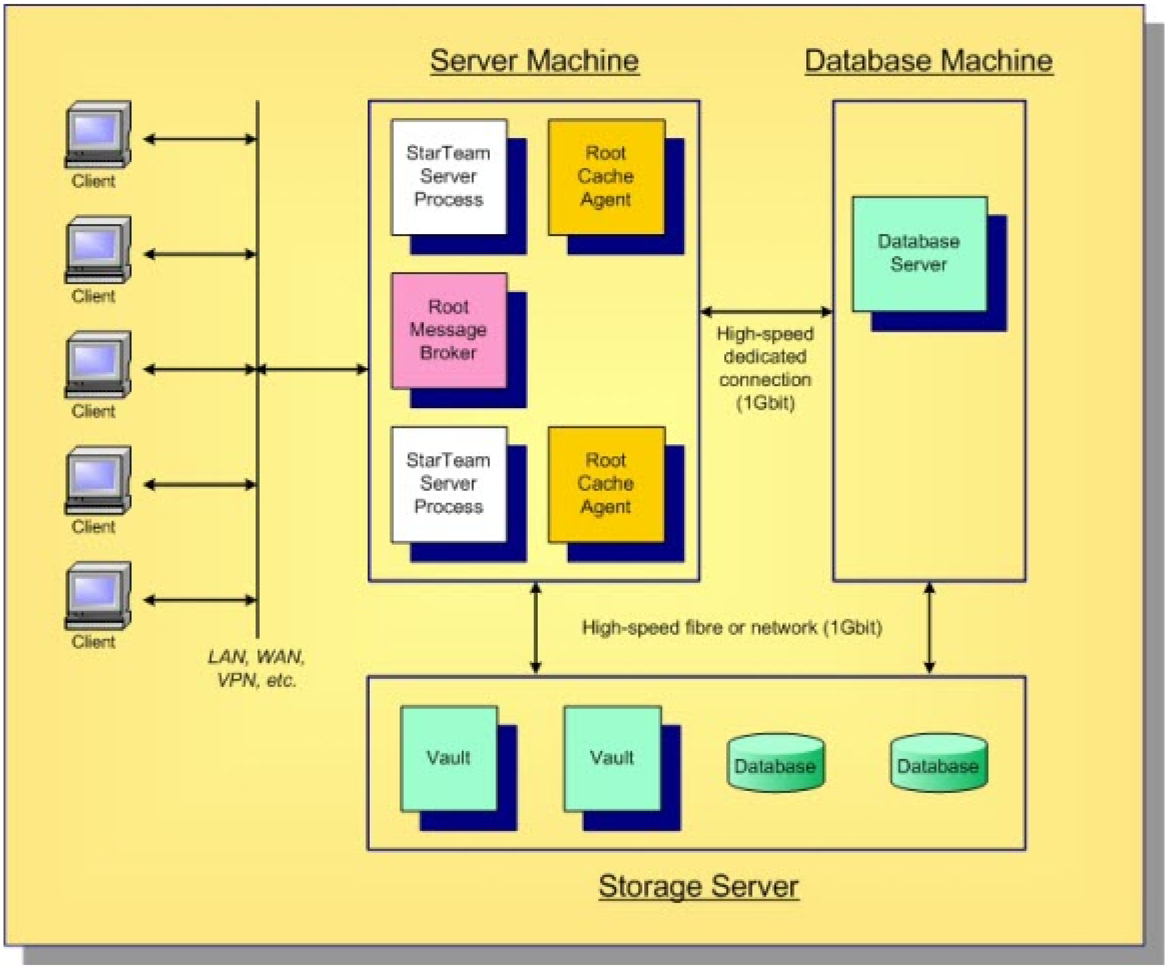

Using a shared storage server for all configuration vaults and databases has several advantages. Depending on the storage system, all important data can be backed-up with a single procedure. Hardware to support mirroring or other RAID configurations can be concentrated in a single place. Many storage systems allow additional disks to be added dynamically or failed disks to be hot-swapped.

Large Configurations

Micro Focus considers a large configuration to be one that supports 200 concurrent users or more during peak periods. For these configurations, place the StarTeam Server process on its own system. The database process should also execute on its own machine. Though not strictly necessary, the root MPX Message Broker and MPX Cache Agent processes can also benefit by executing on yet another machine. Especially when concurrent users rise to 200, 300, or more, moving the processes to their own machine can remove network traffic and other resource contention from the StarTeam Server machine. A typical deployment of multiple large configurations is shown below.

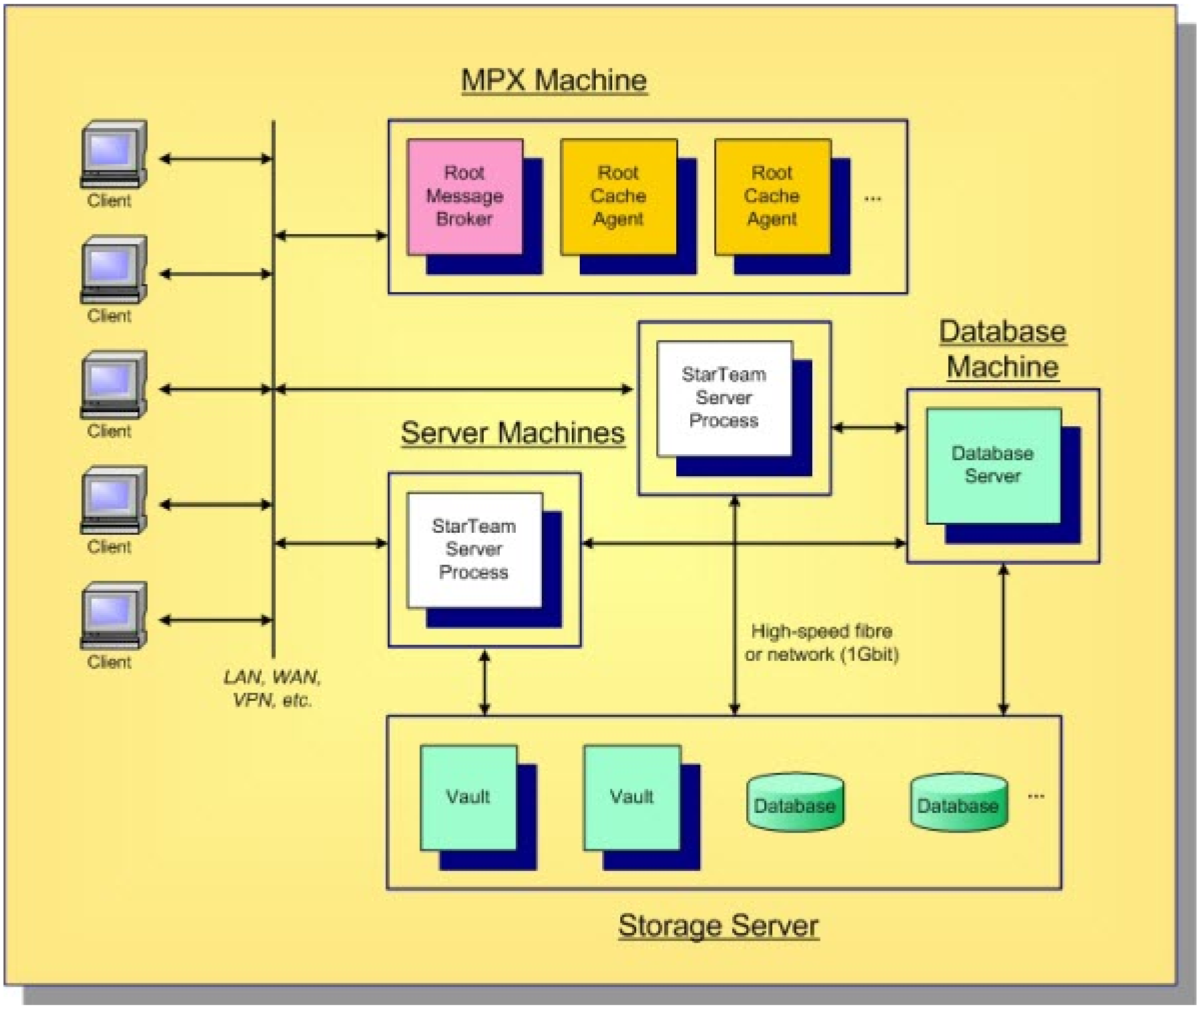

The key points of this multiple, large configuration deployment are:

- The StarTeam Server process for each configuration executes on its own machine. This is typically a high-end machine with a multi-core CPU and at least 16 GB of memory running on a 64-bit OS. If you have more than 100 concurrent users we recommend you use a 64-bit version of the StarTeam Server.
- The database server executes on its own machine. Multiple StarTeam configurations can share the same database server. (Micro Focus has seen up to eight configurations use the same database server without a performance issue.) Each StarTeam configuration uses its own “schema instance”. Each StarTeam server machine should have a high-speed dedicated connection to the database machine.
- The root MPX Message Broker and root MPX Cache Agents can all execute on a single machine. Each root MPX Cache Agent requires access to the appropriate vault, but a high-speed dedicated connection is not necessary. File access over the network (for example, using UNC paths) is sufficient. If you utilize the StarTeam Notification Agent, you can put it on the machine as well.
- A shared storage server such as a SAN server can be used for all StarTeam vaults and database partitions. Depending on the hardware, an interface (for example, “host” card) may be needed for each StarTeam server machine in order to access the SAN.

Active/Passive Clustering

StarTeam works with *active/passive clustering*, in which a “warm standby” node is maintained for quick fail-over. One general rule to remember is that only one StarTeam Server process can be active for a given configuration at one time. However, StarTeam configuration files can be copied to multiple machines along with all the necessary software. Also, multiple machines under the control of Failure Management Software (FMS) can be connected to the same database (which may be clustered itself), and they can be connected to the same shared storage server for vault access.

Active/passive clustering works like this: the StarTeam Server process on one node in the cluster is started, making it the active node for that configuration. The IP address of the active node is mapped to a virtual “cluster address”, which is the address to which clients connect. If the active node fails, the FMS takes care of fail-over: it starts the StarTeam Server process on a passive machine, making it the active node, and remaps the cluster address to the new active node’s IP address. Running clients receive a disconnect message and have to reconnect, but in most cases the fail-over will occur quickly, so clients can immediately reconnect.

When you have multiple StarTeam configurations, you can “pair” machines so that an active node for one configuration is the passive node for a second configuration and vice versa. Hence, both machines are actively used, and only in a fail-over scenario one machine must support the processing of both configurations. An example of active/passive cluster configuration is shown below.

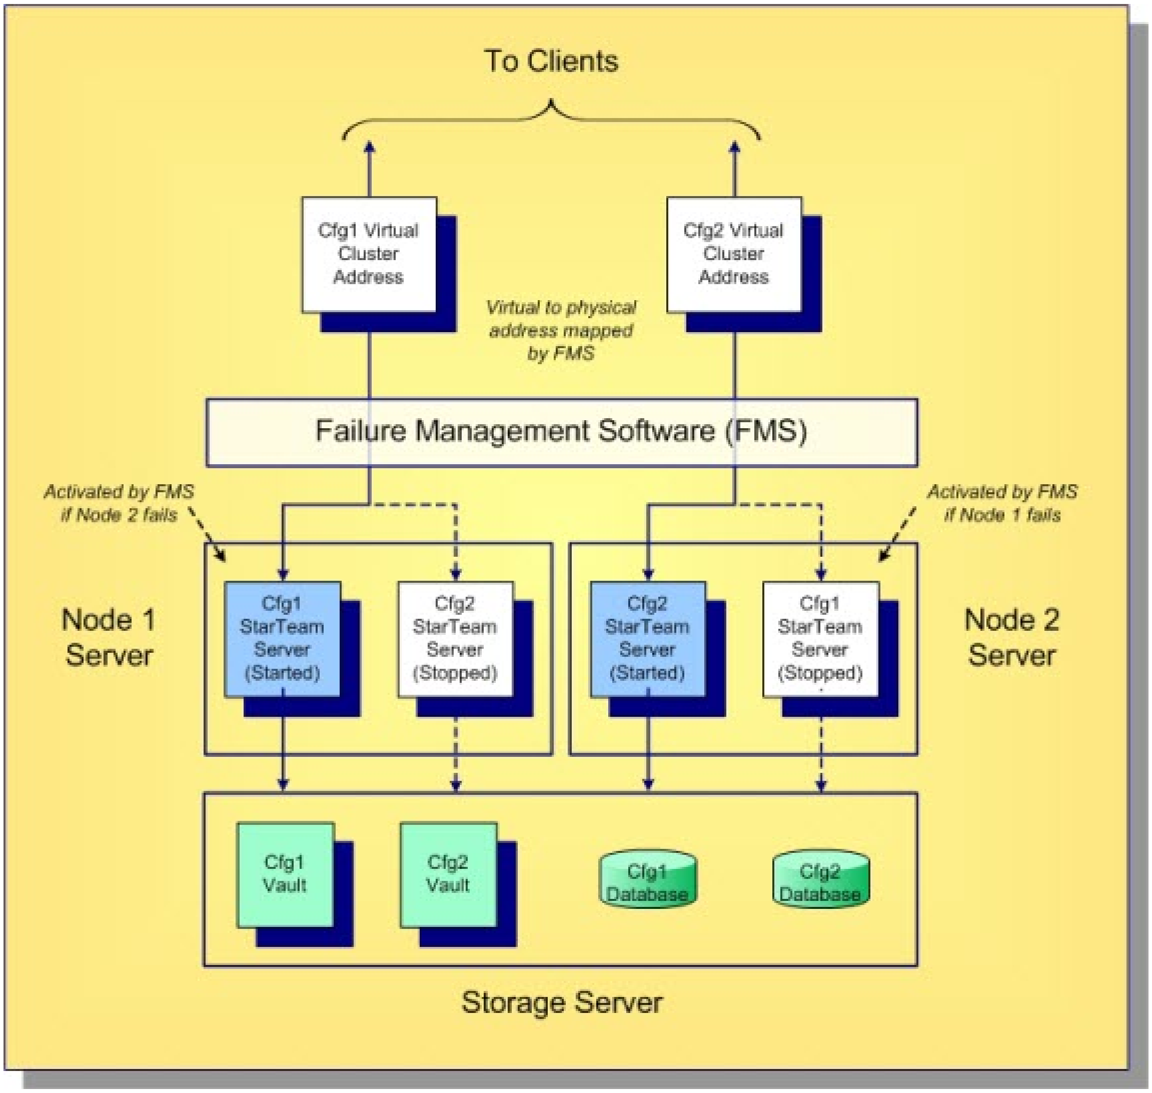
To Clients
Cfg1 Virtual Cluster Address
Cfg2 Virtual Cluster Address
Virtual to physical address mapped by FMS
Failure Management Software (FMS)
Activated by FMS if Node 2 fails
Activated by FMS if Node 1 fails
Node 1 Server
Cfg1 StarTeam Server (Started)
Cfg2 StarTeam Server (Stopped)
Cfg2 StarTeam Server (Started)
Cfg1 StarTeam Server (Stopped)
Node 2 Server
Cfg1 Vault
Cfg2 Vault
Cfg1 Database
Cfg2 Database
Storage Server

In this example, the StarTeam configurations `Cfg1` and `Cfg2` are “paired”, hence one node is active and one node is passive for each one. (The database process is not shown – it might also be deployed on a cluster.)

Server Administration

StarTeam Server Administration Overview

The StarTeam Server manages data for all its client applications. It is maintained by an administrator who is familiar with the complexities and details of the its operation. Client applications, such as the Cross-Platform Client, connect to the StarTeam Server to access data. As an administrator who initially installs the StarTeam Server, you may perform some or all of the following actions:

- Install the StarTeam Server.
- Configure the StarTeam Server.
- Register (license) the StarTeam Server.
- Create and start a new StarTeam Server configuration (an instance of the StarTeam Server).
- Set up StarTeamMPX for the new StarTeam Server configurations.
- Add new users and groups to the StarTeam Server configuration.
- Set up Directory Server and use LDAP QuickStart Manager to add users.
- Set up a password policy for non-LDAP users.
- Create projects and views for the StarTeam Server configuration.
- Set up access rights for projects.
- Enable StarTeam Server configuration diagnostics.
- Set up email notification and customize automatic email notification with your own text or HTML-based email message templates.
- Set up system policy, that is, manage passwords, logon failures, access rights, and security events for the server configuration.

The StarTeam Server creates new projects with only the `File` type pre-selected as a default for new views. Users can still change the project properties after the project is created, and they can change the item types included for any given new view. However, if the user changes nothing, by default new views will only include files when they are created.

A StarTeam Server can manage any number of projects. Each project has one root view and any number of child views. The root view and every child view has one application folder as a root folder. An application root folder can have any hierarchy of child folders. This is called the folder hierarchy. When an administrator creates a project, that project’s root view and the root view’s root folder are created automatically and given the same name as the project. For example, if the project’s name is `Great App`, the root view’s name is initially `Great App`, and the root folder’s name is initially `Great App` (although the administrator can change these names).

Your first task as an administrator is to install, configure, and register the StarTeam Server, as explained in the *StarTeam Installation Guide*. Next, you must create an instance (known as a StarTeam Server configuration) on the computer on which the StarTeam Server is installed. A StarTeam Server configuration must be running before you and your team members can access the application.

Server Administrator Assumptions

We assume that StarTeam Server administrators are familiar with:

- Creating and modifying relational databases.
- Working with the features of their operating system, such as creating files, running executable files, and managing access rights.
- Basic software configuration management concepts.

We also assume that StarTeam Server administrators will never modify database contents or vault files other than through a client or **Server Administration** tool.

Warning: Direct database manipulation is not supported.

StarTeam Server Configuration Sample Data

StarTeam provides sample configuration data that you can download from within the **Server Administration** tool. It contains a Visual C++ sample application and related materials. It has sample files, change requests, topics, tasks, and it also includes a `StarFlow Extension` project. You can use the sample repository to experiment with and learn more about StarTeam.

The StarTeam Server creates new projects with only the file type pre-selected as a default for new views. You can change the project properties after the project is created, and change the item types included for any given new view. However, if you do not change anything, by default new views will only include files when they are created.

Licensing the Server

Setting Up License Servers

You have a choice between using the Borland License Server (BLS) and the native StarTeam licensing found in this and earlier releases. If you use the license server, users must use their network logon names as their StarTeam user names. This section explains the steps the administrator follows to set up a license server.

1. As the StarTeam administrator, you should receive licensing information from Borland via email (a sales representative should put this in motion).
2. Install the license server (the license server documentation explains how to do this).
3. Save the license files (this involves accessing a Micro Focus web site and downloading Micro Focus license files called slips).
4. Place the slip files in the `/License` folder, a subdirectory of the `StarTeam Server <version>` installation folder.
5. Configure the license server for users (this is covered in the license server documentation).
6. Use the Server Administration tool to:
   - Change user names to network logon names.
   - Assign users to specific licenses by setting the license options.

When StarTeam Server starts, it checks for slips and stores information about them in memory. It does not recognize new slips until the next restart.

When a user logs in from a StarTeam client, the StarTeam Server tells the client what features are available to its user based on the license assigned to that user. If the user is assigned a license from a slip, but that slip is no longer in the `/License` folder, StarTeam Server displays an error message. If the user license type is **Unassigned**, the user is not logged on and StarTeam Server returns an exception.

Note: If you are using a license server, concurrent licenses are released immediately by StarTeam Server, but the license server might not find that out for a few minutes. StarTeam Server updates the license server about license usage at an interval specified in the licensing slip. The license server will know that a license has been released only when the next update for that license occurs.

Using Evaluation Licenses

The first time you run StarTeam Server, an evaluation license is created for StarTeam Enterprise Advantage, which is the edition of StarTeam with the largest feature set.

Before the 30-day product review period expires, be sure to register the product or extend the evaluation period. Otherwise, when clients access a server configuration managed by a StarTeam Server that has expired, no components (such as the **File** or **Change Request** components) are available and in the StarTeam Cross-Platform Client, the upper and lower panes have no tabs.

1. Obtain an evaluation extender key by contacting Borland at *http://www.borland.com/products/starteam/try/StarTeamwin.aspx*.
2. Click **Start** > **Programs** > **Micro Focus** > **StarTeam Server <version>** > **StarTeam Server** . The Server Administration Tool opens.
3. Choose **Help** > **About** . The **About StarTeam Server** dialog box appears.
4. Select the **License** item in the left pane.
5. Click **Extend Evaluation**. The **Extend Evaluation** dialog box appears.
6. Type the evaluation key.
7. Click **OK**.

Using Native Licenses

The first time you run StarTeam Server, an evaluation license is created for StarTeam Enterprise Advantage, which is the edition of StarTeam with the largest feature set. Before the 30-day product review period expires, be sure to register the product or extend the evaluation period.

Note: You cannot license StarTeam Server while any of its server configurations are running as a Microsoft Windows service. If you change the registered license while a StarTeam project is open on a user's workstation, the licensing takes effect for that user by closing and reopening the project window. If you license StarTeam Server as Enterprise after using an evaluation license which is for the Enterprise Advantage edition, the feature set changes. For example, if you created requirements during the evaluation and license the server as anything other than Enterprise Advantage, the requirements tab disappears.

Registering a Native License Using the Server Administration Tool

1. Click **Start** > **Programs** > **Micro Focus** > **StarTeam Server <version>** > **StarTeam Server** . The Server Administration Tool opens.
2. Choose **Help** > **About** . The **About StarTeam Server** dialog box appears.
3. Select the **License** item in the left pane.
4. If you have yet to enter a license, you must delete the evaluation key by selecting it from the right pane of the dialog and clicking **Delete**.
5. Click **Register**. The **Server Registration** dialog box opens.
6. Type the correct numbers in the **Serial Number** and **Access Key** fields.

   Note: Serial numbers are case-sensitive, access keys are not.
7. Click **OK**.

Registering a Native License at the Command Prompt

1. Open a command prompt, and navigate to the home installation folder for StarTeam Server. For example:

   ```
   C:\Program Files\Micro Focus\StarTeam Server <version>
   ```
2. At a command prompt type:

   ```
   starteamserver -serial number -access key
   ```

Saving License (.slip) Files

After you receive licensing information in a license certificate email, you need to install the license server and host the licenses. This involves accessing a Micro Focus web site and downloading license files called `.slip` files as described here.

1. From the Micro Focus web site using the link in the license certificate email, download all of the `ZIP` files containing the `.slip` files per the instructions provided on the web.
2. Copy each `concurrent_.slip`/`named_.slp` file into your `C:\Program Files\Micro Focus\StarTeam Server <version>\License` directory.
3. Copy each `server_.slip` file to `C:\Micro Focus\BLS4\conf`.

   Note: `BLS4` folder name might change depending on the version of License Server that you are using. Also, if you are using FLEXlm instead of Borland License Server, the files might need to be installed elsewhere. Check your FLEXlm server documentation.

When the StarTeam Server starts, it checks for slips and stores information about them in memory. It does not recognize new slips until the next restart.

Once the slips have been placed in the `\License` folder and the StarTeam Server has been restarted, the **User Manager** in the Server Administration Tool can display information about the slips and an administrator can assign licenses from those slips to users.

Assigning Licenses to Users

To be able to work with StarTeam, users must have a named user license or a concurrent license. By default, users are assigned to use concurrent licensing. The StarTeam Server Administrator uses the **User Manager** in the Server Administration Tool to assign licenses to existing users or to new users.

From the **User Manager**, you can use the **User Properties** dialog box or context menu to assign licenses. Generally, use the context menu for bulk actions and the **User Properties** dialog box for assigning licenses to users one at a time. Licenses can also be assigned using LDAP Quickstart Manager (usually done in bulk but can be done one-by-one). Which of these you choose depends on what you are doing as illustrated in these scenarios:

- If you are adding a new user and filling in all the data about him/her, you can assign a license as part of the process. This would most likely be done in the **New User Properties** dialog box.
- If you are upgrading and need to assign a group of existing users to a new license slip for the new release, you can multi-select the appropriate users and assign them to a slip simultaneously. You would do this from the context menu.
- If you have been evaluating StarTeam and now have purchased native licenses or licenses to be used with a license server, you can select all the existing users (from the evaluation period) from your production server configuration and assign them to a license type or a license slip. You would do this from the context menu.
- If a group of people have been laid off and you no longer want them to use StarTeam, you can select them all and change their license type to **Unassigned**.

Note: The named and concurrent user licenses are the same as the licenses in earlier StarTeam releases.

Assigning Licenses to Existing Users

1. Click **Start** > **Programs** > **Micro Focus** > **StarTeam Server <version>** > **StarTeam Server** . The Server Administration Tool opens.
2. Select a server configuration from the list of servers. If you have not yet logged on, you will be asked to do so.

3. Click the **Accounts** bar and then click (**User Manager**). The **User Manager** tab opens.
4. Select one or more users.
5. Right-click to display the context menu and choose **Properties** . The **User Properties** dialog box opens.
6. Select the license type from the **License** list:

    (optionally) The license number of one or more license server slip files for either a named or concurrent license.

    StarTeam Named The user has a particular license assigned to them.

    StarTeam Concurrent The user is assigned one of the “floating” licenses when they log on to StarTeam.

    Unassigned Select this “license type” when a user has no license.

7. Type the rest of the data on the **General** and other tabs as appropriate. Remember to use the network logon name for the **User Name** field on the **Logon** tab.
8. Click **OK**.

Note: The status line at the bottom of the **User Manager** tab provides licensing statistics including the number of named user licenses that are currently available.

Assigning a License to a New User

1. Click **Start** > **Programs** > **Micro Focus** > **StarTeam Server <version>** > **StarTeam Server** . The Server Administration Tool opens.
2. Select a server configuration from the list of servers. If you have not yet logged on, you will be asked to do so.
3. Click the **Accounts** bar and then click (**User Manager**). The **User Manager** tab opens.
4. Click **New User**. The **New User Properties** dialog box opens.
5. Select the license type from the **License** list:

    (optionally) The license number of one or more license server slip files for either a named or concurrent license.

    StarTeam Named The user has a particular license assigned to them.

    StarTeam Concurrent The user is assigned one of the “floating” licenses when they log on to StarTeam.

    Unassigned Select this “license type” when a user has no license.

6. Type the rest of the data on the **General** and other tabs as appropriate. Remember to use the network logon name for the **User Name** field on the **Logon** tab.
7. Click **OK**.

Managing Named User Licenses

Users can have either named user or concurrent licenses. A named user license (formerly called a fixed license) can be used only by the user who has been assigned that license whereas concurrent license users share the licenses and can log on as long as there are concurrent licenses available. Users who receive the named user licenses are guaranteed access to the StarTeam Server.

You can add as many users as you choose, but access to the server is granted only to users with named user licenses or to users who receive concurrent licenses as they log on. If you have named user licenses, you must assign them to specific users in the Server Administration Tool **User Manager**. An anchor appears before the name of users with named user licenses. Before assigning named licenses, you must add the users.

The StarTeam Server Administrator is automatically assigned a named user license which cannot be removed. This free license is not counted against the number of named user licenses you have available. After the server is licensed, named-user licenses can be assigned.

Tip: The **User Manager** status bar indicates how many named user licenses and how many concurrent licenses are in use.

1. Click **Start** > **Programs** > **Micro Focus** > **StarTeam Server <version>** > **StarTeam Server** . The Server Administration Tool opens.
2. Select a server configuration from the list of servers. If you have not yet logged on, you will be asked to do so.
3. Click the **Accounts** bar and then click (**User Manager**). The **User Manager** tab opens.
4. Select the user to whom a named user license will be assigned.

   Note: If the user is not displayed, you might need to select a particular group, or select the **Show Users in All Descendant Groups** check box.
5. Right-click and choose **Assign License** > **Add Named User License** .

Note: If you have downloaded named or concurrent license files from the license server, the context menu contains the license number for each file.

After a named user license is assigned to a user, an anchor appears before the name of the user.

Note: When you change the type of license a user has, the change does not take effect until the user logs on the next time. To make the license change effective immediately, you need to force-logout the users affected by the change.

Note: To remove a named user license, go to the **User Manager**, select a user, and choose **Assign License** > **Remove Named User License** . Removing a user named license automatically changes the user to a concurrent use license.

Working with Server Configurations

About Server Configurations

Server Configuration Overview

Before using StarTeam Server, you must decide what database to use and where to store the database and file revisions. Then you must create at least one StarTeam Server configuration (an instance of the StarTeam Server). This topic discusses StarTeam Server configurations and their storage *hives*.

StarTeam Server Configurations

A StarTeam Server configuration defines:

- The set of options, including endpoints (the TCP/IP port) and encryption levels, used for server access.
- Location of the database that stores project data and other related information.
- Locations for the repository and repository-related folders.

Any number of projects can be stored in the database associated with a particular server configuration. However, the database must be configured properly to store the amount of data produced by those projects. For more information about specific databases supported by StarTeam Server, refer to the *StarTeam Release Notes*.

You can create a server configuration by using the **Server Administration** tool. A server configuration defines a specific database as the repository for its data. To prevent corruption, that database can be associated with only one server configuration. However, that database can be used by other applications. The application stores all projects on the Server, which may contain numerous server configurations.

To access an existing project, you must first add its server configuration to your system. The StarTeam Server can be accessed from any of its clients. Each client must have a user name and the correct access rights to access the selected server configuration. Your company or team may store its data on several server configurations on one or more computers. Any of these configurations can be accessed from a number of clients.

More than one instance of the StarTeam Server may be running on the same computer. For example, you might run one server configuration with a sample project and another with a software development project, both on the same computer. Each server configuration has a different name and a different port or endpoint for each protocol. When a configuration is in use, another session using that configuration cannot be started.

Before creating a server configuration, you need to decide upon a unique name for the configuration. This name is case insensitive and cannot contain : \ /, but can contain blanks or apostrophes (').

The StarTeam Server places server log files in the location designated as the server configuration's repository path. When you first start a new server configuration, the StarTeam Server creates the `Attachments` folder, `HiveIndex`, and other folders in the same location. These folders are maintained by the StarTeam Server. Do not delete them.

Tip: Once you have created a server configuration, you can change the path to the `Attachments` folder from the **Server Administration** tool's **Configure Server** tab.

Other server configuration settings control where, when, how, and by whom the data is accessed. Some initial settings that you provide for the server configuration are properties that are necessary to start it. For example, if the user name and password that allow StarTeam Server to access the database are not accurate, StarTeam Server cannot run. Before starting StarTeam Server, you can change these properties to meet your requirements.

Native-II Vaults/Hives

Native-II is a vault architecture that provides greater scalability for all server configurations created with StarTeam and for server configurations converted to Native-II vault format with StarTeam. Server configurations have one or more *hives*. A hive is a logical disk container of files that includes an *Archive area* and a *Cache area*. The archive area consists of a folder tree in which unique file revisions are stored. The cache area consists of a folder tree that stores uncompressed file revisions on a temporary basis. Hives can hold an unlimited number of files, providing increased storage capacity, larger file revisions, more locations to store archives, and faster, more efficient performance. A single server configuration can have several hives, each of which has its own archive and cache path.

Note: StarTeam supports only the Native-II vault format for hives.

The initial hive used for storage of the server configuration’s archive files is created along with the server configuration. You must supply an archive path and a cache path to this hive when creating the server configuration. The default paths are `repository_path\DefaultHive\Archives` and `repository_path\DefaultHive\Cache`. If desired, the location of these paths can be changed later by using the **Hive Manager** dialog found in the **Server Administration** tool.

Native-II vaults store each file revision in its entirety (even though the archive file may be compressed). But the revisions can be spread over many volumes by the use of hives for storage. If one hive fills up, you can add another, without changing any data locations or moving any archive files. When a server configuration has multiple hives, StarTeam adds files to each hive in turn before reusing the first hive's archive path.

When you create a server configuration, it automatically has at least one hive (either the default or a custom hive). To increase the amount of available space for a server configuration, you can add one or more new hives with the **Hive Manager**. You can create hives while the server configuration is running, because the configuration already has an initial path, if only to a `Default Hive` in the repository path.

The main purpose of the **Hive Manager** is to create new hives for an existing StarTeam configuration, to increase the amount of available space.

Server Configuration Guidelines

In terms of initial planning, one of the most important decisions your organization must make is how many StarTeam configurations it will use. While distributing projects across multiple StarTeam Servers will increase administrative costs, it will also increase project independence and improve performance and availability. By estimating project growth and considering interdependencies ahead of time, you can avoid having to split up a configuration that has become too large. Below are some strategies to consider when developing the server deployment plan for your organization.

Advantages of Shared Server Configurations

Transactional integrity	Because a configuration uses a single database, all data within the same configuration is transactionally consistent. That is, a configuration represents a data consistency boundary. If you backup and later restore a configuration, all information within the configuration will be restored to the same point in time.
Linking	Items in the same configuration can be linked, even if they are in different projects. StarTeam does not allow cross-configuration linking.
Sharing and moving	An item can be shared or moved to any folder, view, or project within the same configuration. Moving or sharing items across configuration boundaries is not allowed.
Administrative simplicity	Administrative tasks such as adding users and groups, applying security, performing backups, and so forth are done at the configuration level.
Shared customizations	Many StarTeam resources such as filters, queries, custom forms, and work-flows can be defined at the configuration level and shared by all projects. (However, custom forms and workflow can also be customized per project or per view.)
Shared server components	All data in the same configuration utilize a single server process, database, vault, and root MPX Cache Agent. New projects can be added dynamically without adding any new server-side components.

Advantages of Separate Server Configurations

Performance	Larger configurations take longer to start, use more resources, and tend to return larger command responses. Conversely, smaller configurations have less data and fewer concurrent users, so they tend to perform better in these regards.
Managing growth	Even if you initially place multiple configurations on a single machine, you can easily move a configuration to its own machine if you need to.
Maintenance schedules	Separate configurations can be independently started and stopped for installing patches, upgrading hardware, etc. When a configuration is offline, all projects it contains are unavailable.
Custom fields	Custom fields are added at the “type” level, which has configuration-level scope. This means that if you add a custom field to a CR, all CRs in that configuration will have a value for that field. Hence, if different teams or business units have competing interests in custom fields, this argues for placing their projects in separate configurations.

Other Server Configuration Considerations

The next sections describe additional factors to consider when developing the server deployment plan for your organization.

Business Unit Divisions

When multiple business units require their own StarTeam projects, it often works well to define StarTeam Servers along organizational boundaries. That is, deploy a separate StarTeam Server for each major business unit or department, allowing each to access its own projects. Dividing along business unit lines isolates separate (and sometimes competing) requirements for security, backup processes, and other administrative issues. Separate servers can also help mitigate ownership or “turf” issues.

Where development life-cycle processes cross server configurations, clients can open multiple projects in a single StarTeam client. “Deploying” interrelated artifacts from one project to another can also be used to address cross-configuration integration needs.

Leverage StarTeam Support for Distributed Teams

Team members that require access to the same artifacts should share a single StarTeam Server. Dividing a StarTeam Server solely due to geographically dispersed teams is not necessary. StarTeam was designed to work well with distributed teams. StarTeam emphasizes a centralized configuration approach with publish/subscribe messaging and MPX Cache Agents to support distributed teams.

Avoid Partitions for Internal/External Access

In many situations, teams both behind and outside the corporate firewall require access to the same StarTeam configuration. A common practice in this scenario is to deploy the StarTeam Server process in the DMZ area of the firewall, placing the database server and storage server behind the firewall. Depending on the capabilities of the firewall, it may be appropriate to configure a dedicated port to the StarTeam Server. Alternatively, you can install two network interface cards (NICs) on the StarTeam Server machine: one “outward” facing and one “inward” facing. In this scenario, StarTeam allows specific inbound IP addresses (or address ranges) to be configured with different connection security requirements.

StarTeam provides SSL-like encryption for the command API, preventing eavesdropping on client/server traffic. All MPX Message Broker and MPX Cache Agent traffic is also encrypted, making data private across public links. To limit access to specific teams, you can use reference views or StarTeam’s security ACLs to limit access to specific projects, views, folders, and even individual artifacts. Other security features, such as strong password management and automatic account lockouts, further increase the viability of using the same StarTeam configuration for both internal and external users.

Plan for Growth

In planning how many StarTeam configurations to create, take a long-term view: at least three to five years. If you can estimate concurrent user usage, this is the best metric for capacity planning. On today’s hardware, StarTeam readily supports up to 300 concurrent users. Some customers have configurations that peak at over 400 concurrent users, and one customer has seen peaks of 600 concurrent users. But at these concurrency levels, the application types become important (that is, batch applications tend to demand more than online clients). Even a 300-concurrent user load may drive down responsiveness unacceptably if a substantial number of users are running high-demand applications.

Another way to gauge configuration scalability is with command rates. You can measure the command rates of an existing configuration by using the server trace functionality. The StarTeam Server can be tuned to provide adequate performance with command rates from 200,000 to 300,000 commands per hour (56 to 83 commands per second). Command rates of 400,000 per hour (111 per second) or more with adequate performance have been observed with good network infrastructure (low latency). Attempts to drive a single configuration higher than this tend to produce unacceptable response times.

If you cannot project user concurrency rates or command rates, you can use “defined” users, but the server load is less predictable using defined users alone. In geographically-distributed user communities, we typically see a defined-to-concurrent ratio around 10:1. So, we would expect 1,000 named users to yield about 100 concurrent user sessions during peak periods. In less-distributed topologies, where users are concentrated in one or two time zones, we expect the defined-to-concurrent ratio to be closer to 5:1. If you don’t have better data, use these approximations to estimate your peak concurrent user rate.

After estimating your three-to-five year projection, you should have an idea of how many StarTeam configurations will be needed to support your user community.

Project Structure

An instance of the StarTeam Server controls the storage of your files. Each StarTeam Server instance runs a server configuration. Here's an overview of the project structure controlled by an instance of StarTeam Server.

Server A server is a computer running the StarTeam Server software. StarDisk enables you to connect to the server. The StarTeam Server controls the repository, which is a storage place for file revision archives, and a database that contains information about files, such as their descriptions, the number of revisions, and so on.

Project A project is a way to group all the materials needed to accomplish some goal. Large, complex projects have many folders and files that are worked on by many team members. A project is the collection and organization of all these files and folders. A project might contain the files that comprise a software program, a technical publication, a legal case, a financial forecast, a building, an aircraft, or anything involving numerous files, each of which may undergo many revisions as the job progresses.

View A view, also called a project view, is a way of looking at a project. It enables users to see the parts of the project they need to see, without the confusion of seeing the entire project. Users might use several different views of a single project, or views of several different projects, depending on the files they must use to do their work. Each project has only one root view, which is created automatically when the project is created. The root view may have several child views, each of which may have several child views of their own. A view that has child views can be referred to as a parent view.

Folder Each view has one root folder. That folder can have any hierarchy of folders. Usually those folders have names that indicate their contents, such as `Marketing Materials`, `Product Documentation`, and `Source Code`.

Below are some diagrams illustrating how all these pieces fit and work together.

Server-level Hierarchy

The server can manage any number of projects.

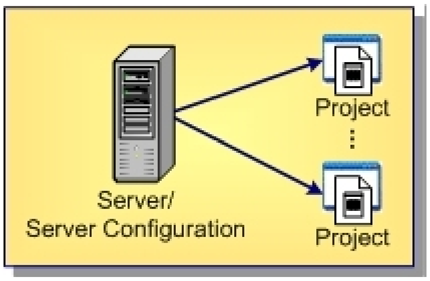

Project-level Hierarchy

Each project has one root view and any number of child views.

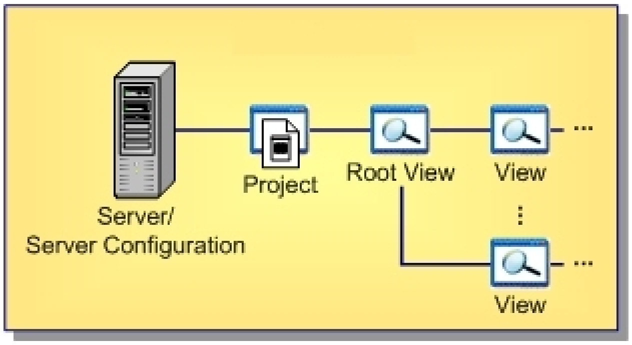

View-level Hierarchy

The root view and every child view has one application folder as a root folder.

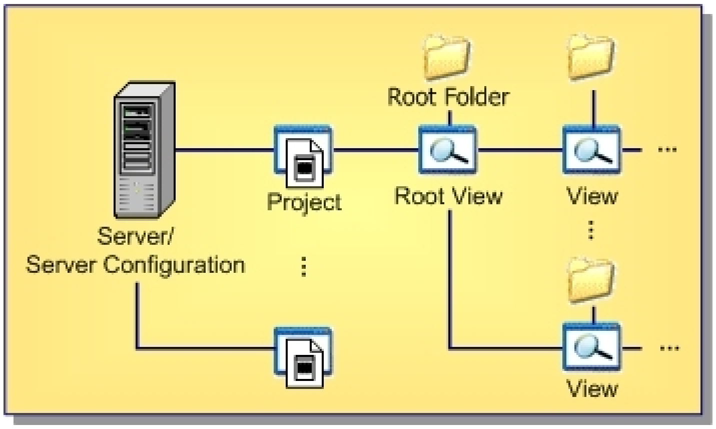

Folder-level Hierarchy

An application root folder can have any hierarchy of child folders. This is called the folder hierarchy.

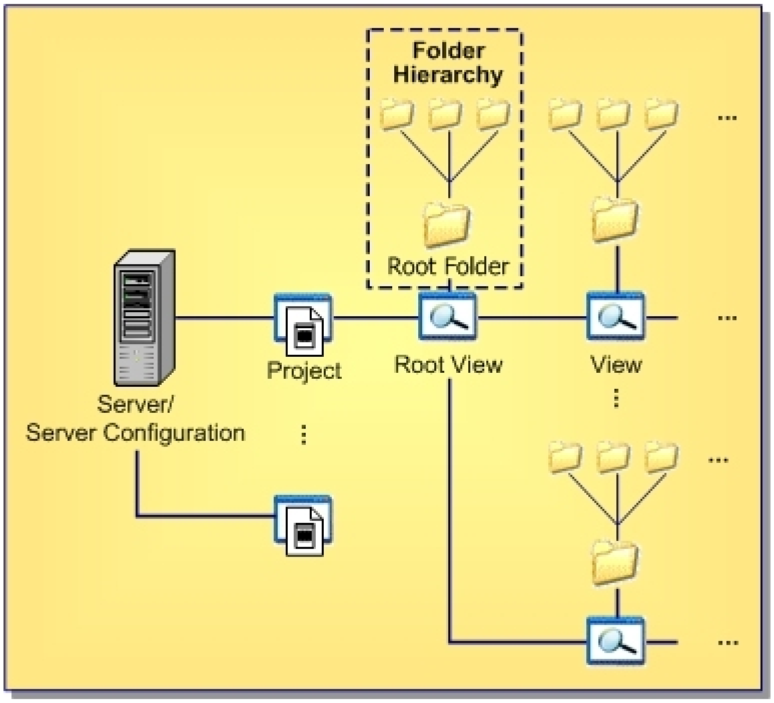

Using a Test Server

A simple but often overlooked measure you can take to smooth out administrative operations in your environment is to deploy a StarTeam Server configuration as a test server. Your test server can use lower-cost hardware than your production server, but it should be capable of running on a backup copy of your production server. With this capability, your test server can provide many useful benefits, including:

- You can test new SDK applications, workflow rules, release procedures, and so forth on the test server without fear of unwanted side-effects to your production server.
- You can use the test server to stage new releases of StarTeam and simulate upgrade and migrate operations before applying them to the production server.
- You can use the test server for training new developers and administrators.
- You can test backup and recovery procedures for your organization. Once you are sure your emergency procedures are functional, you can use the test server as a backup machine in the event of a catastrophic failure to the production machine.

Opening the Server Administration Tool

Before you use the **Server Administration** tool to administer a server configuration, you must have administrative privileges for that configuration and the configuration must be running. You can start the **Server Administration** tool from the command prompt or from the **Start** menu

The **Server Administration** tool can be used to manage server configurations running on the computer on which it is installed or multiple computers running the StarTeam Server. Connection information for server configurations is stored in the `starteam-servers.xml` file.

Starting the Server Administration Tool from the Start Menu

1. Click **Start** > **Programs** > **Micro Focus** > **StarTeam Server <version>** > **StarTeam Server** .
2. If you have installed the Server Administration tool with the StarTeam Cross-Platform Client, select **Start** > **Programs** > **Micro Focus** > **StarTeam Cross-Platform Client <version>** > **Server Administration** . This is available with custom installations only.

These actions run the `AdminTool.stjava` file opening the **Server Administration** tool. The **Server Administration** tool on the StarTeam Cross-Platform Client is similar to that which you run with the StarTeam Server except that it can be used to administer remote servers only. Some functions, such as migrating a database, can be performed only from the Server Administration tool which is installed with the StarTeam Server and only when a server is shut down.

Starting the Server Administration Utility from a Command Prompt

1. Open a command prompt window.
2. Change directories to the StarTeam Server folder, for example:
   ```
   cd C:\Program Files\Micro Focus\StarTeam Server <version>
   ```
3. Type the following at the command line:
   ```
   serveradministration
   ```
   The Server Administration tool opens.

Server Administration Tool

Server Administration Tool UI

When you need to administer your server configurations, you use the **Server Administration** tool. The **Server Administration** tool is a Java application that enables administrators to create and manage StarTeam Server configurations and the repositories they access. It is automatically installed with the

StarTeam Server and can be run only from a computer on which StarTeam Server resides. From the StarTeam Server, this tool can administer both local and remote configurations, as it can access the `starteam-server-configs.xml` file.

If you choose a custom installation, you can also install this tool with the StarTeam Cross-Platform Client. However, from the client installation, the StarTeam Server can administer remote StarTeam Server configurations only. With the **Server Administration** tool, an administrator can perform all operations on either remote or local server configurations, including the following:

- Create, enable, disable, or delete a server configuration.
- Display or modify the session options for a server configuration.
- Start or shut down a server configuration.
- Set or remove a server configuration as a Windows service.
- Review the status and execution mode of all server configurations running on this computer.
- Access the Hive Manager.

You can also perform the following tasks on remote server configurations from clients on which you have installed the **Server Administration** tool:

- Log onto a server as a different user.
- Add and manage user accounts.
- Set the security policy for a server configuration.
- Assign access rights to users and groups for a server configuration.
- Add, modify, or delete connections to a server configuration.
- Set or modify the configuration options for a server configuration.
- Display the server log file (`Server.locale.Log`).
- Lock or unlock a server configuration.

The rest of this topic describes the numbered components in the above diagram.

Main Menu

The main menu consists of the **Server**, **Actions**, **Tools**, and **Help** menus. The **Tools** menu provides a cascading menu separating administrative and user account commands.

The **Server Administration** tool enables or disables menu commands depending on the status of your server configuration. For example, when you are not running a server configuration the **Server Administration** tool does not enable the **Actions** > **Logon As Shutdown Server** main menu commands.

Toolbar

Frequently used main menu commands corresponding to the **Server** and **Actions** menus have corresponding buttons on the toolbar. Fly-over text displays when you hover your mouse over the toolbar buttons. The Server Administration tool enables or disables toolbar buttons depending on the status of your server configuration.

Server Pane

The **Server** pane lists the servers that are present in the `starteam-servers.xml` file. Choosing **Server** > **Add Server** and proceeding through the **Add Server** dialog box updates this file.

Shortcut Pane

The shortcut pane displays quick access buttons corresponding to the cascading menus provided under the **Tools** menu for the administrative and user account commands. The shortcut pane is divided into the **Administration** and **Accounts** areas enabling you to access frequently used main menu commands.

Display Pane

When accessing commands from the **Tools** menu or from the shortcut pane quick access buttons, the **Server Administration** tool displays the dialog boxes for these commands in the display pane.

Tip: Expand the **Server Administration** tool window to enlarge the dialog boxes presented in the display pane.

Toolbar

The following are the toolbar buttons on the Server Administration Tool toolbar:

Click **New Configuration...** to create a server configuration.

Click **Add StarTeam Server** to add a server.

Click **Configuration Properties** to show the **<Configuration Name> Properties** dialog box.

Click **Delete** to delete an item.

Click **Logon As** to log on as a different user.

Click **Start Server** to start a server configuration.

Click **Start with Override...** to open the **Start with Override** dialog box to start the server with overrides.

Click **Set to Run as Service** to have the server configuration run as a service.

Click **Enable/Disable Server** to mark the server enabled or disabled. A disabled server configuration cannot be started.

Click **Shut Down Server** to stop the server configuration.

Click **Lock** to lock the server configuration

Click **Unlock** to unlock the selected server configuration.

Click **Migrate Database** to start the database migration process.

Click **Upgrade Database** to upgrade your database.

Click **Catalog Export** to export the following application tables: `Catalog_Tables` and `Catalog_Fields`.

Click **Statistics monitoring** to run the **Statistics Monitoring** tool. To enable, click **Tools** > **Administration** > **Configure Server** and then use the **Diagnostics** tab.

Click **Load Samples** to load sample data into your server configuration. An empty configuration is required.

Click **Configure Maintenance Tasks** to open the **Maintenance Task Scheduler**. Use this to automate running scripts.

Server Configuration Status Icons

When using the Server Administration tool, you will notice that icons display to the left of the server configurations to indicate their status. These icons are described below.

A running server configuration.

A server configuration in the process of starting.

A server configuration running as a Microsoft Windows service.

A server configuration that is not running.

A new server configuration that is not running but is enabled.

An enabled server configuration that is not running but set up to run as a Microsoft Windows service.

A disabled server configuration.

A server configuration in the process of shutting down.

Server Administration Keyboard Shortcuts

The following are the keyboard shortcuts for the **Server Administration** tool:

New Group	**Ctrl+G**
New User	**Ctrl+U**

Creating, Running, and Managing Server Configurations

Creating Server Configurations

Before creating a new server configuration, you need to decide upon a unique name for the configuration. This name is case insensitive and cannot contain `: \ /` but can contain blanks or apostrophes `'`. You must also set up the database to be used with the server configuration. A database can contain only one server configuration. However, other applications can share a database with StarTeam.

Tip: You can also create a server configuration from the command line.

1. Open the Server Administration tool.

   Note: You must access the Server Administration tool on the computer where StarTeam Server is installed.
2. Click **Server** > **New Configuration** . The **Create a New Configuration** wizard opens.
3. To set the **General** options for the new server configuration, do the following:
   a) Type a unique name in the **Configuration name** field.
   b) In the **Repository path** field, enter or click **Browse** for the location in which the StarTeam Server will create the server configuration files.
   c) Select a **Database type** from the list. You cannot change the database type once the server configuration has been created.
   d) Select the **Message Broker type**. Choose `None`, `Legacy StarTeam MPX`, or `ActiveMQ MPX`.
   e) Check or clear **Create new StarTeam database**. The wizard selects this option by default.
   f) If you want your clients to use the search functionality, verify that **Enable search for StarTeam Server** is checked.
4. Select the **Default** or **Custom** hive option for the **Initial Hive Settings**.

   Note: If you select the **Default**, changing the repository path changes the default hive settings. Changing the repository path does not have this effect if you select **Custom**.

   If you select **Custom**, you can override the default hive settings by modifying any of the following fields:

Name	Unique name for the hive. `DefaultHive` is the default.

Archive path	Path to the `Archives` folder for the new hive. The default path is `<repository path>\DefaultHive\Archives`.
Cache path	Path to the `Cache` folder for the new hive. The default path is `<repository path>\DefaultHive\Cache`.
Maximum cache size	Maximum number of MB of hard disk space that the cache can use. The default is 20% of the disk space available when the option is set.
Cache cleanup interval	Seconds between cache cleanup/refresh operations. The default value is 600. The range is 60 (1 minute) to 3153600 (1 year).
Storage limit threshold	Percentage of total disk space allowed for hive. When this percentage has been reached, StarTeam does not add any more archives to the hive. The default is 95% of total disk space.

5. Click **Next** when the information is complete. The second page of the wizard opens. The information that you must enter for this page of wizard varies according to the database selected.

   For Microsoft SQL Server/Microsoft SQL Server Express and PostgreSQL databases:

   a) In **Database Server Name**, type or browse for the names of the computer and the database on your network that should be used.
   b) If the database is running on an alternate port, check the **Edit Database Port** field to enable the **Port** field. Type the port number in the field.
   c) Type the password for the system administrator in the **Sys Admin (sa) password** field. If this is a Microsoft SQL Server Express instance installed with StarTeam Server, the initial default system administrator password is `StarTeam123` or blank.
   d) Click **Verify Connection** to make sure that you can properly connect to the database.
   e) To keep the name of the server configuration, the database name, and the database login name the same, all of the **New database name** and the **New database login name** fields default to the name you provided for the server configuration in the first page of the wizard. Change these values if you prefer to have different names.
   f) Type and confirm a database password.

   For Oracle databases:

   a) Type the Oracle database server name in the **Database server name** field.
   b) Type either the service name or the SID. Select the option to specify one or the other.
   c) If the database is in a different port than the default port, check the **Edit Database Port** checkbox and type the new port number in the field.
   d) Type the database system password in the **System password** field.
   e) Click **Verify Connection** to make sure that you can properly connect to the database.
   f) To keep the name of the server configuration and the schema user the same, the **New schema user name** field defaults to the name you provided for the server configuration in the first page of the wizard. Change these values if you prefer to have different names.
   g) Type and confirm a password for the schema user name.

6. Click **Next**. The final page of the wizard, **Create Data Files and Transaction Logs** opens. Again, the information that you can enter for this page of wizard varies according to the database selected.

   For Microsoft SQL Server/Microsoft SQL Server Express databases:

   a) Review the information in the dialog box.
   b) If you have fewer than 15 users and expect to store 1 GB or less data, the default settings are appropriate for your use. If you are very familiar with Microsoft SQL Server and Microsoft SQL Server Express databases, you may choose to make some changes by first clearing the **Use default configuration** check box and then altering sizes and locations for data files and log files. To avoid fragmentation, make the data files as large as possible, based on the maximum amount of data expected in the database. Use at least 3 data files and at least 3 transaction log files when creating a database, because Microsoft SQL Server and Microsoft SQL Server databases use a proportional fill strategy. This way all the files tend to become full at about the same time.

c) When you are satisfied with your configuration settings, click **Finish**. A message requests: *Please make sure the required disk space (x MB total for both data files and transaction log files) is available on the database host machine.*

d) Click **OK**. After a `Please wait` message disappears, the Server Administration tool displays, showing your new server configuration as a child of the `Local` node.

> **Note:** Microsoft limits the size of a Microsoft SQL Server Express database, by license, to 2048 MB. If you require a larger database, you must purchase a license for Microsoft SQL Server.

For Oracle databases:

a) Review the information in the dialog box.

b) The tablespace name defaults to the name of your server configuration, but you can change that.

c) If you have fewer than 15 users and expect to store 1 GB or less of data, the default settings are appropriate for your use. If you are very familiar with Oracle schema users, you may choose to alter the names, sizes, and locations of the data files. To avoid fragmentation, make the data files as large as possible, based on the maximum amount of data expected in the database. Use at least three data files when creating a tablespace because there is a size limit of 2GB per data file, and fewer files can result in slow response times when insert activity is heavy.

d) Click **Finish**. The Server Administration tool displays, showing your new server configuration as a child of the `Local` node.

> **Tip:** If you are creating a custom component, do not start the server until you've completed all steps for creating components and their layouts.

7. Select the configuration from the server pane.
8. Click ▶ (**Start Server**). StarTeam Server then initializes the database and creates the files and folders for the server configuration. The initialization process may take a few minutes. When the server finishes this activity, the status icon to the left of the server configuration name changes from ☰ (**New**) to ✓ (**Running**).

In addition to creating the server configuration, StarTeam Server adds information about the new server configuration to your `starteam-server-configs.xml` file.

After the server configuration has been created, you can modify the default server configuration options, which enable you to fine-tune server configuration performance.

> **Note:** On a double-byte operating system (such as Japanese or Chinese), see your database administrator to manually set the collation sequence to `Latin1_General_CI_AS`.

Loading Server Configuration Sample Data

You can download sample data from within the Server Administration tool.

1. Create a new server configuration and select it.
2. Click ▶ (**Start Server**).

   > **Tip:** The server configuration must be empty (no existing projects) before you can load the samples.

3. After the server has started and is running, click **Actions** > **Load Samples**.
4. Click **Yes** to load the sample data. You may be prompted to log on to the server. A **Loading Samples** dialog box displays showing the progress of the download. The download may take several minutes.
5. Users can now connect to this new server configuration and access the sample data.

The sample data is installed in `C:\Program Files\Micro Focus\StarTeam Server <version>\Samples`.

Starting and Stopping Server Configurations

You can start a server configuration using the Server Administration tool or from the command prompt using the `starteamserver` command. You can also run the server configuration as a Microsoft Windows service.

Starting and Logging into a Server Configuration Using the Server Administration Tool

1. Open the Server Administration tool and select the server configuration.

   Note: You must access the Server Administration tool on the computer where the StarTeam Server is installed.

2. Click ▶ (**Start Server**).

   Note: The StarTeam Server uses TCP/IP port 49201 as the default starting port for the server configuration.

   The server configuration begins its startup operation. The first time you start a new server configuration, the StarTeam Server performs several startup tasks. It creates and initializes the database to be used by the server configuration, installs the stored procedures for that database type, and creates the repository folders and the hive used by the configuration. This process may take several minutes.
3. After the server configuration finishes its startup procedure, the status icon to the left of the server configuration changes to → (**Running**).
4. Login to the server configuration by performing any menu/toolbar commands. The **Log On** dialog box opens.
5. Type the **User name** and **Password**.

   Tip: Unless it has been changed or deleted by a server administrator, a default username/password, `Administrator/Administrator`, exists for every server configuration.

   Tip: To login as a different user, click **Actions** > **Logon As** .

Starting a Server Configuration with Different Configuration Options

1. Open the Server Administration tool and select the server configuration.

   Note: You must access the Server Administration tool on the computer where the StarTeam Server is installed.

2. Click ▶ (**Start With Override**). The **Start With Override** dialog box opens.
3. Modify the fields as appropriate and click **OK**. The server configuration information in the `starteam-server-configs.xml` file is update accordingly. If you are already using the default endpoint (49201) for another server configuration, the first time you start a new configuration you may want to use an override for the endpoint. This action sets the endpoint to the one that you will want to use in the future.

Note: You can override certain server configuration values with the `-restart` option.

Starting a Server Configuration Using Defined Values from the Command Line

1. Open a command prompt window and navigate to the installation folder of the StarTeam Server.
2. Type the following at the command prompt:

   ```
   starteamserver -start "ConfigurationName"
   ```

   Note: You can also start a server configuration to override the defined values using:

   ```
   starteamserver -start "ConfigurationName" [options]
   ```

Although you do not need to shut down a server configuration to perform a backup, you may need to do so to perform other maintenance tasks.

Note: If you have an Enterprise Advantage license and if you are running the StarTeam Server as a service and the Notification Agent as a dependent service, you cannot shut down the server configuration unless the Notification Agent service is shut down first.

Shutting Down a Server Configuration Using the Server Admin Tool

1. Open the Server Administration tool and select the server configuration.

   Note: You must access the Server Administration tool on the computer where the StarTeam Server is installed.
2. Click (**Shut Down Server**). This opens the **Server Shutdown** dialog box which displays a message asking you to confirm that you want to shut down the server configuration.
3. Click **Yes**.

Shutting Down a Server Configuration Using the Command Line

1. Open a command prompt window and navigate to the installation folder of the StarTeam Server.
2. Type the following at the command prompt:

   ```
   starteamserver -stop "ConfigurationName"
   ```

Server Configurations and Microsoft Windows Services

You can start a server configuration using the Server Administration tool or from the command prompt using the `starteamserver` command. You can also run the server configuration as a Microsoft Windows service.

Setting a Server Configuration to Run as a Microsoft Windows Service

Note: If a server configuration is newly-created, you must start it once, shut it down, and then set it to run as a service.

1. Open the Server Administration tool and select the server configuration.

   Note: You must access the Server Administration tool on the computer where the StarTeam Server is installed.
2. If the server configuration is running, click (**Shut Down Server**).

   Note: A configuration cannot be set to run as a Microsoft Windows service if the server includes a remote hive using a mapped drive.
3. Click (**Set to Run As Service**). The **Log On Service As** dialog box opens.
4. Check **Local System account** to use the local system account, or to use a specific user account, do the following:
   a) Clear the **Local System account** check box.
   b) Type an account name. The usual format is `DomainName\UserName`. If the account belongs to a built-in domain, you can use `.\UserName`.

The next time you start the server configuration or restart your computer, the server configuration runs as a Microsoft Windows service.

To determine whether a server configuration is running as a Windows service, locate the server name in the left panel of the Server Administration tool. Beside the name, an icon indicates whether the server is enabled and/or running as a service.

If you want to discontinue running a server configuration as a service, you must first stop the server configuration, and then remove the service using the Server Administration tool.

Shutting Down a Server Configuration that is Running as a Windows Service

1. Open the Server Administration tool and select the server configuration.

   Note: You must access the Server Administration tool on the computer where the StarTeam Server is installed.

2. If the server configuration is running, click ⏹ (**Shut Down Server**).

   Note: A configuration cannot be set to run as a Microsoft Windows service if the server includes a remote hive using a mapped drive.

3. Click ⚙ (**Set to Run As Service**). The server configuration will no longer run as a service.

If a server configuration that is set as a service fails for any reason or has been shut down, Windows records that information in the **Event Viewer Application** log.

Troubleshooting a Server Configuration that is Running as Microsoft Windows Service

1. From the computer on which StarTeam Server is installed, select **Start** > **Settings** > **Control Panel** > **Administrative Tools** > **Event Viewer** from the Windows Start Menu. The **Event Viewer** opens.
2. Click the **Application** node. The log information displays in the right pane of the Event Viewer.
3. Double-click the log entry to view the **Event Properties** dialog box.

Disabling and Enabling Server Configurations

You can disable or enable a server configuration from the Server Administration tool. Disabling a server configuration enables you to take a server configuration “out of service” and ensure that it is not started by accident. For example, if you migrate a server configuration, you should disable the prior server configuration. After you are sure that the new server configuration and database are working properly, you can delete the prior server configuration. You can also reactivate a disabled server configuration.

1. Open the Server Administration tool and select the server configuration.

   Note: You must access the Server Administration tool on the computer where the StarTeam Server is installed.

2. If the server configuration is running, click ⏹ (**Shut Down Server**).
3. Once the server has shut down, click ⊜ (**Disable Server**).
   - If the server configuration is currently enabled, it becomes disabled.
   - If the server configuration is currently disabled, it becomes enabled.

Tip: The icon to the left of the server configuration indicates its status.

Locking and Unlocking Server Configurations

Locking a server configuration enables you to limit access to that configuration while you perform backup procedures or database maintenance. When a server configuration is locked, only server administration commands are accepted. For any other command, such as checking out files, the StarTeam Server sends an exception message stating that the server configuration is unavailable.

To unlock your server configuration is locked, click 🔓.

1. Open the Server Administration tool and select the server configuration.

   Note: You must access the Server Administration tool on the computer where the StarTeam Server is installed.

2. Click 🔒 (**Lock Server**).

3. On the resulting dialog, indicate whether you want to:

Lock server This option allows minimal administrative options, primarily, start, shutdown, lock and unlock operations. It is usually done for backup operations in environments where server activity is not 24/7.

Lock the server for exclusive use by <user name> This option, which displays the logon name for the user, allows a user to lock the server for his or her use only.

A dialog box opens indicating that the server configuration is locked.

4. Click **OK**.

Note: If a locked server configuration is restarted, it will become unlocked.

Managing Users and Groups

User and Group Configuration Overview

You can use LDAP QuickStart Manager to import information about people from a directory service or LDIF file into a StarTeam or Caliber Server as user properties. You can also manually add new groups and users to a server configuration. When users log onto the application, they can be validated by a password that has been entered in or imported to the application or obtained from Microsoft Active Directory Services (the LDAP server). This operation is possible only when the server is on a trusted domain in relation to the LDAP server.

Understanding the Default Groups

New server configurations come with predefined default groups: All Users, Administrators, System Managers, and Security Administrators. These groups come with default privileges but you can assign privileges in accordance with your company policy.

The users in the Administrators group initially have all available privileges, giving them complete access to the system unless the system is set up to ignore privileges. The All Users, System Managers, and Security Administrators groups initially have no privileges.

All Users All users are members of the All Users group because All Users is the root group in the User manager and because all members of a child group are members of its parent group. Therefore, all users inherit any rights and privileges assigned to this group.

Administrators This group initially contains the Server Administrator user. You may want to add others who have administrative privileges.

StarTeam Server comes with a user named “Administrator” who has the password “Administrator”. Because this is common knowledge, you will want to change that password.

System Managers The users in this initially-empty group receive email (at the address specified for them in the User Manager) whenever an error is added to the server log.

Security Administrators The users in this group can receive email about users who attempted to log on unsuccessfully. This group initially contains only the user who has been designated as the Server Administrator.

Tip: Never have only one user account with administrative privileges. If you are logged on using the only user account with administrative privileges and you become locked out, you have no way to unlock your own account.

Group Membership

A user can be a member of more than one group. If users belong to multiple groups, they can perform operations at the highest level permitted by any of their group privileges. For example, suppose that User A belongs to both the `All Users` group and the `Administrators` group and that the **Delete Item** privilege is granted to the `Administrators` group but not to the `All Users` group. User A can then delete any item in the server configuration projects.

Membership can be explicit or implicit. Membership in a group is explicit if:

- The group was selected at the time the user was created.
- The name of the group was selected from the **Group Membership** tab in the **User Properties** dialog box of the Server Administration tool.

The group hierarchy determines implicit membership. If a user is a member of a child group, the user is also a member of the parent group, even if the name of the member does not appear in the user list when you select the parent group. For a selected group that has child groups, you must select the **Show Users in All Descendant Groups** check box to see the complete list of members.

A user who is a member of a parent group and also a member of a child group within that group will have both implicit and explicit membership in the parent group.

Directory Service Support

StarTeam allows password verification with Active Directory. It allows centralized, secure management of an entire network. To validate users against the directory server, the StarTeam Server must be on a trusted domain in relation to that server.

On the **Directory Service** tab of the **Configure Server** dialog box, you must also select the **Enable directory service** option and enter the location and port number of the directory server. For each individual who will be validated against the directory server, you must select the **Validate with directory service** option on the **New User Properties** or **User Properties** dialog boxes and enter a **Distinguished name** (used to uniquely identify a directory services user).

Even if the settings are correct, the user will not be able to log on if the directory server is unavailable. Although directory service support is off by default, it can be activated at any time. The server cannot be running at the time you enable or disable the support. When the user supplies a logon name and an Active Directory password, StarTeam Server recognizes that the user is set up for directory service password validation and uses the **Distinguished name** and password as it contacts Active Directory. If the password is verified, the user is allowed to access the server configuration.

Overview of Security Strategies

By default, all users initially have access to everything in the client. To avoid accidental deletions and other problems, administrators must set access rights as soon as possible.

The following sections cover access rights and provide general security guidelines:

- General security guidelines.
- Server-level access rights.
- Project, view, folder, and item-level access rights.
- Component, filter, and query-level access rights.
- Access rights for promotion states.

General Security Guidelines

Until you become familiar with access rights, we recommend that you follow the guidelines suggested in this section.

From the StarTeam Server

On the StarTeam Server, the **User Manager** dialog allows you to create users and groups for each server configuration while that configuration is running. Use the following guidelines:

- Do not change the privileges for the All Users, Administrators, System Managers, and Security Administrators groups.
- Do not create additional groups under the Administrators group.
- Create the groups that you need under All Users or under each other. For example, you may need to create the following groups: Developers, Testers, and Writers.
- Create users and assign them to groups. Make sure that at least two users are administrators, in case one administrator becomes locked out.

From the Client

Use the following guidelines:

- Although you can deny rights as well as grant them, it is best only to grant them.
- If you do deny rights, observe both of the following rules:
  1. Never allow any node on an **Access Rights** dialog box to have only deny rights records.
  2. Always make sure that the deny rights records for any node precede any records that grant rights for that node.
- When you set access rights for a node, remember that any user who does not have access rights for the node (individually or in a group) is denied all rights at this level for this node (unless that user has privileges that allow access).
- Set access rights at the project level first. Set them for every group (except the All Users group) for every node. The nodes are Project, View, Child Folders, File, Change Request, Requirement, Task, and Topic. Depending on which version of the client your company uses, you may not see all of these nodes. The most important nodes to set at this level are the Project and View nodes. The Project node is the only location in which you can set project access rights. The View node controls view-level access to all views. Newly created views start out with only the view access rights set here for all views. Initially, they have no view-level access rights.
- Set access rights at the view level next. Set rights for every user and/or group that needs access at this level for every node. (The nodes are View, Child Folders, File, Change Request, Requirement, Task, and Topic).
- Set up access rights at the folder level only if you really need to have access rights for the folders. Remember that these settings go with the folder when it is moved or shared and when it becomes part of new views (until the folder branches in the new view). Remember that folders branch only when their properties change, and that their properties tend to change infrequently.
- Avoid setting access rights on root folders because those rights can conflict with those set at the project or view levels.
- Avoid setting access rights on items. Remember that these settings go with the item when it is moved or shared and when it becomes part of new views (until the item branches in the new view).

Server-Level Access Rights

Server-level access rights allow users to perform server administration operations, such as modifying server configurations and viewing logs. Additional rights at the server level include the rights to create projects, create custom fields, control component-level access rights, and perform certain operations specific to the Notification Agent.

The server-level rights you assign to users and groups authorize them to perform specific operations in a particular server configuration. One of the options determines who can and who cannot create projects when the server configuration is running.

Note: Server-level access rights can be assigned only when a server is running.

By default, the Administrators group is assigned all project and server rights. By default, the All Users group has the rights to create projects and review the server configuration and the server log.

Project, View, Folder, and Item-Level Access Rights

Initially, any user who can see a project, view, folder, or item can set the access rights for it. However, project-level, view-level, folder-level, and even item-level rights function hierarchically and may be affected by group privileges.

As users log onto a server configuration, they are identified by their user names and as members of the groups to which they belong. This information is stored as an access token for each user. When users perform operations on objects (projects, views, folders, and items), the client examines these tokens and the access rights for the objects on which the users are performing the operations.

Determining Object Access Rights and Tokens

The StarTeam server checks access rights in layers. The right to access an object begins with the **System Policy** tab which can be reached by choosing **Tools** > **Accounts** > **System Policy** in the Server Administration Tool.

Similarly, unless privileges are being ignored, the privileges granted to groups also override and take precedence over the access rights configured elsewhere. Privileges are group properties that are set by using the **Privileges** tab of the **Group Properties** dialog.

A user is granted the same privileges as the group to which he or she belongs. If the user belongs to two groups and one group is granted certain privileges and the other group is denied the same privileges, the user is granted the privileges because at least one group to which he or she belongs has those privileges.

After checking privileges, the client checks the access rights granted for specific objects. Settings on the **Access Rights** dialogs for projects, views, folders, and individual items grant or deny users or groups the ability to perform operations at those levels.

Note: If rights are granted to any user or group at a given level in an **Access Rights** dialog, those users who are not granted rights at that level are effectively denied the rights. Ultimately, if a user can see an object and no deny records stop the user from performing an operation, the user can do anything that a grant record allows him or her to do, whether as an individual user or as a member of a group. The only exception involves issues of privileges.

To summarize, the client performs the following checks to determine whether a user can perform an operation:

1. If the user belongs to a group that has a satisfactory privilege and privileges are not being ignored, access is granted. Note that privileges, when not ignored, take precedence over access rights wherever access rights are set. If users belong to a group that has the correct privileges, they can be granted access rights that are specifically denied to them in the client.
2. If the user or any group to which the user belongs has been granted satisfactory access rights for the object on which the operation will performed, access is granted. If the object has access rights set, but none are satisfactory, the user is denied access.
3. If the object has no access rights set, the client checks the next higher level. For example, if the operation is on a file, change request, topic, task, or child folder, the client checks the access rights for the parent folder. If the operation is on a root folder, the client checks the access rights for the view. If the operation is on a view, the client checks the access rights for the project. If the operation is creating a project, the server access rights are checked.
4. If none of the levels has access rights set, access is granted.

Administrators can override group privileges by setting options from the server configuration **System Policy** dialog. Use this option with caution, because it changes the steps used by the StarTeam Server to check every user (including administrators) for access to all objects in the repository. If you ignore privileges, only access rights determine who can and cannot perform operations on objects in the repository.

Group Privileges for Objects

The privileges assigned to a group may allow members of that group to access objects and perform operations that they are otherwise not allowed to do. In other words, group privileges override the access rights settings.

If you choose **Tools** > **Accounts** > **User Manager** from the Server Administration Tool, notice that the server configuration comes with some default groups: All Users, Administrators, System Managers, and Security Administrators. The default user named Administrator belongs to both the Administrators and the Security Administrators groups. By default, the Administrators group has all group privileges. Also by default, the other groups have none of these privileges.

All members of a group have the same privileges on every project managed by this server configuration. The privileges apply to all levels equally: projects, views, folders, and items within folders. If users belong to more than one group, they have the maximum amount of privileges, regardless of which group provides them with those privileges.

Understanding Object Access Right Levels

Access rights are defined for individual users or groups at the following levels:

Project level Access rights can be defined for the project itself. You can also define access rights that apply to all its views, child folders and items, unless a object has access rights set specifically for it. There are View, Child Folders, and other nodes at this and other levels.

View level You can define access rights for the view itself. You can also define access rights that apply to all its child folders and items, unless a specific object has access rights set specifically for it.

Folder level You can define access rights for the folder itself. You can also define access rights that apply to all its child folders and items, unless a specific object has access rights set specifically for it.

Item level You can define access rights to a specific file, change request, requirement, task, or topic. (It is unusual to set rights at this level.)

Note: Project access rights can be set only at the project level, because that is the only level with Project node in the access rights hierarchy. You can set view access rights at either the project or the view level, because both of those levels have a View node. You can set folder access right at the project, view, or folder levels, and so on.

Opening Projects and Views

A project is indistinguishable from its initial view and also from the root folder of that view. In fact, any view of a project is indistinguishable from its root folder. Therefore, a user will not be able to open a project if you deny that user (or all groups to which the user belongs) any of the following:

- Ability to see the project.
- Ability to see the initial project view.
- Ability to see the root folder of the project’s initial view.

A user will not be able to open a particular view of a project if you deny that user (or all the groups to which the user belongs) any of the following:

- Ability to see that view.
- Ability to see that view’s root folder.

Component, Filter, and Query-Level Access Rights

The client components (file, change request, requirement, task, and topic) are server-wide objects. For example, the change request component appears in every project view and has the same filters and queries in every view.

Component-level access rights govern the use of filters and queries for each component. They determine the users who can create public filters and queries in that component, who can use the filters and queries, and so on. A server-level access right named **Administer component-level access rights** allows users to set these rights.

Individual filters and queries also have access rights. These rights override the general access rights set for filters and queries.

The right pane contains a tree of access rights subcategories. When expanded, each subcategory displays a set of access rights as its children.

Each filter or query resides in a particular component (such as the Change Request component or the File component) and can be applied to that component's type of data only in any project view managed by a specific server configuration.

Any user can create and use private filters and queries, but public filters and queries have access rights, individually and per component. Rights set on a specific filter or query take precedence over access rights set at the component level.

To apply a public filter or query, a user must be able to access the data type for the component in some open project view. When you apply the filter or query, it affects the type of data that visible in the open project view.

Users can apply any public filters and queries that they can view. In general, users can see any public filters and queries for which they have access rights.

Access Rights for Promotion States

Each view has its own set of promotion states. Access to these states is controlled by:

- The "Define promotion level" right.
- Access rights that govern access to individual promotion states.

The Define Promotion Level Right

The **Define promotion level right** is available from the **View** node of the **Access Rights** dialog at the view and project levels. A user with the Define promotion level right can do anything to the promotion model:

- Create and delete states.
- Edit their properties.
- Promote a label from one state to another. Promotion is a subset of editing properties. Anyone who can edit the properties of a state can also promote that state.
- Reorder the states within the view.

Promotion State Access Rights

Promotion state access rights govern access to individual promotion states. These Generic object rights and Promotion state specific rights are available from the Promotion State node of the Access Rights dialog at the view and project levels. They also appear on the access rights for individual promotion states.

The rights for an individual promotion state are checked at the state level; if necessary, the checking continues at the view level and eventually the project level. If a user is granted a given right at one level, there is no need to check the next.

When a right is granted at the view level, it applies to all states in the view, unless access is denied at the state level.

When a right is granted at the project level, it applies to all the states in all the views within the project, unless access is denied at the state or view levels.

Component Access Rights

If you have the server-level access right to “Administer component-level access rights,” you can set component-level access rights from any open component. Below is a description of the Component Access Rights.

Generic item rights

Create public filters	Create public filter for this component.
Create public queries	Create public queries for this component.

Setting Component-Level Access Rights

If you have the server-level access right to **Administer component-level access rights**, you can set component-level access rights from any open component.

1. Open any project view to which you have access.
2. Select the correct tab (file, change request, and so on) for the component.
3. Select **<Component Type>** > **Advanced** > **Component access rights** from the main menu. The **<Component Type> Component Access Rights** dialog box opens.
4. Select an appropriate node:
    - To control who can create public filters and queries for the component, use the **Component** node.
    - To control who can use public filters for the component, use the **Filter** node.
    - To control who can use public queries for the component, use the **Query** node.
5. Click **Add**. The **Assign Access Rights To** dialog box opens.
6. Select a user or group. Users are listed by their user names and groups are listed by their paths (excluding the `All Users` group).
7. Select **Grant** and click **OK**.

    Caution: Never select **Deny** unless you are creating an exception. Deny records must be created before grant records.
8. Click **OK**.

    Caution: Never select **Deny** in the **Assign Access Rights to** dialog box unless you are creating an exception.
9. Select/clear the appropriate check boxes. Selecting or clearing the check box for a category, such as **Generic object rights** for a project, selects or clears all the access right check boxes for that category. The category check box has only two states. When it is cleared, the access right check boxes for that category are either all cleared or mixed: some selected and some cleared.

    Caution: Clicking **Delete** removes the selected user or group from the **User and Groups** list in the **Access Rights** dialog box. The selected user or group loses any previously set access rights to the Server.
10. Click **OK**.

Component-level Filter Access Rights

The following describes the filter access rights at the component level:

Generic item rights

See object and its properties	See public filters for this component in the **Filters** list (on the toolbar) and view their properties in the **Filters** dialog box.

Modify properties	Change public filter properties for this component. The properties that can be modified for a filter are its list of displayed fields, its sorting and grouping rules, the query associated with it, and its context (the items of the component to which it can be applied).
Delete object	Delete public filters for this component from its list of filters.
Change object access rights	Change access rights for public filters for this component.

Component-Level Query Access Rights

The following table describes the query access rights at the component level:

Generic item rights

See object and its properties	See public queries in the **Queries** dialog box and view their properties in the **Edit Query** dialog box.
Modify properties	Change public queries properties for this component. The properties that can be modified are the query's name and its conditions.
Delete object	Delete public queries for this component from its list of queries.
Change object access rights	Change the access rights for public queries for this component.

Configuring Server-level Access Rights

The server-level access rights you assign to users and groups authorize them to perform specific operations in a particular server configuration. One of the options determines who can and who cannot create projects when the server configuration is running.

Note: You can assign server access rights only when a server configuration is running.

1. Open the Server Administration tool and select the server configuration.
2. Click the **Accounts** bar and then click (**Access Rights**). The **Access Rights** dialog box opens.
3. Click **New**. The **Add a User or Group** dialog box opens.
4. Select the user or group to be assigned access rights.
5. Check **Grant**, and click **OK**.

   Caution: Never check **Deny** unless you are creating an exception.
6. Select a user or group from the **User and Groups** list. This enables the various check boxes in the **Access Rights** dialog box. You can select or clear the appropriate check boxes as needed. If you cannot view the entire **Access Rights** dialog box, resize the Server Administration tool window.

   Click **Select All** and **Clear All** as necessary to speedily check or clear all of the check boxes in the **Access Rights** dialog box.

   Caution: Clicking **Delete** under the **Users and Groups** list removes the selected user or group from the list. As a result, the user or group loses any previously set access rights to the server.
7. Click **OK**.

Project Access Rights

The project-level rights you assign to users and groups authorize them to perform specific operations in a particular project.

By default, the `Administrators` group is assigned all project rights. The project access rights are briefly described below.

View/edit project access rights section
Access rights set on a project/view can be viewed and edited from the server administration tool. Access rights report can be generated with options to include server ACLs and include folder level access rights. The report can be saved to disk or printed. Editing project access rights is only available on remote server administration tool shipped with the cross platform client.

User access rights report section
Effective access rights test performed to check if a given user has specific permissions on a project/view. A report can be generated detailing the results of the test (if the user is granted /denied rights and an explanation for the result). The report can be saved to disk or printed.

Clone Project rights section
Most times administrators require to grant similar set of access rights across multiple projects. Access rights setup can be complicated and cumbersome to be replicated across multiple projects. This feature allows administrators to clone project level access rights easily from one project to another.

Clone view rights section
Most times administrators require to grant similar set of access rights across multiple views. Access rights setup can be complicated and cumbersome to be replicated across multiple views. This feature allows administrators to clone view level access rights easily from one view to another under the same project.

Server Access Rights

The server-level rights you assign to users and groups authorize them to perform specific operations in a particular server configuration. One of the options determines who can and who cannot create projects when the server configuration is running. Server rights can be assigned only when a server is running.

By default, the `Administrators` group is assigned all server rights. By default, the `All Users` group has the rights to create projects and review the server configuration and the server log. The server access rights are briefly described below.

Server operations

View server log
Review, but not change, server log information.

View statistics and licensing information
Review, but not change, statistics information (StarTeam Server 5.4 and earlier). Create license usage files.

View system configuration
Review, but not change, the server configuration options.

Modify system configuration
Change the server configuration options.

Remotely administer server
Lock/unlock the server, restart the server from the client, shut down the StarTeam Server from the client, access the **Start/Stop Conversion** and **Hive Manager** vault buttons.

Administer user accounts
Add groups and users.

View system policy
Review, but not change, the password and logon failure options for the server configuration.

Modify system policy
Change the password and logon failure options for the server configuration.

Change server security settings
Set Server access rights. If you change this setting, be sure that you remain one of the users who can change access rights.

View security log	Review, but not change, server log information

StarDisk Operations

Create new users	Add new users to sample project.

Replication Support

Change user/operation time	Manipulate creation times and user names when using special clients, such as StarTeam Notification Agent.

Project Operations

Create projects	Create projects when the Server is running the server configuration.

Customizations

Add/modify database schema	Create customized fields as item properties, or modify a field for an item that can be modified.

Component operations

Administer component-level access rights	Designate the users and groups who can create and apply filters and queries for a specific component in the server configuration.

Setting Up Users

If you have the appropriate access rights, you can add users to a server configuration from either the Server Administration Tool or a client. Initially, you add a user to a specific group, such as `Developers` or `Testers`. The user becomes an explicit member of this group and an implicit member of any of this group's parent groups, such as the All Users group. This operation can be performed only when the server is running.

Caution: Creating a user account with the name `StarTeam` has been known to cause problems when using the command line `stcmd` server-mode command to lock or unlock the server configuration. The command requests a password even when the user has a blank password or when a password has already been provided.

Note: If you are using the Server Administration tool installed with the client, you can administer remote servers only.

Adding a User

1. Open the Server Administration tool and select the server configuration.
2. Click the **Accounts** bar and then click (**User Manager**). The **User Manager** tab opens.
3. Select a group from the **Groups** tree and click **New User**. The **New User Properties** dialog appears.
4. Type the user's name in the **Full Name** field and optionally type the user's e-mail address in the **E-Mail** field.
5. Optionally, type the user's phone number, voice mail number, pager number, fax phone number, and street address in the appropriate fields.
6. Click the **Logon** tab.

a) Type the name to be used to log onto the application in the **User Name** field. If you enter a user name that already exists, the following message displays after you click **OK**: `A user with a given user name already exists.`

b) Select the **Validate through StarTeam Server** button if you want to validate the user against the server. Type a StarTeam password for the user in the **Password** field and again in the **Confirm** field. Asterisks appear in the text box instead of the password itself. If the password's minimum length can be zero, you do not have to enter a password. If you are using strong passwords, be sure to follow the rules for those passwords.

c) (For Microsoft Active Directory or OpenLDAP) To validate the user against your organization's directory server, select the **Validate through directory service** button and type the **Distinguished Name** for the user. An alphanumeric value of up to 254 characters, this value is used to uniquely identify the directory services user. To use directory service validation, the Server must be on a trusted domain in relation to the LDAP server.

7. Optionally, select the **Access Policy** tab and specify when this user can access the server configuration. Select one of the following option buttons:

   - **Access not restricted** (the user can access the server configuration at any time)
   - **Standard five day work week** (the user can access the server configuration Monday through Friday from 8 A.M. to 5 P.M.)
   - **Custom access hours** (to set one or two time periods per day when the user can or cannot access the server configuration)
   - Select a day of the week from the Day list.
   - Select the **No Access on That Day** check box to deny access, or clear it to allow access on that day.
   - Use the **From** and **To** boxes to set one or two time periods when access is either allowed or denied.

   If the user's workstation is not in the same time zone as the computer on which the server configuration is running, select the **Adjust for Workstation Time Zone** check box, and type the number of hours from Greenwich Mean Time (GMT) in the **hours from GMT** field.
8. Add the new user explicitly to other groups, as appropriate. Remember that a user is implicitly already a member of the current group's parent groups, but you must explicitly add a user to groups that are not parents of the current group.

Reviewing a User's Explicit Group Memberships

1. Open the Server Administration tool and select the server configuration.
2. Click the **Accounts** bar and then click (**User Manager**). The **User Manager** tab opens.
3. Select the user from the **User** list.

   If the user does not appear in the **Users** list, you can display a list of all users by:

   1. Selecting the **All Users** group in the **Groups** tree.
   2. Selecting the **Show Users in All Descendant Groups** check box.
4. Right-click and select **Properties**. The **User Properties** dialog box appears.
5. Select the **Membership** tab. The list displays the groups in which this user has explicit membership.

Changing Group Memberships

1. Open the Server Administration tool and select the server configuration.
2. Click the **Accounts** bar and then click (**User Manager**). The **User Manager** tab opens.
3. Select the user from the **User** list.

   If the user does not appear in the **Users** list, you can display a list of all users by:

   1. Selecting the **All Users** group in the **Groups** tree.

    2. Selecting the **Show Users in All Descendant Groups** check box.
4. Right-click and select **Group Membership**. The **Group Membership** dialog box appears.
5. Select the groups to which you want to add this user explicitly.
6. Click **OK**.

Removing Users from Groups

1. Open the Server Administration tool and select the server configuration.
2. Click the **Accounts** bar and then click (**User Manager**). The **User Manager** tab opens.
3. Select the user from the **User** list.

   If the user does not appear in the **Users** list, you can display a list of all users by:

   1. Selecting the **All Users** group in the **Groups** tree.
   2. Selecting the **Show Users in All Descendant Groups** check box.
4. Right-click the user’s name, and select **Group Membership**. The **Group Membership** dialog box opens.
5. Deselect the option next to the group from which you want to remove the user.
6. Click **OK**.

Removing User Accounts

1. Open the Server Administration tool and select the server configuration.
2. Click the **Accounts** bar and then click (**User Manager**). The **User Manager** tab opens.
3. Select the user from the **User** list.

   If the user does not appear in the **Users** list, you can display a list of all users by:

   1. Selecting the **All Users** group in the **Groups** tree.
   2. Selecting the **Show Users in All Descendant Groups** check box.
4. Right-click the user’s name and select **Delete Account**. The system displays the following message:

   ```
   Do you want to delete username’s user account?
   ```
5. Click **Yes**. This action permanently removes the user from the server configuration.

Checking Account Status for Users

1. Open the Server Administration tool and select the server configuration.
2. Click the **Accounts** bar and then click (**User Manager**). The **User Manager** tab opens.
3. Select the **All Users** group in the **Groups** tree.
4. Select the **Show Users in All Descendant Groups** check box.
5. Review the information about the specific user that displays in the **Users** list.

   Tip: To ensure that the information in the Users list box is current, click **Refresh**.

Suspending User Accounts

1. Open the Server Administration tool and select the server configuration.
2. Click the **Accounts** bar and then click (**User Manager**). The **User Manager** tab opens.
3. Select the user from the **User** list.

   If the user does not appear in the **Users** list, you can display a list of all users by:

    1. Selecting the **All Users** group in the **Groups** tree.
    2. Selecting the **Show Users in All Descendant Groups** check box.
4. Right-click the user's name and select **Suspend Account**. The account status in the Users list changes to `Suspended`, and access to the server is denied after the user logs out.

Reactivating User Accounts

1. Open the Server Administration tool and select the server configuration.
2. Click the **Accounts** bar and then click (**User Manager**). The **User Manager** tab opens.
3. Select the user from the **User** list.

   If the user does not appear in the **Users** list, you can display a list of all users by:

    1. Selecting the **All Users** group in the **Groups** tree.
    2. Selecting the **Show Users in All Descendant Groups** check box.
4. Right-click the user's name and select **Reactivate Account**. The user account is reactivated.

Configuring Privileges

As an administrator, you can override group privileges by setting the option for the server configuration in its **System Policy** dialog box. Use these options with caution, because they change the steps used by the Server to check every user (including administrators) for access to all objects in the repository. If you ignore privileges, only access rights determine who can and cannot perform operations on objects in the repository.

Note: You can modify this option only on running server configurations.

1. Open the Server Administration tool and select the server configuration.
2. Click the **Accounts** bar and then click (**System Policy**). The **System Policy** tab opens.
3. Select the **Access Rights** tab.
4. Check or clear **Ignore Group Privileges**. When cleared, the server configuration checks for privileges.
5. Click **OK**.

Setting Up Groups

Users who can log onto a server configuration can be organized into groups. Creating and using groups simplifies the task of managing security on a project, because each group can be assigned a set of privileges that apply to all the users in that group, rather than setting privileges on a user-by-user basis.

The status bar on the **User Manager** dialog box displays the number of users in the selected group who have access to the server configuration, the number of users connected to the server configuration, and the number of users logged on. The number of users connected to the server configuration and the number of logged on users differ when individual users log on more than once.

These operations can be performed only when the server is running.

Adding a Group

1. Open the Server Administration tool and select the server configuration.
2. Click the **Accounts** bar and then click (**User Manager**). The **User Manager** tab opens.
3. Select a group from the **Groups** tree.

   Note: We recommend that, initially, you select the **All Users** group when adding a new group. Subsequent groups can be added to any group listed under the **All Users** group. Avoid adding new groups to the administrative and management group. If a user is a member of a child group, it is also implicitly a member of the parent group—even if the member's name does not appear in the list when you select the parent group. You must select the **Show Users in All Descendant Groups** check box to see the complete list of members for a selected group that has child groups.

4. Click **New Group**. The **New Group Properties** dialog box appears.
5. Type the group name in the **Name** field.
6. Type a description of the group in the **Description** field.
7. Select the **Privileges** tab.

   The privileges selected on the **Privileges** tab can override any Access Rights that have been previously set for any user in the privileged group. However, the privileges are not a substitute for Access Rights. If you have not set up Access Rights, you have no security system.

   The privileges set on the **Privileges** tab apply to all objects in all projects in a server configuration. For example, if you give a group the Delete Item privilege, any user in that group can delete any project, view, folder, child folder, or item from the server configuration, regardless of what the Access Rights are for deleting these items.
8. Set privileges as appropriate.
9. Click **OK**. The new group appears in the **Groups** list.

Changing the Parent of a Group

1. Open the Server Administration tool and select the server configuration.
2. Click the **Accounts** bar and then click (**User Manager**). The **User Manager** tab opens.
3. Select a group from the **Groups** tree.
4. Right-click and select **Change Parent Group** from the context menu. The **Change Parent Group** dialog box appears.
5. Select a new parent group, then click **OK**.
6. Click **OK**.

Determining the Members of a Group

1. Open the Server Administration tool and select the server configuration.
2. Click the **Accounts** bar and then click (**User Manager**). The **User Manager** tab opens.
3. Select a group from the **Groups** tree.
4. Select the **Show Users in All Descendant Groups** check box to also display the implicit members of the group in the **Users** list.

Removing an Empty Group

1. Open the Server Administration tool and select the server configuration.
2. Click the **Accounts** bar and then click (**User Manager**). The **User Manager** tab opens.
3. Select a group from the **Groups** tree.
4. Right-click and select **Delete**. The system displays the following message:

   ```
   Do you want to delete group groupname?
   ```
5. Click **Yes**.

   - If the group is empty, it is removed from the **Groups** list.
   - If the group contains users, the system displays the following message:

     ```
     The group you want to delete contains user accounts. Please delete these
     user accounts or move them to another group prior to deleting a group.
     ```
6. Click **OK**. Then either delete the users in this group or move them to another group.

Configuring Group Privileges

The privileges assigned to a group may allow members of that group to access objects and perform operations that they are otherwise not allowed to do. In other words, privileges override the access rights settings.

In the **User Manager** dialog box, you will notice that the server configuration comes with some default groups (`All Users`, `Administrators`, `System Managers`, and `Security Administrators`). The default user named `Administrator` belongs to both the `Administrator` and the `Security Administrators` groups. By default, the `Administrator` group has all group privileges. Also by default, the other groups have none of these privileges. All members of a group have the same privileges on every project managed by the this server configuration. The privileges apply to all levels equally—projects, views, folders, and items within folders. If users belong to more than one group, they have the maximum amount of privileges, regardless of which group provides them with those privileges.

Note: You can modify privileges in the **User Manager** dialog box only on running server configurations.

1. Open the Server Administration tool and select the server configuration.
2. Click the **Accounts** bar and then click (**User Manager**). The **User Manager** tab opens.
3. Add or select a group in the **User Manager** dialog box.
4. Add users to the group, if necessary.
5. Right-click the name of a group in the **Groups** tree and choose **Properties** . The **Group Properties** dialog box opens.
6. Select the **Privileges** tab.
7. Check or clear the check boxes to grant privileges to the group or take them away.
8. Click **OK**.

Group Privileges

The privileges assigned to a group may allow members of that group to access objects and perform operations that they are otherwise not allowed to do. In other words, privileges override the access rights settings.

If you select **User Manager** from the **Server Administration** tool, you will notice that the server configuration comes with some default groups: **All Users**, **Administrators**, **System Managers**, and **Security Administrators**. The default user named Administrator belongs to both the **Administrators** and the **Security Administrators** groups. By default, the **Administrators** group has all group privileges. Also by default, other groups have none of these privileges.

All members of a group have the same privileges on every project managed by this server configuration. The privileges apply to all levels equally: projects, views, folders, and items within folders. If users belong to more than one group, they have the maximum amount of privileges, regardless of which group provides them with those privileges.

Generic item rights

See object and its properties

See all projects, views, folders, items, and their properties. This privilege overrides the similarly named access right found in the **Generic Object Rights** in the **Access Rights** dialog boxes.

Modify object properties

Modify the properties of any projects, views, folders, or items. This privilege overrides the similarly named access right found in the **Generic Object Rights** in the **Access Rights** dialog boxes.

Delete object	Delete any projects, views, folders, or items. This privilege overrides the similarly named access right found in the **Generic Object Rights** in the **Access Rights** dialog boxes.
Purge object (delete permanently)	This privilege is not supported at this time.
Change object access right	Change access rights for any projects, views, folders, or items. This privilege overrides the similarly named access right found in the **Generic Object Rights** in the **Access Rights** dialog boxes.
Create object and place it in a container	Create new objects and put them in containers. When this privilege is set, the group can add new views to a project, new folders to a view, and new folders and items to a folder. This privilege overrides the similarly named access right found in the Generic Object Rights in the Access Rights dialogs. It does not override the server-level access right that allows users to create projects.
Grant all specific class-level rights for all classes of objects	Perform any operation not covered by the preceding privileges. For example, this privilege allows group members to check out files, break locks, perform linking operations, and perform labeling operations. This privilege overrides some of the access rights found in the Generic Object Container Rights and all of the access rights in the <item>-specific Rights in the Access Rights dialog.

Password Use

Passwords are required for the server administrator and users to access StarTeam Server configurations. When the server configuration is created, a server administrator account is created by default with both the user name and password set to `Administrator`. This password should be changed immediately. When the server administrator adds a user, a unique user name is created and a password is assigned according to the password properties specified for this server configuration.

The server administrator specifies password properties for each server configuration in the **Tools** > **Accounts** > **System Policy** dialog on the **Passwords** tab. Whatever is specified as the system policy for passwords applies to all users accessing this server configuration.

Password properties include the password expiration time limit, the minimum length, and use of strong passwords.

About Strong Passwords

The server administrator can specify that a strong password is required for users accessing a server configuration. If the system policy for this server configuration requires a strong password, the password must:

- New password must be different from the old password.
- New password must be different from the user name.
- New password must be mixed case, containing at least one lowercase and at least one uppercase alphabetical character. (This is the English alphabet as determined by the ASCII value of the character.)
- New password must contain at least one non-alphabetical character.

By default, the strong password option is turned off.

Password Property Changes

If the system administrator changes the password properties for a server configuration, when the changes take effect depends on the property.

Changes made to the password length properties take effect immediately, but apply only to new user accounts or new passwords. For example, if you change the minimum password length from eight characters to ten, all new users must have a password that is a minimum of ten characters long. However, existing users will still be able to use their eight character passwords.

Changes made to the expiration time limit take effect after the appropriate time interval. For example, if you change the password expiration time limit to thirty days, user accounts are suspended if their passwords have not been changed before the time expires. Users are prompted to change their passwords two weeks before the suspension takes place. The only user account not subject to expiration is the Administrator account.

If the strong password option is turned on, it applies only to new users and users who change their passwords. Until such a change is made, their old “weak” passwords continue to work.

Note: The system administrator can force a password change if they want users to immediately conform to a password property change or if a project security breach has occurred.

LDAP for Password Verification

StarTeam can use directory services (either Microsoft Active Directory or OpenLDAP) to perform password authorization. As users log on, they enter their StarTeam user name and their directory service password. Before allowing the users to access the server, StarTeam then checks a directory service for valid passwords.

LDAP Quickstart Manager is a utility that allows you to import information about people from a directory service or LDIF file into a StarTeam Server as user properties. LDAP Quickstart Manager makes it easy to maintain the DNs and other directory service information that you choose to store in StarTeam Servers.

To set up directory service authentication in StarTeam, you set options on the **Directory Service** tab of the **Configure Server** dialog. These options enable directory service support and provide information about accessing the service. In addition, you use the User Manager to set options for the individual users whose passwords are to be authenticated. Not all users need to use this feature.

The distinguished name (DN), a unique identifier, is used by Micro Focus servers as they communicate with the directory service. For example, StarTeam must send each user’s distinguished name (DN) to the directory service in order to verify the user’s password. DNs can be long and not very intuitive. Also, some organization’s change DNs occasionally, and updating these changes by hand can be very tedious.

When you import users using LDAP Quickstart Manager, you indicate whether new users will have their passwords authenticated by the StarTeam Server or by a directory service by selecting either the **Validate Password Through Directory Service** or the **Validate Password Through StarTeam Server** option button. StarTeam Servers request directory service validation of user passwords if the server configuration both allows directory service validation and has the correct connection settings for the directory service.

Enabling Directory Service Support

StarTeam allows password verification with Microsoft Active Directory.

1. Open the Server Administration tool and select the server configuration.
2. Click **Tools** > **Administration** > **Configure Server** . The **Configure Server** page opens.
3. Select the **Directory Service** tab.
4. Check **Enable directory service**. By default this option is not selected.
5. Type the **Host** name and a secure (SSL) or non-secure **Port** number for the directory server. By default the Server Administration tool specifies port 636. You must specify both values to enable directory service support.
6. You can optionally check the option to **Use a secure port**. This is the recommended default setting.
7. Click **OK**. The system displays a message instructing you to reboot the server. You must do this to enable directory service.

Note: Remember that a user cannot be authenticated by the directory server unless the **Validate through directory service** option is selected on the **Logon** tab of the **New User Properties** or **User Properties** dialog boxes and a **Distinguished name** is entered for that user.

Changing User Passwords

In addition to setting or changing a user’s password, you can specify how long a password is usable, how many characters a password must have, and whether strong passwords are required. This operation can be performed only when the server is running.

1. Open the Server Administration tool and select the server configuration.
2. Click the **Accounts** bar and then click (**User Manager**). The **User Manager** tab opens.
3. Select the user from the **User** list.

   If the user does not appear in the **Users** list, you can display a list of all users by:

   1. Selecting the **All Users** group in the **Groups** tree.
   2. Selecting the **Show Users in All Descendant Groups** check box.
4. Right-click, and select **Properties**. The **User Properties** dialog box appears.
5. Select the **Logon** tab.
6. Verify that the **Validate through StarTeam** button has been selected.
7. Type a new password for the user in the **Password** field.
8. Type the password again in the **Confirm** field.
9. Click **OK**.

Configuring Password Constraints

Changes made to the password length properties take effect immediately, but apply only to new user accounts or new passwords. For example, if you change the minimum password length from eight characters to ten, all new users must have a password that is a minimum of ten characters long. However, existing users will still be able to use their eight character passwords.

Changes made to the expiration time limit take effect after the appropriate time interval. For example, if you change the password expiration time limit to thirty days, user accounts get suspended if their passwords have not been changed before the time expires. Users will be prompted to change their passwords two weeks before the suspension takes place. By default, the strong password option is turned off. When this feature is turned on, as users change their passwords, they must provide strong passwords. Until such a change is made, their old “weak” passwords continue to work.

1. Open the Server Administration tool and select the server configuration.
2. Click the **Accounts** bar and then click (**System Policy**). The **System Policy** tab opens.
3. On the **Passwords** tab, select a password expiration option:

   - **Passwords never expire**
   - **Passwords expire after ___ days**. With this option, you must enter the number of days a password will be valid. StarTeam counts the days from the time the password was created.
4. In the **Password configuration** group, select the **Require Strong Passwords** check box to require passwords to meet all of the following criteria:

   - New password must be different from the old password.
   - New password must be different from the user name.
   - New password must be mixed case, containing at least one lowercase and at least one uppercase alphabetical character. (This is the English alphabet as determined by the ASCII value of the character.)
   - New password must contain at least one non-alphabetical character.

Selecting this check box also changes the value in the **Minimum password length** field to 3. You can increase it if you choose.

5. Optionally, type a number for the minimum password length. The default, zero, allows passwords to be blank. The maximum password length is 32 characters.
6. Click **OK**.

Forcing Password Changes

It may be necessary to force users to change their StarTeam passwords if a project security breach has occurred. This operation can be performed only when the server is running. You can set the password expiration time limit, the minimum length, and require the use of strong passwords. These password properties apply to all user accounts on the server configuration.

1. Open the Server Administration tool and select the server configuration.
2. Click the **Accounts** bar and then click (**System Policy**). The **System Policy** tab opens.
3. Select the user from the **User** list.

   If the user does not appear in the **Users** list, you can display a list of all users by:

   1. Selecting the **All Users** group in the **Groups** tree.
   2. Selecting the **Show Users in All Descendant Groups** check box.
4. Right-click the user's name, and select **Force Password Change**. The **Account Status** column in the **Users** list changes to `Password change required`. The user will be asked to change his or her password at the next log on. If the change is not made, the user is allowed access to the server configuration and the projects it contains, but will be locked out of the server configuration at the next log on. An error message warns the user that this will happen.

   Note: The accounts of users who fail to change their passwords can be reactivated by administrators.

Forcing Users to Log Off

You may have to force a user to log off for any number of reasons, including code violations or disaster recovery. When you force a user to log off, the user's next operation displays the following error message: `You are no longer logged on.`

Depending on the reason for your action, you may need an additional method, such as e-mail or the telephone, to notify users to stop accessing the application.

To log on again, the user must exit the application and restart the client. Most integrations between StarTeam and another application require the user to restart the application being used. However, these users are not usually notified that their connections to the server have been terminated. This operation can be performed only when the server is running.

1. Open the Server Administration tool and select the server configuration.
2. Click the **Accounts** bar and then click (**User Manager**). The **User Manager** tab opens.
3. Select the user. If the user you want to work with does not appear in the **Users** list, you can display a list of all users by doing the following:

   1. From the **Groups** tree, select the **All Users** group.
   2. Select the **Show Users in All Descendant Groups** check box.
4. Right-click the user's name and select **Force Logoff** from the context menu. The user is immediately denied access to the server configuration and to all projects residing in this server configuration.

   Note: You cannot force your own logoff.

Reactivating Administrative Accounts

It is possible for any user, even users with an administrative account, to be locked out of a server configuration when the number of retries with the wrong password has been exceeded. The lockout period for the main administrative account (`Administrator`) is 24 hours. However, you can unlock the administrative account before the 24 hours have elapsed by using the following procedure:

1. Shut down the server configuration and disconnect its network connection to keep remote users off.
2. Start the server configuration using the command line in foreground mode from the StarTeam Server installation folder.

   ```
   starteamserver -start StarClient -fg
   ```

   The configuration name specification is case-sensitive with the command line. The command prompt must be left open until the server configuration is shut down.

   Note: In the above example, `StarClient` is the name of the server configuration.
3. Set the system clock one day ahead.
4. Log in as `Administrator` and log off. This action will reactivate the `Administrator` account.
5. Set the clock back one day to its original time.
6. Shut down the server configuration by entering “X” and clicking `Type`, which is how the server is shutdown in foreground mode.

Configuring the Server

Online Purge

Online Purge (in the **Administration** group) allows administrators to purge deleted views and data from a server while it is running. A purge process deletes unwanted data from the database and removes deleted archives from the vault. This operation can be performed only if the server configuration is running.

Online Purge contains a simple **Start/Stop** button and a log content pane in the lower half which displays the progress of the purge as it deletes the data, and which can be refreshed at any time. The **Online Purge** process can be started and stopped using the **Online Purge** button in the **Server Administration** tool. You can start and stop **Online Purge** on a remote StarTeam Server as well as a local StarTeam Server.

Using **Online Purge** while StarTeam Server is running prevents the costly downtime of an Offline Purge, which could be anywhere from a few hours to a few days. **Online Purge** not only eliminates this costly downtime, but is much faster than an Offline Purge.

In **Online Purge**, newly deleted data will be available to purge only after a Server restart. **Online Purge** is an interactive process which can be stopped and restarted anytime when StarTeam Server is running. **Online Purge** records its current execution state and provides the ability to restart from the exact point where it stopped. After StarTeam Server starts, **Online Purge** has to be restarted manually.

Starting and Stopping Online Purge

This topic covers how to purge deleted views from a server configuration when the server is running. This is called an Online Purge. To use Online Purge, you must already have data that has been deleted from one or more views. Online Purge only purges data that has been deleted from a server.

Online Purge is started and stopped from the Server Administration tool. The **Online Purge** tab in the Server Administration tool enables application administrators to remove deleted views from a server configuration database and vault and rebuild the indexes in that database. If the deleted view has items that are active in another view, these items are not deleted. For example, if two views share a file and you delete one view, the shared file is not deleted. It is recommended that you perform a purge after deleting one or more views from a project.

The Online Purge operation takes much less time to complete than an Offline Purge if a large number of records need to be deleted or moved.

Note: Before you start any purge process, be sure to backup the database before using the purge feature since the process is irreversible. You should also start the server configuration from which the view was deleted at least once before using the purge feature. Purge is available for Oracle, Microsoft SQL Server, and PostgreSQL databases. You must have installed the database client application on the same computer as the Server for the purge to work properly. This operation can be performed only if the server configuration is running.

1. Open the Server Administration tool.
2. In the **Servers** list, select the server which contains the data you want to purge.

   Note: You must access the Server Administration tool on the computer where the Server is installed.
3. Click the **Online Purge** icon in the **Administration** section at the bottom left of **Server Administration** window.

   This opens the **Online Purge** tab on the right side of the **Server Administration** window.

   Note: Please note the information at the top of the **Online Purge** tab.
4. Press the **Start** button to start the Online Purge.

   The **Start** button is only available if the **Status** is `Ready`. Once the Online Purge starts, the button changes to a **Stop** button, and the **Status** changes to `In Progress`.

   As the Online Purge proceeds, a log of what is being deleted is displayed below the button. You can refresh the log at any time to see the current status of the purge.

   Note: If the Server is stopped for any reason during an Online Purge, you will have to restart the server, and manually restart the Online Purge. The Purge will start over from the beginning.
5. If you need to pause the Online Purge, press the **Stop** button.

   Note: Stopping the Online Purge pauses the process until you start it again by pressing the **Start** button. If the server has been running during this time, the Online Purge will continue from the place it was stopped.

When the Online Purge is complete, the button returns to a disabled **Start** button, and the **Status** is “Completed”.

Note: You cannot start another Online Purge on this server until more data is deleted from the server.

Customize VCM Tab

The **Customize VCM** tab allows an administrator to create new customized View Compare/Merge types based on the default merge types of `Rebase`, `Promote`, `Replicate`, and `Compare`. The administrator can specify at the server, view, or project level which merge types will be available to the user of that particular context. The administrator can also specify what the default merge action will be for each difference type found in the session.

Using the **Customize VCM** tab, administrators can simplify the View Compare/Merge process by presetting the View Compare/Merge operation settings in the **View Compare/Merge Wizard**, eliminating the need for users to view and set all of the **View Compare/Merge Wizard** options to start a VCM session. Along with setting the default merge actions, the administrator can also specify which Included **Item Types** and **VCM Options** to display to the user.

Note: Before you can create a custom VCM merge type, you must create a `StarFlow Extensions` project the server. First create a `StarFlow Extensions` project, then create the `Projects` folder under the root folder in the view. Otherwise the save operation will fail in the **Customize VCM** tab

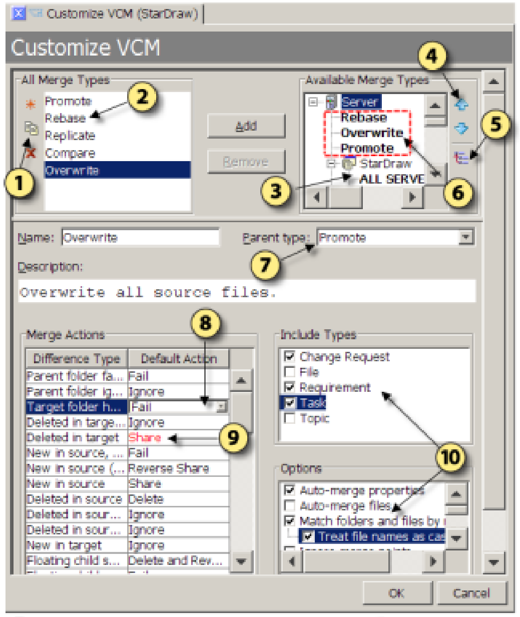

Name: Overwrite

Parent type: Promote

Description:

```
Overwrite all source files.
```

Merge Actions

Difference Type	Default Action
Parent folder fa...	Fail
Parent folder ig...	Ignore
Target folder h...	Fail
Deleted in targe...	Ignore
Deleted in target	Share
New in source, ...	Fail
New in source (...	Reverse Share
New in source	Share
Deleted in source	Delete
Deleted in sour...	Ignore
Deleted in sour...	Ignore
New in target	Ignore
Floating child s...	Delete and Rev...

Include Types

☑ Change Request
☐ File
☑ Requirement
☑ Task
☐ Topic

Options

☑ Auto-merge properties
☐ Auto-merge files
☑ Match folders and files by
☑ Treat file names as cas

1. New, Edit, and Delete Merge Type buttons
2. Default Merge Types cannot be deleted (Promote, Rebase, Replicate, Compare)
3. If context node does not have specific Merge Types added, it inherits all types from its parent node
4. Up/Down arrows used to change display order of Merge Types
5. Parent Merge Types button sets node to use only the Merge Types specified in its parent node
6. Specific Merge Types added to the server, project, or view level for the custom Merge Type are the only ones available in the VCM wizard at this level
7. All Merge Types must be based on one of the default Merge Types
8. Drop-down list editing for Default Action on each Difference Type
9. Red text indicates that the default merge action is not the same as the default merge action for this difference type in the parent Merge Type (the Merge Type on which the custom Merge Type is based)
10. Default item types to include and VCM options to use for this Merge Type in a VCM session

Available Merge Types

The administrator can control which custom merge types are available at the context level, such as the server, project, or view level. The **Customize VCM** tab provides a hierarchical context tree from the server down through the projects and views on each server. Custom merge types are specifically added to each desired level of the context tree. In the **Available Merge Types** tree:

- The nodes with icons are the context nodes which represent the server, project, and view levels.
- The nodes in bold text define what merge types will be available to the user when they are in that context.

By default, if StarTeam cannot find settings for a feature at the current view, it looks up the tree at the parent view. If there are no settings at the parent view, StarTeam will continue moving up the tree until it gets to the server level. When you add a custom merge type to a particular context view node, it becomes available for all the child nodes under it.

Note: Once you add specific merge types to a level in the **Available Merge Types** tree, only explicitly added merge types will be available in the **View Compare/Merge Wizard** for VCM sessions at that level. You must specifically add any default merge types back to the level if you want to still make them available. Use the **Parent Merge Types** button to quickly reset a level to use only its parent merge types.

The order you add merge types to a context level is the order they display in the **View Compare/Merge Wizard**. You can change the order using the **Up** and **Down** arrows to the right of the **Available Merge Types** tree.

Default Difference Type Actions

In the compare phase of a View Compare/Merge session, VCM uses the default merge actions for the type of merge selected to resolve any differences. The server administrator can control what default actions

View Compare/Merge will take for each **Difference Type**. The **Merge Actions** section allows the administrator to change which default action to take by selecting a different one from the drop-down lists in the **Default Action** column.

Note: A merge action that has been changed from the default parent action is displayed with red text.

Include Types

A user can limit the item types to include in a View Compare/Merge session using the **Include Selected Items** page of the **View Compare/Merge Wizard**. By checking specific item types in the **Include Types** section of the **Customize VCM** tab, the administrator can customize what item types appear in the **View Compare/Merge Wizard** for user selection.

Options

The **Options** section lets the administrator specify which compare/merge options to display as the defaults on the **Set Options** page of the **View Compare/Merge Wizard**. The options selected on this page of the wizard are performed when the View Compare/Merge session begins the compare phase.

Creating a Custom View Compare/Merge Merge Type

As a server administrator, you can create new, custom merge types for View Compare/Merge based on the standard merge types of Promote, Rebase, Replicate, or Compare.

Note: Before you can create a custom View Compare/Merge merge type, the server must have a StarFlow Extensions project. First create a StarFlow Extensions project, then create the **Projects** folder under the root folder in the view. Otherwise the save operation fails in the **Customize VCM** tool.

1. Open the **Server Administration** tool and click the **Customize VCM** icon in the **Administration** section.
2. In the **Customize VCM** pane, select one of the standard merge types in the **All Merge Types** section.
3. Click the **New Custom Merge Type** icon (the red asterisk in the left margin of the **All Merge Types** section.)

   Tip: You can create a new merge type by copying an existing custom merge type. Select the custom merge type in the list and click **Copy Custom Merge Type**.

   Clicking **New Custom Merge Type** or **Copy Custom Merge Type** displays the customization options in the lower area of the **Customize VCM** pane.
4. Give the new merge type a **Name** and **Parent Type**. All custom merge types must have a **Name** and a **Parent Type** because they are all derived from either Promote, Rebase, Replicate, or Compare.
5. In the **Available Merge Types** section, select the context level in the tree to which you want this merge type to be added, and then click **Add**.

   Note: The **Available Merge Types** tree is hierarchical. When you add a merge type to a node, it becomes available at all the child node levels under it. To make the new merge type available at all levels, add it to the **Server** node at the top of the tree.

   Tip: You can change the order in which these merge types will display in the **View Compare/Merge** wizard. Select a merge type in the tree and click the **Up** or **Down** arrow on the right to move it.
6. Enter the description you want to display in the **View Compare/Merge Wizard**.
7. Choose the default **Merge Action** you want for each **Difference Type**.

   The default merge actions are used in the compare phase, are not visible to the user in the **View Compare/Merge Wizard**, and the user cannot change them.

Note: A merge action that has been changed from the default parent action is displayed with red text.

8. In the **Include Item Types** section, optionally check which item types to pre-select for this merge type. The user can view and change these types on the **Include Selected Items** page of the **View Compare/ Merge** wizard.
9. In the **Options** section, optionally check which options you want to pre-select for this merge type.

   The options selected are performed by default when the View Compare/Merge session begins the compare phase. The user can view and change these options on the **Set Options** page of the **View Compare/Merge** wizard.
10. Click **Save** when you are finished. The new merge type now appears in the **Available Merge Types** list of the **Customize VCM** tool.

Vault Verify for Verifying File Revisions

Vault Verify utility is a Java application that reports on and optionally attempts to resolve integrity issues with a StarTeam Native-II vault. It requires a StarTeam configuration name and, if the `starteam-server-configs.xml` file is not in the current folder, the path name of the folder containing this file. Vault Verify opens the corresponding database via JDBC, but it does not modify the database. Vault Verify also parses the `hive-index.xml` file to learn the location of vault hives.

This topic contains the following information:

- Checks Performed by Vault Verify.
- Vault Verify Requirements.
- Tips and Best Practices for Using Vault Verify.

Checks Performed by Vault Verify

Vault Verify is a command line utility that performs checks for corrupt, missing, or stray files for Native-II vaults. Optionally, Vault Verify can attempt to repair archive files based on what problem it finds with each file. For example, this utility will locate stray files and move them to a specified location. The administrator may then archive them off or delete them (after verifying their results).

Corrupt Files Check

This check validates all files in archive folders. For each file found in an archive folder, Vault Verify ensures that:

- The name of the file is a valid archive filename.
- The file is located in the correct folder based on its name.
- The file can be opened and read.
- The actual MD5 for the file matches its filename.
- If it is a compressed (.gz) file, its format is a valid GZIP format.

Note: If the `repair` option is requested, *corrupt* files are moved to the default or a configured *corrupt files* folder. Once moved, the corrupt file is classified as missing if it is referenced in the database.

Missing Files Check

This check ensures that all archive files defined in the database are present on disk. If the `repair` option is requested, Vault Verify will attempt to recover missing files from vault caches or other archive files.

Note: If you specify the `useca` (use Cache Agent) option, Vault Verify attempts to recover missing files from a remote Cache Agent.

Stray Files Check

This check ensures that all archive files in the vault are represented by corresponding database records. If the `repair` option is requested, *stray* files are moved to the default or a configurable *stray files* folder.

Vault Verify Requirements

Vault Verify requires the following:

- Vault Verify must have read access to the database used by StarTeam Server.
- You must download and install the Oracle JDBC driver for Oracle configurations. Go to *http://www.oracle.com/technology/software/index.html*, and scroll down to the **Drivers** section and click `JDBC`. Click the latest JDBC driver link. Following the download instructions, a page displays a list with `JAR` files. Download the `JAR` file which corresponds to the JDK version you are using.

  Note: You must have an Oracle.com user name and password before downloading the JDBC driver. If you do not have an account, you can create one from the Login page. Save the `JAR` file in the `VaultVerify` installation folder.
- Vault Verify must have read access to `starteam-server-configs.xml` and `hive-index.xml`.
- Vault Verify requires read access to the archive files for each hive and write access to the folders for each hive if you use the `repair` option.

Tips and Best Practices for Using Vault Verify

The following are tips and best practices for working with Vault Verify:

- You should run Vault Verify using the tailored batch file, `VaultVerify.bat` (or the shell script version on Linux) to ensure that the proper version of Java is used. The batch file (or shell script) is located in the Vault Verify installation folder.
- You must install the Vault Verify utility on the same system where you are running StarTeam Server. Vault Verify installs in its own Vault Verify folder under the StarTeam Server installation folder. For example, on a Windows system, Vault Verify installs in the `C:\Program Files\Micro Focus\StarTeam Server <version>\Vault Verify` folder.
- Vault Verify must have read access to the database used by StarTeam Server. By default, it uses the same userid as the StarTeam Server to access the database. If the password to that userid is not blank, it must be explicitly passed to Vault Verify. An alternate database userid can also be passed. Note that for Oracle configurations on Linux, Vault Verify requires the Oracle JDBC driver, which must be downloaded and installed by the customer.
- It is recommended that you run Vault Verify at least once per quarter and as often as once a month. It is also recommended that you run Vault Verify on a restored copy of the production database and the vault backup on a test box. Running Vault Verify on a test box not only ensures that the backup/restore procedure is working, but it offloads the I/O that Vault Verify does from the production server.
- If you are running Vault Verify against a mid- to large-size database, you should pass the Java `-Xmx1024m` parameter to avoid running out of memory.
- When using the *corrupt* check (this check opens and reads every archive file), Vault Verify returns results at 3 to 30 GB/hour depending on the system hardware and the size of the vault. When also using the `missing` and `stray` checks (these checks are much faster and perform file existence tests--they do not open or read files), each check adds another 5-30 minutes to the run time depending on the system hardware and the size of the vault.
- The requested check options are performed in the following order: `corrupt`, `missing`, and `stray`. Consequently, if `repair` is used along with the `corrupt` and `missing` checks, a corrupt file will first be moved to the corrupt files folder and then treated as a missing file.
- The specified StarTeam Server configuration can be in use when Vault Verify is running. However, the `stray` check and the `repair` option will be ignored if the StarTeam Server configuration is in use.
- All reporting, including problem files, displays in the console window (if so desired, you can pipe this information to a text file). If you request the `repair` option, the

results of any repair attempts are also displayed. The `verbose` option provides additional progress and diagnostic reporting. Vault Verify uses a stored procedure for reporting the share paths (project/view/folder path) of each valid archive file that is corrupt or missing. If this procedure is not present, the file name of problem files is reported, but share paths are not.

- The Vault Verify utility is contained in a set of jar files. The *main* file is `VaultVerify.jar`. It requires JRE 1.5 or newer. To get help text for Vault Verify, you can enter `java -jar VaultVerify.jar -help`. Usage text is also available in this help system.
- StarTeam Server always looks for the `starteam-server-config.xml` file in its own installation folder to determine whether the server is running. Be cautious about this fact if you decide to copy this file to a different location and then indicate to Vault Verify its new location with the `path` option. If you have indicated in the copied version of `starteam-server-config.xml` that the server is not running and use the `stray` and `repair` options in Vault Verify, these options are not ignored if StarTeam Server is running.
- The server configuration name passed to Vault Verify is case-sensitive, and if it includes spaces, you must pass the server configuration name to Vault Verify in quotation marks.
- By default, Vault Verify uses the same userid as the StarTeam Server to access the database. If the password to that userid is not blank, it must be explicitly passed to Vault Verify. An alternate database userid can also be passed.
- By default, the output from Vault Verify is output to the command window. We recommend that you pipe the output to a file so that if needed, you can send the information to Technical Support.

Verifying File Revisions with Vault Verify

The Vault Verify utility installs by default in the `C:\Program Files\Borland\StarTeam Server <version>\VaultVerify` folder on a Windows system.

In general, you can run Vault Verify from the command line as follows: `Vault Verify [options] "server configuration"`.

1. Open the Server Administration tool and shut down the server configuration you want to verify.

   You can use the specified StarTeam Server configuration when Vault Verify is running.

   Note: The `stray` check and the `repair` option are ignored if the server configuration is in use.

2. At the command prompt, navigate to the `VaultVerify` folder and type the following command:

   ```
   VaultVerify.bat -check all -cf C:\test -path
      "C:\Program Files\Micro Focus\StarTeam Server <version>" “My Server
   Configuration”
   ```

   On a Windows Server 2008 machine, type the following command instead:

   ```
   VaultVerify.bat -check all -cf C:\test –dbname <database name>  -dbuser
   <database user>  -dbinstance <instance name> -dbhost <host name>  “My
   Server Configuration”
   ```

   Tip: To view command-line options for the `VaultVerify` command, navigate to the `VaultVerify` folder and type `VaultVerify.bat - help`. Optionally, you can use –?, or –h instead of `–help`.

Exporting Database Information

The **Catalog Export** utility exports two application tables, `Catalog_Tables` and `Catalog_Fields`, into comma-delimited files. This tool is useful for database administrators because **Catalog Export** translates database tables and column names into identifiers used by the StarTeam Server. You can import and view the exported data in any application that supports comma-delimited fields. For example, if you save the file with a .csv extension, it will open in Microsoft Excel.

If you examine a column of data in the exported field catalog and find that one record has a surprising value (for example, all other records have a -1 in a column, but this record has a 16-digit number), the record may have been corrupted. However, we do not recommend that you delete any records, especially if you have not backed up the database.

1. Open the Server Administration tool and select the server configuration.
2. If the server configuration is running, click (**Shut Down Server**).
3. Click (**Catalog Export**). The **Catalog Export** dialog box opens.
4. Type, or browse for, the target path and location for the table catalog in the **File name for exported table catalog** field.
5. Type, or browse to, the target location and path for the field catalog in the **File name for exported field catalog** field.

   Note: Be sure to type the appropriate file extension for the application to which you want to export the files. By default, the utility specifies a .csv file.
6. Click **OK**.

Changing Server Configuration Properties

The session options for each server configuration are stored in the `starteam-server-configs.xml` file. You can modify a number of these options from the Server Administration tool or the command prompt with the `starteamserver` command.

When the server configuration is not running, you can modify the following session options using the Server Administration tool. Any changes that you make take effect the next time you start the server configuration. You can also change certain configuration options by using the **Start With Override** toolbar button.

1. Open the Server Administration tool and select the server configuration.

   Note: You must access the Server Administration tool on the computer where the StarTeam Server is installed.
2. If the server configuration is running, click .
3. Click (**Configuration Properties**). The **<Configuration Name> Properties** dialog box opens.

Change the name	Type a new name in the **Configuration name** field.
Change the log file path	1. Click **Change Path**. 2. Select a new folder for the server log file (`Server.locale.Log`).
Change the database or schema user	1. Select the **Database Connection Information** tab. 2. Type a new Host or TSN service name in the appropriate field. 3. Type a new user name and password in the **User name** and **User password** fields. If the server configuration uses an Oracle database, these boxes are named **Schema user name** and **Schema password**. 4. Click **Verify Connection**.

4. Click **OK**.
5. Restart the server configuration.

Setting Security Options

Server Time-Out Options

You can set the following time-out options for the server (**Tools** > **Administration** > **Configure Server** > **General**):

Server Logon Sequence Time-Out

The logon sequence time-out setting applies to both a client and the server configuration. This is the amount of time the client has to make the connection to the Server. If this time expires and a connection was not made, the user must try to log on again.

You use the **Logon sequence timeout** option on the **Configure Server** tab to set the logon sequence time-out value. This operation can be performed only when the server is running.

Server Inactivity Time-Out

The inactivity time-out is a security feature that automatically logs users off when they have been inactive for the length of time specified by the administrator. If a client has no communication (either automatic or manual) with the server configuration for that length of time, the server drops the connection. If the user's session has no other server connections, the session is deleted from the server. If the user has a concurrent license, that license is automatically returned to the pool of concurrent licenses. The user must then do a full login to reconnect.

You use the **Inactivity timeout** option on the **Configure Server** tab to set the inactivity time-out value. To allow named users (that is, users with a fixed license) to remain logged on even if they exceed the inactivity time-out limit, administrators can select the **Exclude named users** option after selecting the **Inactivity timeout** option and entering a time-out value.

Even if an inactivity time-out value is set, users will not time out if their system notifications are set for a period of time that is shorter than the inactivity time-out. For example, suppose a user has notification set to automatically check for new change requests every ten minutes and the inactivity time-out is set for 60 minutes. In this case, because of automatic communication between the client and the server, the user will never time out.

Server Reconnect Time-Out

If a client loses its network connection, users are disconnected from the server. The reconnect time-out option determines the amount of time the client has to reestablish the connection. The client attempts to reconnect only if the user is trying to send a command to the server. A reestablished connection contains the full context of the lost connection.

If the client successfully reestablishes its connection to the server within the window of time set in the Reconnect time-out, users can simply continue working in the application. They do not have to close their projects, log in again, and reestablish their view settings. However, if the **Reconnect time-out** has expired, you must either close the client, or log onto the server again.

You use the **Reconnect timeout** option on the **Configure Server** dialog box to set the reconnect time-out value. The reconnect time-out can be changed only on a server that is running. It does not work when the server has been restarted.

Note: When a server must be restarted, the client cannot automatically reconnect to the server.

When setting the **Inactivity timeout**, set it to a value greater than the **Reconnect timeout**. Otherwise, if the **Reconnect timeout** and the **Inactivity timeout** are both enabled and the **Inactivity timeout** is shorter, the user is logged off before the client can reestablish the connection. That is, if the **Reconnect timeout** is longer than the **Inactivity timeout** and both are turned on, then the **Inactivity timeout** acts before the **Reconnect timeout** time period has expired.

Number of Logon Attempts

You can increase the security of your projects by entering a logon failure setting and duration. One cause of logon failure is hackers trying to figure out passwords for users. In such cases, you should consider changing the IP address of the system to make it more difficult for attackers to locate the server configuration and repeat their efforts. You may also want to change the user names of all users in the system.

You choose **Tools** > **Accounts** > **System Policy** and then use the **Logon failures** tab to specify how to handle logon failures and the length of a lockout if one is applied. You can also specify that the server configuration notify members of the Security Administrators group by email about logon failures and lockouts. This operation can be performed only when the server is running.

It is possible for any user, even users with an administrative account, to be locked out of a server configuration when the number of retries with the wrong password has been exceeded. The lockout period for the main administrative account (Administrator) is 24 hours. However, you can unlock the administrative account before the 24 hours have elapsed (see *Reactivating Administrative Accounts*.)

Changing Server Timeout Options

Use the methods in this section to change time-out options for a server configuration.

Changing the Logon Sequence Time

1. Click **Start** > **Programs** > **Micro Focus** > **StarTeam Server <version>** > **StarTeam Server** . The Server Administration Tool opens.
2. Select a server configuration from the list of servers. If you have not yet logged on, you will be asked to do so.
3. Click **Tools** > **Administration** > **Configure Server** . The **Configure Server** page opens.
4. Select the **General** tab.
5. Type the number of seconds users have to log on in the **Logon sequence timeout** field.
6. Click **OK**.

Setting an Inactivity Timeout for Users

1. Click **Start** > **Programs** > **Micro Focus** > **StarTeam Server <version>** > **StarTeam Server** . The Server Administration Tool opens.
2. Select a server configuration from the list of servers. If you have not yet logged on, you will be asked to do so.
3. Click **Tools** > **Administration** > **Configure Server** . The **Configure Server** page opens.
4. Select the **General** tab.
5. Check **Inactivity timeout**.
6. Type the number of minutes in the **Inactivity timeout** field
7. Optionally, if you want to allow named users (that is, users with a fixed license) to remain logged on, even when they exceed the **Inactivity timeout** limit, check **Exclude named users**.
8. Click **OK**.

Changing the Reconnect Timeout

1. Click **Start** > **Programs** > **Micro Focus** > **StarTeam Server <version>** > **StarTeam Server** . The Server Administration Tool opens.
2. Select a server configuration from the list of servers. If you have not yet logged on, you will be asked to do so.

3. Click **Tools** > **Administration** > **Configure Server** . The **Configure Server** page opens.
4. Select the **General** tab.
5. Check **Reconnect timeout**.
6. Type the number of minutes in the field to set the reconnect timeout value. The default time is 30 minutes.
7. Click **OK**.

Configuring the Number of Logon Attempts

You can increase the security of your projects by entering a logon failure setting and duration. One cause of logon failure is hackers trying to figure out passwords for users. In such cases, you should consider changing the IP address of the system to make it more difficult for attackers to locate the server configuration and repeat their efforts. You may also want to change the user names of all users in the system.

You can configure the server configuration to notify members of the security administrators group by email about logon failures.

Note: You can perform this operation only on a running server configuration.

1. Click **Start** > **Programs** > **Micro Focus** > **StarTeam Server <version>** > **StarTeam Server** . The Server Administration Tool opens.
2. Select a server configuration from the list of servers. If you have not yet logged on, you will be asked to do so.
3. Select **Tools** > **Accounts** > **System Policy** . The **System Policy** dialog box opens.
4. Select the **Logon Failures** tab.
5. Select one of the following **Logon failures**:

Ignore	This selection disables the logon failures option.
Lockout account after ___ failures	Type the number of logon failures you want to allow.

6. Select one of the following **Lockout duration** options:

Forever	With this option selected, only an administrator can reinstate the user.
Keep locked for ___ minutes	Type the number of minutes for the duration of the lockout. The user will be able to log on again after the designated timeout period.

7. To notify members of the security administrators group that users attempted to log on unsuccessfully, check **By e-mail**.
8. Click **OK**.

Setting an Encryption Level

Encryption protects files and other project information from being read by unauthorized parties over unsecured network lines, such as the Internet. For TCP/IP connections, you can set a minimum level of encryption for a server configuration for IP addresses that access that server configuration. You can set different encryption levels for an IP address, ranges of IP addresses, or all IP addresses.

Clients can set the encryption level on a per-workstation basis. Users must use at least the minimum level of encryption set for underlying server configuration.

Note: You can perform this operation only on a running server configuration.

Setting an Encryption Level for Transferred Data Regardless of IP Address

1. Click **Start** > **Programs** > **Micro Focus** > **StarTeam Server <version>** > **StarTeam Server** . The Server Administration Tool opens.
2. Select a server configuration from the list of servers. If you have not yet logged on, you will be asked to do so.
3. Click **Tools** > **Administration** > **Configure Server** . The **Configure Server** page opens.
4. Select the **Protocol**tab.
5. Select **Default** in the **TCP/IP encryption levels** list.
6. Click **Modify**. The **Set Encryption Type** dialog box opens.
7. Select the type of encryption you want to use with the server configuration for IP addresses not specified in this list.
8. Click **OK**.

Setting a Different Encryption Level for a Specific Address or Range of Addresses

1. Click **Start** > **Programs** > **Micro Focus** > **StarTeam Server <version>** > **StarTeam Server** . The Server Administration Tool opens.
2. Select a server configuration from the list of servers. If you have not yet logged on, you will be asked to do so.
3. Click **Tools** > **Administration** > **Configure Server** . The **Configure Server** page opens.
4. Select the **Protocol**tab.
5. Click **Add**. The **Set Encryption Type** dialog box opens.
6. Type the starting IP address in the **Starting IP** boxes.
7. Type the ending IP address in the **Ending IP** boxes.
8. Select the type of encryption to be used with the server configuration for these addresses.
9. Click **OK**.

Enabling Server Auto-reconnect

If a client loses its network connection, users are disconnected from the Server. The reconnect time-out option determines the amount of time the client has to reestablish the connection. The client attempts to reconnect only if the user is trying to send a command to the server. A reestablished connection contains the full context of the lost connection. If the client successfully reestablishes its connection to the server within the window of time set in the **Reconnect timeout** option, users can continue working in the application. They do not have to close their projects, log in again, and reestablish their view settings.

Note: You can change the reconnect time-out for running server configurations. It does not work when the server has been restarted. When a server must be restarted, the client cannot automatically reconnect to the server. Also, if you enabled the **Reconnect timeout** and the **Inactivity timeout** options and the **Inactivity timeout** time is shorter, the user may be logged off before the client can reestablish the connection.

To change the reconnect timeout

1. Open the Server Administration tool and select the server configuration that you want to modify.

   Note: If you are using the client, you will be able to administer remote servers only.

2. Click the **Configure Server** shortcut in the shortcut pane, or choose **Tools** > **Administration** > **Configure Server** from the main menu.

   This opens the **Configure Server** dialog box.
3. Select the **General** tab.

4. Check **Reconnect timeout**.
5. Type the number of minutes in the text box to set the **Reconnect timeout** value. The default time is 30 minutes.
6. Click **OK** to apply your changes.

Working with Event Handlers

StarTeamMPX has an event transmitter that must be installed on the same computer as the Server. In addition, the Message Broker can be installed on the same or another computer, depending on your needs. If you install the Message Broker, **Unicast On-site** event displays in the **Event Handlers** tab of the **Configure Server** dialog box.

For more information about StarTeamMPX, its XML files, properties, and values, see the *MPX Administrator's Guide*. The following topic describes how to work with event handlers. The **Event Handlers** tab provides a simple interface for editing the `StarTeamMPXTransmitter.XML` files.

Note: You can perform these operations only on a running server configuration.

Creating Event Handlers

The **Event Handlers** tab provides a simple interface for editing the `StarTeamMPXTransmitter.XML` files.

Note: You can perform this operation only on a running server configuration.

1. Open the Server Administration tool and select the server configuration.
2. Click **Tools** > **Administration** > **Configure Server** . The **Configure Server** page opens.
3. Click the **Event Handlers** tab.

   To create a new event handler from scratch:

   1. Click **Add** to open an empty **Event Handler Profile Properties** dialog box.
   2. Type a name and description in the appropriate fields.
   3. Click **Add** to display an empty **Event Handler Property** dialog box.
   4. Type the property name and its value in the fields.
   5. Repeat until you have added all the properties you need.

   To create a new event handler from an existing one:

   1. Select an existing event handler that is very similar in its properties to the new handler that you need.
   2. Click **Copy**. The **Event Handler Profile Properties** dialog box opens displaying the properties for the selected event handler.
   3. Change the name and description in the appropriate text boxes.
   4. Select and modify other properties as appropriate.
4. Click **OK**.

Reviewing or Modifying Existing Event Handlers

For more information about StarTeamMPX, its XML files, properties, and values, see the *MPX Administrator's Guide*. This topic describes how to review or modify existing event handlers. The **Event Handlers** tab provides a simple interface for editing the `StarTeamMPXTransmitter.XML` files.

Note: You can perform this operation only on a running server configuration.

1. Open the Server Administration tool and select the server configuration.
2. Click **Tools** > **Administration** > **Configure Server** .

The **Configure Server** page opens.

3. Click the **Event Handlers** tab.
4. Select an existing event handler.
5. Click **Modify**. The **Event Handler Profile Properties** dialog box opens.

   To change a setting

   1. Select a setting from the **Profile Properties** list.
   2. Click **Modify**. The **Event Handler Property** dialog box opens.
   3. Change the value.

   To add a property

   1. Click **Add**. An empty **Event Handler Property** dialog box opens.
   2. Type a property name and value in the appropriate check boxes.

   To remove a property

   1. Select a setting from the **Profile Properties** list.
   2. Click **Remove**. Be aware that you cannot delete a profile that is currently used as the default profile.

6. Click **OK**.

Assigning Default Event Handlers for the Server/Clients

1. Open the Server Administration tool and select the server configuration.
2. Click **Tools** > **Administration** > **Configure Server** . The **Configure Server** page opens.
3. Select the **Event Handlers** tab.
4. Select an existing event handler.
5. Click one or both of the following:

   Server Default Makes the selected profile the profile for the server. A server icon displays in front of the default server profile.

   Client Default Makes the selected profile the default profile for clients. A green check mark displays in front of the default client profile. As users create server descriptions on their workstations, the profile selection defaults initially to this profile. Users can change from the default to another existing profile. If a profile is both the server and client default, you see only the server icon.

6. Click **OK**.

Note: A file transmitter does not use profiles. It interacts with the event transmitter which uses the **Server Default** profile.

Removing an Event Handler

1. Click the **Event Handlers** tab.
2. Select an existing event handler.
3. Click **Remove**.

StarTeam Offline Proxy Utility

Use the Offline Proxy utility to communicate to StarTeam users that the server is down for maintenance. Previously when StarTeam Server is down due to database issues or maintenance, a user connecting with StarTeam SDK or StarTeam Cross-Platform Client gets a message stating, "Unable to connect to server." Since they do not know what the reason is, they might raise issues with their IT department . When using the Offline Proxy utility, administrators can better communicate the reason for the disruption in the Server, avoiding unnecessary IT service tickets.

StarTeam Offline Proxy utility runs on the same port as the defined configuration, and it responds to any command from clients with an error message to the user.

Run the StarTeam offline utility from command line.

```
StarteamOffline - start ConfigName
```

You can also start and stop multiple proxies by using `-all` parameter.

```
starteamoffline [-help] [-list]
               [-start config]
              [-stop config] [-tcpip endpt]
              [-version]

        -help                   display this help message
        -list                   list all configurations
        -q                      run quietly
        -start name             start offline server with config [name]
        -start -all             start offline server for all configurations
        -stop name              start offline server with config [name]
        -stop -all              stop offline server for all configurations
        -tcpip endpt            TCP/IP sockets endpoint
        -version                display version info
```

If a Server Message Notification has been defined before the server was shutdown, the users will see that message. Otherwise, they will see a default message, "StarTeam Server is currently offline for maintenance."

Starting a StarTeam Server Configuration will automatically stop the offline proxy.

Server Message Notification

You can create a message from the server and send it to users to notify them of server outages, scheduled maintenance, updates, and etc.

Open the Server Administration Tool and navigate to Server Configuration. From within the configuration, you can optionally select a start and end date for the message and type your message text.

Users will see the message when logging in or when the StarTeam Offline Proxy is running.

Email Support Notifications

To take advantage of email notifications, you must enable email support and email notification in the **Email** tab of the **Configure Server** dialog box in the Server Administration Tool.

Note: Client-calculated fields cannot be used in custom email notifications or with Notification Agent.

Email Support

When you enable email for a server configuration, users can email the properties of an item to another user from within the application. The email recipients do not need to be running the application to receive the email.

The application sends automatic email to users when their exclusive locks on items are broken. Users can only break locks if they have the correct access rights and privileges to do so. You can also configure the application to perform automatic email notification when certain other events occur. Depending on the server configuration and system policy options you select:

- Members of the *System Managers* group can receive email whenever an error is added to the server log.
- Members of the *Security Administrators* group can receive email whenever a logon failure occurs.
- All users can receive automatic notifications about items for which they are responsible or for which they are recipients.

Note: If a recipient of an item or notification has an incorrectly formatted email address, an entry is written to the server log indicating that there was a problem sending email to that address. If an email address is formatted correctly but is invalid (as in “junk@place.com”), the email is sent to all valid recipients, and the sender gets an “Undeliverable message” from the email system for the invalid address.

Email Notifications

If you enable email notification, a user automatically receives email if:

- The **Responsibility** field value changes in a change request
- Any field for a requirement or task for which the user is responsible has changed.
- Any field for a topic for which a user is listed as a recipient has changed. (If no recipients are listed for a topic, no one receives notification)

Because email notification is client-independent, your team members do not have to run a client to receive notification messages.

Default messages sent to recipients of automatic email notification are localized, based on the locale of the server. When no translation is available for a locale, the message is in English.

Tip: The language used with a specific server configuration can be changed by adding `NotificationLocale` to the section of the `starteam-server-configs.xml` reserved for the configuration. For example, if you add `NotificationLocale=ja`, the messages are sent in Japanese.

Users may confuse email messages sent by individuals (using the **Send to** command in the client) with email notification messages, because unless you choose to customize the email message templates, they are somewhat similar. Therefore, it is a good idea to let users know when you enable automatic email notification and to explain the differences between the two types of email messages and the two types of notification.

Note: You can dynamically customize the email notifications on a per-server configuration, per-project, or per-component basis. Edit the templates provided in your repository under the `Notifications` folder. You can use fields stored in the StarTeam database within the custom templates.

Custom Email Notifications

You can configure customized, automatic email notifications on a per-server configuration, per-project, or per-component basis. You can design your own text or HTML-based message templates or use the default templates provided in the `Notifications` folder, a subfolder of the server installation folder. All email notification messages (both plain text and HTML) are sent in UTF-8 encoding.

You can define custom email templates to use with email notifications for the following components:

- Change Request
- Task
- Topic
- Requirement

When a server configuration starts for the first time, the contents of the `Notification` folder in the installation directory are copied to the repository for the server configuration in a corresponding `Notification` folder. You can make customizations to the default templates in the `Notification` folder found in the server configuration repository. The predefined email notification files consist of a set of component-level XML configuration files – one for each desired component and an arbitrary number of email message body templates that can have any name that you choose. However, the configuration files must be named as follows:

- `ChangeRequest.xml`

- `Requirement.xml`
- `Task.xml`
- `Topic.xml`

The predefined email message body templates are named:

- `itemTypeAbbr-new.txt`
- `itemTypeAbbr-modified.txt`
- `itemTypeAbbr-new.html`
- `itemTypeAbbr-modified.html`

Where `itemTypeAbbr` corresponds to `cr`, `req`, `task`, or `topic`.

Each time that you start the server configuration, it scans the contents of the `Notification` folder. If the configuration (.xml) files are invalid and you have email notification enabled for your server configuration, the server issues an error message in the server log and fails to start.

You can also make dynamic updates to the message templates in the repository `Notification` folder without restarting the server configuration. The server checks for changes in the configuration and message template files every two minutes and immediately applies valid changes found in the files.

If the server finds a corrupted configuration and/or template file while the server configuration is running, a predefined email message is sent to the Admin group. Email notification becomes unavailable until you restore a valid configuration in the `Notification` folder. While any of the files in the `Notification` folder are in an invalid state, the server sends users the standard email notifications.

Fields Used in Email Notification Message Templates

Within each of the sample templates provided in the `Notification` folder in your repository, you will find fields that you can use in your own customized templates. Three types of fields are used:

- Fields stored in the database
- Client-calculated fields
- Server-calculated fields

Note: You cannot use client or server-calculated fields within customized email notification message templates. The templates recognize only fields stored in the database.

Embedded Images in Email Notification Message Templates

You cannot embed images in the HTML email message templates. However, you can use a URL to an image on a web site that users can access.

Configuring Email Support and Email Notification

This topic assumes that you have the Server Administration Tool open that you have selected and logged onto the server configuration that you want to change. If you are using the client (available with custom installations only), you can administer remote server configurations only.

1. Click **Tools** > **Administration** > **Configure Server** . The **Configure Server** page opens.
2. Select the **Email** tab.
3. Check the option to **Enable e-mail support**.
4. Type the host name for your SMTP (Simple Mail Transfer Protocol) server in the **SMTP server** field.

   You can use an IP address if your site uses only static IP addresses. The application uses SMTP, which traditionally operates over TCP using port 25. It is widely used and is the Internet’s standard host-to-host mail transport protocol.

   Note: For Microsoft Windows environments, the Exchange server is usually the SMTP server.

5. Optionally, type a value in the **TCP/IP endpoint** field if your SMTP server uses a port other than the default value, `25`.
6. For secure ports, select either **SSL** or **TLS**.
7. If you want to use e-mail notification, check the option to **Enable e-mail notification**.
8. If you selected SSL or TLS, you have additional **Security Settings**:

   a) **SMTP Certificate Name**

   Name of certificate file to be used for the secure connection. This setting is required on Windows only .

   StarTeam Server will look for certificates in subfolder: `install path\CACertificates`. You should place the certificates in this folder. On Microsoft Windows, two certificates are pre-installed:

   Equifax `EquifaxCA.pem` - This certificate is useful to connect to GMail SMTP server.

   Verisign `VeriSignClass3PublicPrimaryCA.pem` - This certificate is used to connect to Yahoo SMTP server.

   b) **SMTP HostName**

   SMTP Hostname as it appears in the SMTP server’s certificate. If the hostname in the certificate does not match with the entered one, we will write a warning to the server log similar to: `62 00000002 2016-02-19 10:30:18 SMTP TLS Unrecognized certificate: error code=13 and error message: Certificate name does not match`.

   Hostname not matching is a potential security issue: if the hostname in the certificate is not the expected one, there could be a "man-in the middle" security attack - this means another site is impersonating the SMTP server.

   c) **SMTP Logon Name**, **SMTP Password**

   These are the credentials to be used to login to the SMTP server. These are optional, they should be entered if the SMTP server requires them.
9. Click **OK**.

Configuring Per-project and Per-Component Email Notifications

This topic describes how to configure email notifications on a project-specific and component-specific basis using the change request component as an example. You should have already enabled email support and email notification messages.

1. Navigate to the `Notifications` folder installed with StarTeam Server. The `Notifications` folder installs as a subfolder of your StarTeam Server installation. By default, StarTeam Server is installed in the `C:\Program Files\Micro Focus\StarTeam Server <version>` folder.

   Tip: Make a copy of the `Notifications` folder before making any modifications. You can revert to this copy if you make any undesirable changes.
2. You can edit the component-specific `*.xml` file for the component (change request, task, topic, requirement) that you want to use for project-specific notifications. Open `ChangeRequest.xml` and type the following rules for a specific project:

   ```
   <rule-list>
    <rule project=”MyProject” event=”new” template=”MyProject-cr-new-txt”/>
    <rule project=”MyProject” event=”modified” template=”MyProject-cr-modified-
   txt”/>
   ```

   In the above example, `“MyProject”` corresponds to your specific project name. These entries must go before the following default `<rule project="*" event="new" template="cr-new-html"/>` and `<rule project="*" event="modified" template="cr-modified-html"/>` entries.
3. Enter the template information used for your project under the `<template-list>` tag. For example,

   ```
   <template-list>
     <template-id=”MyProject-cr-new-txt”>
   ```

```
    <subject>New Change Request #~~ChangeNumber~~</subject>
    <body content-type>”text/plain” template-file=”.\MyProject-cr-new.txt”/>
  </template>
  <template-id=”MyProject-cr-modified-txt”>
    <subject>Modified Change Request #~~ChangeNumber~~</subject>
    <body content-type>”text/plain” template-file=”.\MyProject-cr-
modified.txt”/>
  </template>
```

4. Save the changes and close the template files.
5. Copy your new template files and updated `ChangeRequest.xml` file to the `Notifications` folder in your repository.

Designating Endpoints

The default TCP/IP port (endpoint) is `49201`, but you can specify a different port for a server configuration. If you have more than one server configuration running on the same computer, each server configuration must use a unique endpoint. For example, if `Server Configuration 1` uses the endpoint `49201`, `Server Configuration 2` must use a different endpoint. If you attempt to run server configurations that have the same endpoint and computer name at the same time, only the first server configuration you select will start successfully. The remaining server configuration will appear to start, but in fact is ignored by the StarTeam Server.

Note: This operation can be performed only on running server configurations. The changes take effect once you restart the server configuration.

1. Open the Server Administration tool and select the server configuration.

   Note: If you are using the client, you will be able to administer remote servers only.

2. Click **Tools** > **Administration** > **Configure Server** . The **Configure Server** page opens.
3. Select the **Protocol** tab.
4. Type a port number in the **TCP/IP (Sockets)** field to activate a different port.

   The range for port numbers is 1023 through 65535.

5. Optionally, click **Default** if you wish to return to the default endpoint setting (49201).
6. Click **OK**.

   Note: You must restart the server configuration for this setting to take effect.

Server Hooks

Server hooks are scripts that run automatically every time a particular event occurs in the StarTeam repository. Server hooks allow the user to trigger customizable actions at key points in the development life cycle.

StarTeam Server implements one server hook: `post-commit`

The `post-commit` script automatically runs after a Change Package status is set to **commited**. For example, you may want to use `post-commit` to send a notification email or to start a build.

The name of the script file is `post-commit.bat` on Windows and `post-commit.sh` on Linux. The script must be copied to the `ServerHooks` folder under your repository path.

After committing a change package, StarTeam Server checks the `ServerHooks` folder, within the repository directory, to see if a `post-commit` script exists. If yes, it will execute the script in a separate process.

The `post-commit` batch file script will be executed for every committed Change Package with the following arguments:

- - argument 1 - Project Name
- argument 2 - Target View Name
- argument 3 - User Name
- argument 4 - Class Name (i.e., ChangePackage)
- argument 5 - Transaction Time (using machine local time)
- argument 6 - Transaction ID
- argument 7 - Project ID
- argument 8 - Target View ID
- argument 9 - Item ID
- argument 10 - Class ID for ChangePackage
- argument 11 - Source View ID
- argument 12 - Source View Name
- argument 13 - Change Package Name

There are a couple post-commit scripts samples installed in StarTeam Server installation directory under the `ServerHooks` folder:

post-commit - vcm replay.bat.sample	uses View Compare/Merge to rebase every change package to a specific sandbox.
post-commit.bat	sends an email notification for each change package committed.

To help debug post-commit scripts that you create:

- When setting `VerboseLevel` to **1** in `starteam-server-configs.xml` for the configuration, the parameters for each call to post-commit will be written to the server log. It will look something like this:

  ```
   7 00000011 2017-06-22 15:57:24
  Launching c:\repository\test\ServerHooks\post-commit.bat
  ProjectName='test' ViewName= 'Myview' UserName=' joe'
  TransactionID=169859 ProjectID=5 ViewID=40 ItemID=15140 ClassID=89
  SourceViewID=40 SourceViewName='New view2' CPName='Workspace
  Changes for Change Request 1,206 test on 2017/06/22 22:55:35 GMT
  1498172135634'
  ```

- If you redirect the output of your script to a file, make sure the file name is unique. Otherwise, multiple scripts executing in parallel will fail due the inability to access the output file.

Monitoring Server Statistics

The StarTeam Server provides an HTML report to monitor server statistics. This report tracks memory usage, currently executed commands, locking statistics, and so on. By default reports are saved in the `Diagnostics\ServerStatisticsMonitoring` StarTeam Server installation folder.

1. Open the Server Administration tool and select the server configuration.
2. Click **Tools** > **Administration** > **Configure Server** .

   The **Configure Server** page opens.
3. Click the **Diagnostics** tab.
4. Check **Enable Statistics Monitoring. Record every:** and specify the time interval to record the statistics.
5. Click **OK**.
6. Click **Actions** > **Statistics Monitoring** . An HTML report opens in your web browser containing server statistics.

Note: Other diagnostic tests should be specified and generated at the direction of Micro Focus technical personnel.

Troubleshooting Server Configuration Problems

To reduce the amount of time spent diagnosing problems, the application provides tracing and debugging tools for the server. It can create either, or both, trace command files and diagnostic (.dmp) files.

Trace Commands

The trace option creates a file that records single server commands. Commands to be traced must have a duration time that equals or exceeds the number of specified milliseconds. The default time is 0. If you wish to record only commands of longer duration, you should adjust this setting, to avoid taking up unnecessary space in the trace file.

No trace file should generate more than 10 MB of data per day. Typically, users see only a small fraction of this amount of data per day.

Trace data is stored in a `Server.trc` file, which consists of a header followed by an arbitrary number of records. When a trace ends, the server timestamps the existing file as `Server.time.trc`. Trace files are located in the `repostitoryPath\Log\Trace` folder. The next trace file starts when the server configuration is restarted or the trace option is turned on.

Diagnostic (.dmp) Files

The application creates some minidump files automatically, while others are created only when the .dmp options are turned on. Minidump files can be created for either or both:

- Asserts (unexpected conditions). Server log entries with code number 8.
- Exceptions (errors, typically access violations). Server log entries with code number 4.

Minidump files are created in the same location as the server log file. The general naming convention for these files is `prefix-counter-time.dmp`, in which prefix identifies the source of the dump, counter is an integer that increments with each .dmp file to ensure that names are unique, and time identifies the local server time at which the dump was created.

Activating Diagnostic Tests

To reduce the amount of time spent diagnosing problems, the application provides tracing and debugging tools for the server. It can create either, or both, trace command files and diagnostic (.dmp) files. By default, both of these options are turned off. If you encounter a problem, you can simply turn them on and create files that you can review or discuss with a Micro Focus SupportLine representative.

Note: This operation can be performed only when the server configuration is running. For instructions on enabling tracing manually by editing the `starteam-server-configs.xml` file, refer to *Enabling Tracing for Server Configurations Manually*.

1. Open the Server Administration tool and select the server configuration.

   Note: If you are using the StarTeam Cross-Platform Client, you will only be able to administer remote servers.

2. Click **Tools** > **Administration** > **Configure Server** .

   The **Configure Server** page opens.

3. Click the **Diagnostics** tab.
4. To create `Server.trc` files, check **Trace operations that take at least**. If you do not want to use the default milliseconds value, type a different number.
5. To create diagnostic `.dmp` files, check either or both of the following options:

   - **Unexpected conditions (server log entries with code #8)**
   - **Errors (server log entries with code #4)**

6. Click **OK**.

Note: Other diagnostic tests should be specified and generated at the direction of Micro Focus technical personnel.

Change Request Configuration File

This topic describes the `ChangeRequest.xml` configuration file. Your server configuration repository contains a `Notifications` subfolder containing this and other configuration files.

ChangeRequest.xml

```
<?xml version="1.0" encoding="UTF-8" standalone="no" ?>
<notification-config version="1.0">
  <rule-list>
    <rule project="*" event="new" template="cr-new-html"/>
    <rule project="*" event="modified" template="cr-modified-html"/>
  </rule-list>
  <template-list>
     <template id="cr-new-html">
       <subject>New Change Request #~~ChangeNumber~~</subject>
       <body content-type="text/html" template-file=".\cr-new.html" />
     </template>
     <template id="cr-modified-html">
       <subject>Modified Change Request #~~ChangeNumber~~</subject>
       <body content-type="text/html" template-file=".\cr-modified.html" />
     </template>
     <template id="cr-new-txt">
       <subject>New Change Request #~~ChangeNumber~~</subject>
       <body content-type="text/plain" template-file=".\cr-new.txt" />
     </template>
     <template id="cr-modified-txt">
       <subject>Modified Change Request #~~ChangeNumber~~</subject>
       <body content-type="text/plain" template-file=".\cr-modified.txt" />
     </template>
  </template-list>
</notification-config>
```

Note: `template` is a mandatory `rule` attribute. It specifies the message template that you wish to use for the notification and it must correspond to one of the template nodes, such as `<template id=”cr-new-html”>`. `event` and `project` are optional `rule` attributes. `event` specifies whether the item triggering the notification is being created (`“new”`) or edited (`“modified”`) or either of the two (`“*”`). If omitting the `event` attribute, the notification applies to any event. `project` allows the notification to be limited to a certain project or projects only. If omitted, the notification applies to all projects.

Note: You can give the `template-file` attribute an absolute file path or a path relative to the `Notifications` folder.

Examining the Configuration File

When you create or modify a change request, the server searches through the rules defined in the `<rule-list>` element for the first project attribute that matches the name of the current project that consists of more than the asterisk wildcard character. This allows you to use `project=”*”` for any projects where you have not created a more specific rule. For example, using the sample code shown above, a project names `Project2` would use the text templates and a project named `Whitestar` would use the HTML templates.

Next, the server searches the matching rule to determine if it applies to the current `“new”` or `“modified”` change request. The `event` attribute of the rule determines this by the values `“new”` or `“modified”`. If this value is absent or set to `“*”`, all change requests use the rule. If the current change request is new, a rule with an `event` attribute of `“new”` or `“*”` must be found.

Once the server finds the correct rule, the `template` attribute indicates where to look in this .xml file for the subject line to use in the email, the type of template file, and the path to the template file. In the sample `ChangeRequest.xml` file, if the `template` attribute in the selected rule was `“cr-new-txt”` the server uses the following template section to obtain more information:

```
<template id=”cr-new-txt”>
<subject>New Change Request #~~ChangeNumber~~</subject>
<body content-type=”text/plain” template-file=”.\cr-new.txt” />
</template>
```

The subject line of the email reads `New Change Request` followed by the change request number. The template content type is `“text/plain”` and the template file is located at `“.\cr-new.txt”`. You can use the absolute path or a path relative to the *Notifications* folder for the location of the template file. This example uses a relative path to locate the template file within the `Notifications` folder.

Change Request Message Template Syntax

This topic describes the `cr-new.txt` message template file and examines message template syntax in general. You can use HTML or text format for your custom templates. You server configuration repository contains a `Notifications` subfolder containing this file and other message template files in both HTML and text format for change requests, requirements, tasks, and topics.

cr-new.txt

```
This message has been sent to you automatically by StarTeam Server
because the Change Request described below has been created
by ~~ModifiedUserID~~.
The Change Request is located in
Project: ~~project~~
View: ~~view~~
Property Summary:
Type: ~~type~~
Status: ~~status~~
Responsibility: ~~responsibility~~
Priority: ~~priority~~
Severity: ~~severity~~
Platform: ~~platform~~
Entered By: ~~EnteredBy~~
Entered On: ~~EnteredOn~~
Synopsis: ~~synopsis~~
Description: ~~Description~~
You can access Change Request #~~ChangeNumber~~ by following this URL:
~~url~~
```

Examining the Message Template Syntax

Whenever you want to use the current value for a property field, enter the internal name for that field preceded and followed by two tildes (~~). For example, `Status: ~~status~~` displays the word “Status:” in the email followed by the current value of the status property.

The file also uses `~~url~~` to include the StarTeam URL for the change request.

Note: Client-calculated property fields cannot be used in the notification message templates. Refer to the Fields Reference link at the bottom of this topic for a list of all fields. The description for the field includes its internal name and whether it is a client-calculated property.

Displaying Values for Modified Fields

The text of `cr-modified.txt` uses the following syntax:

```
~~isnew.AddressedBy?Addressed By: ~~~~isnew.AddressedBy~~~~isnew.AddressedBy?
~~
```

In these expressions, the `isnew` prefix precedes the internal name of the property field. This allows a template to specify that the property field value and/or additional formatting text should be included in the resulting notification email only if the value of the field has been changed, that is, only if the previous version had a different value for this field.

The text `Addressed By:` (from the first expression, `~~isnew.AddressedBy?Addressed By: ~~`), the value of the `AddressedBy` property (from the second expression `~~isnew.AddressedBy~~`) and a trailing linefeed (from the last expression) display in the email notification message only if the value of the `AddressedBy` property has been changed from the previous revision of the change request.

If you are using an HTML message template instead of the text version, you could display the output of the expression in bold text (only if it was changed for this revision) as follows:

```
Addressed By: ~~isnew.AddressedBy? <b>~~~~AddressedBy~~~~isnew.AddressedBy?</
b>~~
```

The above example behaves as follows in an email notification:

Addressed By:	(Text or formatting always included)
~~isnew.AddressedBy?<b>~~	(Text or formatting included only if the `AddressedBy` field changes in this revision)
~~AddressedBy~~	(Value always included)
~~isnew.AddressedBy?</b>~~	(Text or formatting included only if the `AddressedBy` field changes in this revision)

Definition of Message Template Syntax

In general, you can enter the following information into the message template files:

Note: In the following examples, `propname` could be `type`, `status`, `responsibility`, and so on. Refer to the example `cr-new.txt` file at the beginning of this topic for more examples of property field names.

~~propname~~	Replaced with the value of the specified property.
~~old.propname~~	Replaced with the old value of the specified property.
~~new.propname~~	Replaced with the new value of the specified property.
~~project~~	Replaced with the name of the project containing the item. This is not available as a normal item property.
~~view~~	Replaced with the name of the view containing the item.
~~url~~	Replaced with a StarTeam URL of the item.
~~isnew.propname~~	Replaced with the value of the `propname` property only if the property has a new value. Otherwise, this expression resolves to an empty string.
~~isnew.propname?text-string~~	Replaced with the text-string text only if the `propname` property has a new value. Otherwise, this expression resolves to an empty string.

Enabling Advanced View Types

By default, advanced view types are not available for server configurations. However, you can allow users to create views of advanced types by editing the `starteam-server-configs.xml` file.

1. Open the `starteam-server-configs.xml` file in an editor. By default this is located in the root installation folder for StarTeam Server. For example, on a Microsoft Windows system, you would find this file in the `C:\Program Files\Micro Focus\StarTeam Server <version>` folder.
2. For each server configuration that you want to allow advanced view types for, enter the following:

   ```
   <option name=“DisableAdvancedViews” value=“”/>
   ```

   If you specify the `value` as `“”` , the **Show advanced types** check box appears in the **New View Wizard**, and the **Branch All, Float**, **Branch None**, and **Non-Derived** advanced view types are available in the wizard. If you specify the `value` as `“1”` then the **Show advanced types** check box does not appear.

StarTeam SDK Connection Control

The StarTeam Server allows administrators to fine tune the set of client applications that can connect to the server by customizing a new `app-control.xml` file. This feature prevents unwanted SDK applications from connecting to the StarTeam Server and draining resources.

Note: This is strictly an administrative tool, not a security measure.

app-control.xml Configuration File

The StarTeam Server looks for a configuration file named `app-control.xml` located in the `AppControl` directory under the `StarTeam` repository root directory. When a new configuration is created, StarTeam Server creates this file from a template `app-control.xml` file located in `AppControl` directory under the StarTeam Server installation directory.

The configuration `app-control.xml` file, if present, contains a set of rules. Each rule asks the server to test the incoming client connections to satisfy one or more of the following conditions:

- The StarTeam SDK is greater than or equal to a certain version.
- The application name, connecting user name/client workstation, name must match a specified text pattern.

The StarTeam Server tests each incoming client connection against all the rules present in the `app-control.xml` file until a match is found or until the rule list is depleted. Once a match is found, no more checks are done and the connection handshake sequence is resumed. If no match is found, the connection is refused. If the `app-control.xml` file does not exist in the `AppControl` directory, the StarTeam Server allows all supported client applications to connect.

app-control.xml Structure and Rule Syntax

The root XML element must be named `StarTeamApplications` and have a `version` attribute with a value equal to `1.0`. For example, `<StarTeamApplications version="1.0">`

The server recognizes the following elements directly under the root node:

AppDefault

`AppDefault:` This is an optional element that can be used to specify default values for one of the parameters listed under `AllowedApp`. The syntax of this element is similar to the `AllowedApp` syntax, except that the `Name` attribute cannot have a default value. Default values can be specified for `MinimumSDKVersion`, `WorkStationID`, and `UserName`.

AllowedApp

This is the main rule element. It must have a `Name` attribute that specifies the text pattern for the client application name (such as “client identification string”). The text pattern can have an asterisk character (‘*’) that is used as a wildcard. If an optional parameter is not set, the StarTeam Server does not test the corresponding connection attribute.

Besides the `Name` attribute, this node can optionally specify one or more of the following attributes:

MinimumSDKVersion Specifies a minimum version of StarTeam SDK with which the client application is built. The format of this field is `nn.nn.nn.nn`, where nn is a non-negative number. Not all of the “dot” numbers have to be specified, for example `MinimumSDKVersion=”10.4”` will allow `10.4.x.y` and above (`10.5`, `11.0`, and so on).

WorkStationID If set, specifies text pattern to match the client computer name.

UserName If set, specifies text pattern to match the StarTeam user name.

BlockedApp The `BlockedApp` element provides the ability to block a specific application. It must include the `Name` attribute.

For `BlockedApp` there are 4 recognized attributes: `Name`, `SDKVersion`, `WorkStationID` and `UserName`.

- Attributes `Name`, `WorkStationID` and `UserName` can contain a specific string to be compared with, or a string with an asterisk ("*"). The asterisk in the string will match a pattern in that string.
- Attributes `WorkStationID` and `UserName` can be empty or not specified at all. This will block any `WorkStationID` or any username.
- Missing attributes other than `Name`, will either use values from `AppDefault` attributes if they exist or will assume the value `"*"`.

Attribute `SDKVersion` corresponds to the SDK version to block and has a specific format. It is composed of four-part dot-noted positive numbers `"1.2.3.4"`. This attribute specifies the SDK version to block. If any of the parts are skipped, any version for that part will be blocked. For example:

```
SDKVersion="10" will block 10.1.0.0 but will not block 9.10.5.0
SDKVersion="10.4.50.0" will not block the SDK version specified
SDKVersion="11.4" will block 11.4.5.0 or 11.4.1.0 but will not
block 11.3.0.0
```

Examples

```
<AppDefault MinimumSDKVersion="11" />
            ## defaults to accepting any SDK version with first part 11 or
greater
<AppDefault WorkStationID="*" />
            ## default accepts any WorkstationID. Not specifying a default
accepts any value as well.
<AllowedApp Name="Bulk Checkout Utility" WorkStationID="build-*"/>
            ## Accepts the application with specific string "Bulk Checkout
Utility"
            ## and matching specified worsktation pattern and default
patterns for other attributes
<AllowedApp Name="CPC*" MinimumSDKVersion="8.0" />
            ## Accepts application starting with pattern CPC with minimum
version first part
            ## of 8.0 or greater and that matched other default pattern
<AllowedApp Name="CPC 10.4.1-a" WorkStationID="americas*" />
            ## Accepts application with specific Name
            ## and should match the specified workstation pattern and other
default patterns
<AllowedApp Name="*" UserName="Administrator" />
            ## Accepts any application if login name is "Administrator"
```

```
<BlockedApp Name="CPC*" SDKVersion="10.0" />
            ## Blocks an application starting with pattern CPC with matching
SDK version version 10.0.*.*
```

Configure Server Page

Click **Tools** > **Administration** > **Configure Server** to change your server configuration settings.

Configure Server Page (General Tab)

Tools > **Administration** > **Configure Server**

The **General** tab of the **Configure Server** page allows you to edit the `Attachments` path, set server time-out options, and enable e-mail support.

Server startup log file	Default value is `..\Repository Path\server.log`. This field is read-only. The path specified when creating a new server configuration.
Attachments path	Default value is `..\Repository Path\Attachments`. This is an editable path. The folder created by the Server.
Logon sequence timeout	Default value is 60 seconds. Any logon not completed within this amount of time will fail.
Inactivity timeout __ minutes	Default value is `off`. Automatically logs off users who are inactive for the specified amount of time. Does not apply to users who have set system notification in their client **Personal Options** dialog box for a shorter period of time, because of the automatic communication between the client and the server. Also does not apply to named users, if the **Exclude named users** option (shown below) has been selected. Always set the **Inactivity timeout** to a value greater than the **Reconnect timeout**. Otherwise, if the **Reconnect timeout** and the **Inactivity timeout** are both enabled and the **Inactivity timeout** is shorter than the **Reconnect timeout**, the user is logged off before the client can reestablish the connection.
Exclude named users	Default is disabled. Allows named users to remain logged on even if they have exceeded the **Inactivity timeout** limit. Feature is available only when **Inactivity timeout** is selected and a value entered.
Enable enhanced links for all projects	Allows users to enable and disable enhanced process links for all projects on the server configuration.
Enable enhanced links for new projects	Allows users to enable and disable enhanced process links for new projects created on the server configuration.

Configure Server Page (Audits Tab)

Tools > **Administration** > **Configure Server**

The **Audits** tab of the **Configure Server** page allows you to enable audit log generation and to purge audit logs.

Enable audit generation	Default value is on. Audit log data is stored in the database for the server configuration; if data requires too much space, option can be disabled.
Purge audit entries older than __ days	Default value is `off`. Automatically removes audit entries older than a specified number of days to minimize the amount of log space required. Default is 90 days, if option is enabled. Number of days can be edited. The server configuration must be restarted to purge the audit logs.

Configure Server Page (Database Tab)

Tools > **Administration** > **Configure Server**

The **Database** tab of the **Configure Server** page allows you to view the database type for the server configuration.

Database Server Name	Disabled. Read only. Database server name can be set only when server configuration is created.
Database type	Disabled. Read only. Database type can be set only when server configuration is created.

Configure Server Page (Protocol Tab)

Tools > **Administration** > **Configure Server**

The **Protocol** tab of the **Configure Server** page allows you to set the default starting end point and encryption levels for a server configuration.

Note: Changing the endpoint does not take effect until you restart the server configuration.

TCP/IP endpoint	Default value is `49201`. Selected during creation of server configuration.
TCP/IP encryption levels	Default is set for `No encryption`. Used to set a minimum encryption level for data transferred via TCP/IP. Use **Add**, **Remove**, and **Modify** to add additional encryption levels.

Configure Server Page (Event Handlers Tab)

Tools > **Administration** > **Configure Server**

The **Event Handlers** tab of the **Configure Server** page allows you to assign default event handlers for the server/clients.

Event handler	Default value is `none`. Allows entry or selection of event handler program.
Event handler description	Default value is `on`. Allows description of selected event handler program.

Configure Server Page (Directory Service Tab)

Tools > **Administration** > **Configure Server**

The **Directory Service** tab of the **Configure Server** page allows you to enable directory service support for the server configuration.

Enable directory service support	Default value is `off`. Uses the specified Microsoft Active Directory service to validate user passwords. For a user's password to be validated, the **Validate with directory service** option must also be selected on the **New User Properties** or **User Properties** dialog boxes and the **Distinguished name** from Microsoft Active Directory service entered for the individual. Restart the server configuration to be sure that the connection to the service can be made before setting up the users. The server log contains the connection information. For example: `Connected to Active Directory Server: ldaps://host:port` where host and port are the values you enter on this tab.
Host	Host name or IP address of the Microsoft Active Directory service. It is an alphanumeric value of up to 254 characters. Instead of using a host name or IP address in the **Host** text box, you can use a domain name. When you use a domain name, StarTeam Server can contact any active copy of Active Directory anywhere in the domain so long as that copy uses the specified port. Some companies run more than one copy of Active Directory in case one goes down.

Port Default value is 389 (secure port). TLS or SSL port of the directory server. It's a numeric value.

SSL or TLS Select either one of the secure options or leave nothing selected for unsecure.

Configure Server Page (Diagnostics Tab)

Tools > **Administration** > **Configure Server**

The **Diagnostics** tab of the **Configure Server** page allows you to enable diagnostic tests for your server configuration.

Note: Typically, these options would be enabled when diagnosing a problem with a Micro Focus SupportLine representative.

Trace operations that take at least ___ milliseconds By default, this value is 0 milliseconds. Creates a .trc file that allows commands to be traced. Commands are traced if they have a duration time that equals or exceeds the specified number of milliseconds. If 0 (the default) is used, all commands will be traced.

Enable statistics monitoring. Record every ___ minutes Enables the server to track server statistics such as memory usage, currently executed commands, locking statistics, and so on.

Unexpected conditions Default is `off`. Creates a diagnostic (.dmp) file for asserts (server log entries with code # 8).

Errors Default is `off`. Creates a diagnostic (.dmp) file for exceptions (server log entries with code #4).

Diagnostic file type (default is 0) Use this option only at the direction of Micro Focus technical support. Clicking **Generate Now** creates a diagnostic file (.dmp) and places it in the server configuration's log path. It can take several minutes to generate this file, and the server does no other processing while creating this file.

Note: Other diagnostic file types should be specified and generated at the direction of Micro Focus technical personnel.

Upgrade Distribution through StarFlow

A new Client Upgrade feature allows StarTeam administrators to easily make StarTeam Cross-Platform Client releases available to their user base. To distribute:

Check the StarTeam Cross-Platform Client release build(s) into the StarTeam Server, StarFlow Extensions project, default view, root folder.

The build file names are required to follow specific patterns which are described below:

Installer Type	Build File Name
64 bit windows	`starteam-cpc-win64-{nn.mm.oo.pp}.exe`
32 bit windows	`starteam-cpc-win32-{nn.mm.oo.pp}.exe`
Linux client	`starteam-cpc-ux-{nn.mm.oo.pp}.tar.gz`
Mac Client	`starteam-cpc-mac-{nn.mm.oo.pp}.mpkg.zip`

Auto Client Update

Administrator's can distribute StarTeam Cross-Platform Client updates to each client. To run the update, choose **Download Client Update** from the **Help** menu.

Search

StarTeam allows full text search on all server-wide artifacts. Search components are installed as part of StarTeam Server.

Supported Platforms

Operating Systems	32-bit/64-bit Microsoft Windows
Databases	Microsoft SQL Server Oracle PostgreSQL

Multi-Language Search Restriction

Analyzers/stemmers used are specific to a language and there is no single analyzer supporting all languages. This means that you cannot search multiple languages within the same repository.

Search Index Configuration

To decide which server configurations to index, the search service reads the search configuration XML file installed under `Program Files\Micro Focus\<StarTeam Server Version>\Search\config\starteam-search-configs.xml`

Multiple server configurations specified in the file can be indexed simultaneously. The Server Administration Tool provides an option to enable a server configuration for indexing in one of two ways:

- While creating a new server configuration, the first page of the wizard has an option to enable search.
- Choose an existing server configuration in the Server Administration Tool via **Actions** > **Configure Search**. Check the **Enable Search** checkbox.

  Note: Search can be disabled on a configuration by un-checking **Enable Search**. Disabling a configuration would interrupt any indexing which was in progress and stop servicing clients from getting search results.

Note: The dialog box requires a StarTeam user with administrator privileges. The default user is `Administrator`, but this can be changed. The user ' s account should be validated through StarTeam Server. Users validated through directory service are not permitted to be used for starting search.

Once search is enabled on a server configuration, a new search configuration is created in the `starteam-search-configs.xml` file.

To exclude sandbox views, open the `starteam-search-configs.xml` file and update `<option name="ExcludeSandboxViews" value="1"/>`. `1` enables exclusion. **0** disables exclusion.

A template for the configuration file is available under the same folder.

Customization

The following sections cover how to customize the way text is analyzed with StarTeam search.

Tip: Changes in this section only affect new objects indexed. We recommend that you re-index the whole archive when customizing text analysis.

Multi-Configuration Search

StarTeam allows users of the StarTeam Cross-Platform Client to search across multiple server configurations that may be running on different machines. Each UI provides the user an option to select from a list of available servers from which to search. Appropriate access rights checks are performed on all servers containing matching artifacts.

The search service is a web server process (`StarteamSearchWebService15`) which handles search requests and indexing of artifacts of all the servers located in the same machine. If there are multiple search services running on different machines, a network of search services needs to be defined in `starteam-search-configs.xml`. You do this by (1) defining one as the *Master Broker* and referencing all other search services from there, and (2) pointing to the *Master Broker* from each search service. To do this:

1. Choose the server that will be the *Master Broker*.
2. Open the `<server installation folder>/Search/config/starteam-search-configs.xml` file for editing.
3. Locate `<SearchServiceNetwork>`.
4. Enter an item in the following format for each search service that the *Master Broker* should reference:

   ```
   <SearchService host="01.02.03.04" port="9090"/>
   <SearchService host="05.06.07.08" port="9090"/>
   <SearchService host="09.10.11.12" port="9090"/>
   ```
5. For each search service machine referenced by the *Master Broker*, open its `starteam-server-configs.xml` file to configure the *Master Broker*.

   ```
   <AllConfigurations>
   ..............................................................................................................
   ..............................................................................................................
   <option name="SearchServiceBrokerHost" value="##.###.##.###"/>
   <option name="SearchServiceBrokerPort" value="9090"/>
   </AllConfigurations>
   ```

   Note: If there is only one search service process handling multiple servers on the same machine, the network is not required to be defined. All servers handled by the search service are automatically available for the multi-configuration search.

Filtering Options

The following filtering options are available for non-English words:

Elisions	Use `FilterElisions` to filter elisions from words during indexing. For example, `l'avion` (*the plane*) will be tokenized as `avion` (*plane*). So a search for `l'avion` and `avion` would both work and return the same result.
Accent Marks	Use `FilterAccents` to remove accent marks from Latin characters during indexing. For example, à will be replaced by a and can be searched without accents.

These options need to be configured in `starteam-search-configs.xml.` Use `1` to enable the option. For example:

```
<option name="FilterElisions" value="0"/>
<option name="FilterAccents" value="1"/>
```

Stemming

Stemming is the process of reducing inflected (or sometimes derived) words to their stem, base, or root form.

StarTeam search uses a Porter stemmer for default word analysis and indexing. The current JVM language variable is used to decide which stemmer to apply as follows:

Language "en"	English Stemmer
Language "fr"	French Stemmer
Language "pt"	Portuguese Stemmer
Language "de"	German Stemmer

Chinese and Japanese locales

StarTeam search uses Lucene's `CJKAnalyzer` by default. Analyzers are configurable by editing `starteam-search-configs.xml`. For example, to use Lucene's *SmartChineseAnalyzer*, which is an analyzer for simplified Chinese or mixed Chinese-English text, make the following change:

```
<Analyzers>
<Analyzer name="zh"
value="org.apache.lucene.analysis.cn.smart.SmartChineseAnalyzer"/>
</Analyzers>
```

The string `org.apache.lucene.analysis.cn.smart.SmartChineseAnalyzer` is a class name which is part of the Lucene library.

Stop Words

Stop words are words which are filtered out prior to, or after, processing of natural language data (text).

Each supported language has a corresponding stop words file, `<language>_stopwords.txt`, in the search `config` folder:

- `en_stopwords.txt`
- `fr_stopwords.txt`
- `pt_stopwords.txt`
- `de_stopwords.txt`

These files contain the actual stop words used for each language during indexing. Customize these files by adding or removing words that should not be indexed. Adhere to the file format, which is:

- One word per line.
- Lowercase words.
- Words need to be in UTF8.

Tip: Stop words will be added to the search log.

Troubleshooting

There are several reasons for the search indexing service to fail to work correctly:

- Search indexing logs can be found under `<Server installation folder>\WebServer\logs\starteam-search.log`. If search indexing fails for any reason, the logs should contain error messages detailing the possible issues. To turn on debug mode logging, edit the log levels in `<Server installation folder>\Search\config\starteam-search-configs.xml` to `DEBUG` instead of `INFO`. `<level value="INFO"/>`
- The search indexing process could fail due to database connectivity issues. There is a known issue while working on StarTeam Server using a SQL Server express edition database. When using SQL Configuration manager, make sure that the TCP/IP protocol has been enabled and SQL Browser service is running. If SQL Server express was installed along with StarTeam Server, these are enabled by default. If it was installed separately , these settings are disabled by default.
- Search indexes can consume a lot of disk space depending on the size of the configuration. The search indexes are usually located under `<repository Folder>/searchIndex` directory. This can be

configured in the search configuration file. `<option name="SearchIndexPath" value="C:\repository\Test\searchIndex"/>`. Doing a trial run on a test server with a replicated database and vault should give an idea of the required disk space for index files.

- Search indexing services perform better when MPX is enabled on the StarTeam Server. The service is notified of any changes on the StarTeam Server using MPX events during delta-mode indexing. Without MPX, the search service would work but would query the database more often leading to potential performance issues. By default, the indexing process looks out for changes every 30 minutes. This is defined in `starteam-search-configs.xml` and can be edited if required to a larger value, for example thirty minutes, which means new changes will be available to search only after thirty minutes. `<option name="ExtractDeltaInterval" value="3"/>`
- Restarting the tomcat service `StarTeamSearchWebService<version #>` could sometimes resolve issues related to memory issues. It is also recommended that if StarTeam Server/Message Broker is restarted, the tomcat service `StarTeamSearchWebService<version #>` should also be restarted.

Custom Components

You can create your own custom component to use in StarTeam using the **Custom Component Builder** or manually with an `XML` file.

Custom Component Builder

StarTeam Server's **Custom Component Builder** is used to create custom components analogous with StarTeam's own internal components, such as the *File*, *Change Request*, *Task*, or *Topic*. The **Custom Component Builder** in the StarTeam Server walks you through creating the component, its properties and values and finally, deploying your component in a single click.

When you are finished creating your custom component, use the Layout Designer to lay out the properties that you defined in the `XML` file or via the **Customize** menu in the StarTeam Cross-Platform Client. The layout form you create in the Layout Designer is stored as a file in the `Starflow Extensions` project, and read by the StarTeam Cross-Platform Client when it shows a custom property editor.

Important: To see a new custom component in the StarTeam Cross-Platform Client, go to **Personal Tools** > **Workspace** > **Advanced** and select **Component Order**. Add it to the list of visible components and then re-open the projects.

Creating a Custom Component

1. Open the server configuration in which you would like to create a custom component.
2. In the **Custom Component Builder** tab, click **New Definition**.

   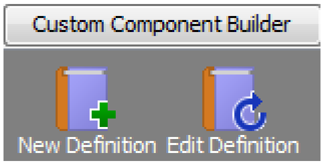

   The **Add Component** dialog box opens.
3. Modify the component definition:

Component Name	Enter the component name. This is the internal name of the component.
Display Name	Enter the display name. This is the name that is displayed on the custom component control label.
Type	Enter the type of component.

Tree Component The *Tree Component* type can have parent child relationships. An example would be a StarTeam *Task* type. When shared into views, this type cannot branch. You need to either pin the new share such that it's frozen in a state in time, or make it float.

Branchable Component A *Branchable Component* is like a StarTeam *Change Request*. It is a type that does not have parent-child relationships. When shared into new views, the item can branch when changes are made. Therefore, the original item and the new item can have different sets of edits.

4. Click **OK**.

Creating Properties for Custom Components

1. Open the custom component to which you want to add properties.
2. Click **New**. The **Add Property** dialog box opens.
3. Enter a **Field Name** for the property. This is the internal name.
4. Enter a **Display Name** for the property. This is the name on the property label.
5. Select the data **Type** for the property.

Data Type	Description
Boolean	Boolean property fields are used to define true/false property values. You can modify the possible values. **Note:** You can modify the `True/False` values in a boolean data type using the **Advanced** > **Customize...** command in the StarTeam Cross-Platform Client.
Content	Content property type fields can be used to store rich text values, such as HTML MIME types.
Date	Supports a date field.
Date/Time	Supports a date and time field.
Enumerated	Allows you to create a custom enumerated field. You can modify all possible values and set a default value.
Group	Use this type to store the user group values defined on the server. You can only select a single group.
Group List	Use this multi-select data type to store the more than one user group value defined on the server.
Integer	Support a numeric *integer* field
Map	Use this type to store structured values on items where the value types are not required to be the same or exist on every item instance
Real	Supports a numeric *real* field.
Text	Allows you to enter text.
Time Span	Allows you to enter a time span in days, hours, minutes, and seconds.
User	User data types can be used to store the user values defined on the server. You can only select one.
User Groups	Same as the User data type but you can choose multiple items.

6. For `Text` or `Map` property types, enter the **Length**.
7. For `Enumerated` data types, add values using the **Possible Values** list.

   1. Click **Add**. The **Add Value** dialog box opens.
   2. Enter a unique numeric key for the value in the **Code** field. This value must be an integer between 100 and 1,000,000.
   3. Enter a value in the **Name** field.
   4. To make a value a child of another value, click the **Make child of** option and select the parent value in the list.
   5. Click **OK**.
   6. Repeat these steps for as many values as you want to add.

8. For the `Enumerated` data type, select the **Supports multi-select** option to allow multi value selections.
9. Enter the **Default value** if it applies:

   Boolean True or False or whatever values you may have changed them to.

   Enumerated Select one of the values that you created.

   Integer Enter a valid `Integer` number.

   Real Enter a valid `Real` number.

   Time span Check the **Time Span** option. Enter the number of days and select the hours, minutes, and seconds.

Editing a Custom Component

1. Open a server configuration.
2. Click **Tools** > **Custom Component Builder** > **Edit Definition**. The **Open** dialog displays.
3. The default directory for saved custom components is `<server installation path>\customcomponents\staging`. Select the custom component you want to modify and click Open. The custom component is opened into the editor.

Cloning a Component Definition

1. Open a server configuration.
2. Click **Tools** > **Custom Component Builder** > **Clone Component Definition**. The **Clone Component definition** dialog displays.
3. The custom component is opened into the editor, loaded with the properties of the chosen component type. Properties can be added/deleted/updated based on the requirement for the new component.

Exporting a Component Definition

1. Open a server configuration.
2. Click **Tools** > **Custom Component Builder** > **Export Component Definition**. The **Export Component Definition** dialog displays.
3. Choose the type whose XML definition is required to be exported and the location on disk to save to.

Deploying a Custom Component to Clients

After you create your custom component, you need to deploy the component to your clients.

Important: Before you can deploy your custom component, all clients must be locked out of the server.

1. Open the custom component.
2. Click **Deploy**.
3. You will receive an option about whether you have locked all clients out of the server. Once they are locked out, click **Create custom component now**.

4. Click **OK**.

Manually Creating Custom Components

Create a server configuration, but do not start the server. If adding custom components to an existing server configuration, stop the server.

You can create custom components using the XML file located in the server directory you created or from the SDK. This procedure below discusses the XML file method. For instructions on how to create custom components with the SDK, refer to the SDK documentation.

1. Create a server configuration or stop the existing server before adding custom components. During the server configuration creation, you designate a repository path. When the configuration is created, files are placed within that directory. The custom component XML file is located in the `CustomComponents` folder contained within the directory.
2. Create the XML file in an editor such as Notepad.
3. Save the XML file in the `CustomComponents` folder in your server configuration.
4. Find the `<component name="">` tag and type the component name between the quotes. The `Component Name` can be no more than 20 characters and consist of ASCII A-Z and a-z and numeric 0-9 characters. Additionally, the `Component Name` must be unique for the server configuration.
5. Find the `<ClassDisplayName default =" ">` tag and type the component name between the quotes. The `Class Display Name` can be no more than 20 characters and consist of ASCII A-Z and a-z and numeric 0-9 characters. Additionally, it must be unique for the server configuration.
6. Type a default language for the `Class Display Name`. Additional language translations are options and will appear as `<translation language = "fr-FR">Exigence de hardware</translation>`, for example.
7. Find the `<ComponentType></ComponentType>` tags and type `0` or `1` between the tags to define the component type.

    - `0` - Tree component includes the following system properties:

    ```
    "ID"
    "CreatedTime"
    "CreatedUserID"
    "DeletedTime"
    "DeletedUserID"
    "ModifiedTime"
    "ModifiedUserID"
    "EndModifiedTime"
    "RevisionNumber"
    "RevisionFlags"
    "ShortComment"
    "CommentID"
    "ParentID" -  system property, type is eLong
    ```

    - `1` - Branchable component includes the following system properties:

    ```
    "ID"
    "CreatedTime"
    "CreatedUserID"
    "DeletedTime"
    "DeletedUserID"
    "ModifiedTime"
    "ModifiedUserID"
    "EndModifiedTime"
    "RevisionNumber"
    "RevisionFlags"
    "ShortComment"
    "CommentID"
    "RootObjectID"
    "ParentObjectID"
    ```

```
ParentRevision"
"PathRevision"
"ViewID"
"DotNotation"
```

The following properties are created for all components:

```
Attachment Service Properties
 "Attachment count" -  system property, type is eLong
 "AttachmentIDs" -  system property, type is IDArray
 "Attachment names" - server calculated property, type is text
Notification Service
 "Notification count" - system property, type is Long
 "NotificationIDs" - system property, type is IDArray
 "ReadStatusUserList" - server calculated property, type is IDArray
Bookmark Service
 "FlagUserList"  - server calculated property, type is IDArray
Component Object Identifier
 ComponentName + "ID" - system property , type is eLong
```

8. In between the `<Properties></Properties>` tags, define each property for the component. Create a name and type for each properties as well as values, order length, and such, depending on the property type. For example:

```
 <Properties>
- <Property name="Name">
  <DisplayName default="Name"/>
  <Type>8</Type>
  <Length>250</Length>
  </Property>
- <Property name="AutoType">
  <DisplayName default="Automobile Type"/>
  <Type>2</Type>
  <DefaultValue>"100"</DefaultValue>
  <Flags multiselect="false"/>
- <Enum code="100">
  <EnumName default="4 door"/>
  <Indent>0</Indent>
  <Order>1</Order>
  <Flags selectable="true"/>
  </Enum>
- <Enum code="101">
  <EnumName default="2 door"/>
  <Indent>0</Indent>
  <Order>2</Order>
  <Flags selectable="true"/>
  </Enum>
  </Property>
```

Tip: The property types are defined at the top of the XML file.

9. Save your changes.
10. Start the server. If there are any errors while parsing the XML file or while creating the custom components, locate the errors in the server log. The server starts without creating the new component.

Import/Export Manager

The **Import/Export Manager** provides performance and scalability improvements by allowing you to selectively move projects from one StarTeam Server to another StarTeam Server. There are many reasons to have multiple servers. In this case, you may want to archive old projects from your active server or simply separate your different business units. Having the ability to move projects around will provide limitless opportunities for you to keep your active servers lean and full of only the most important projects.

Refer to the *Business Rules* topic for important information on how your data will be migrated.

Limitations

- This version of **Import/Export Manager** supports Microsoft SQL Server, Oracle, and PostgreSQL databases. You cannot copy data between Oracle and Microsoft SQL Server or PostgreSQL but you can copy between PostgreSQL and Microsoft SQL Server.
- This version of **Import/Export Manager** does not offer a command-line option.
- This version of **Import/Export Manager** does not offer the option of re-starting. If the import operation fails due to unexpected reasons, the target database may be left in a bad state. It is very important to take a backup of the target database and the vault before importing. In case of failure, the **Import/Export Manager** will write the error messages and the stack-trace detailing the reason for the failure. After resolving the issue, the user should not use the copy of the database involved in the failed operation but should use a new target database backup from pre-import time.

Import/Export Manager Business Rules

The following rules are used when you run the **Import/Export Manager**:

Artifact Number Mapping

All the artifact numbers (change request number, requirement number, etc) are remapped to a new number in the target. If the change request has been shared across multiple views or sand boxes, the change request number will remain intact. The older change request number will be stored in a custom field for cross referencing purposes.

Broken shares and links

- The export process evaluates all the shares and links that cross the selected project boundaries. Shares and links that cross project boundaries, but are within the list of exported projects, are not broken by **Import/Export Manager**.
- For those shares whose parent or root is in a different project which is not being exported, **Import/Export Manager** will make this item its own root. It will reset the branching behavior to model a root item.
- Links with either the source endpoint or the target endpoint outside the project scope will be ignored and will not be exported.

Cross project shares and Links

Cross project shares are shares whose parents reside outside the selected projects boundaries. These shares and links are broken by **Import/Export Manager** and reset to be root items. The branching behavior of these items is reset to model a root item. If the source item is pinned, the pinned behavior is carried over to the target.

Custom Components

Custom component names are matched by their internal name during import. If there is a match found on the target server for a given component name, the source and target components are considered to be the same. If a match is not found, a new custom component will be created on the target server and all its associated data, properties, and enumerations will be copied over. It is the responsibility of the user to review their source and target components and see if there are matching custom components which are not related but happen to have the same internal name. In this case, the user will have to rename either the source/target component before import.

Custom Properties

For components on the source and target server that match, the import operation iterates through all the custom properties available and checks to see if a custom property with the same internal name for the component exists on the target server. If it is not found, this property will be created on the target server.

Microsoft SQL Server has limitations with the row length and importing a large number of custom properties could potentially cross the row length limit. One way to avoid this issue is to exclude disabled properties from being carried over. The **Exclude disabled properties** check box is provided during export. This will ensure that all the catalog field/enumeration entries and the values for the disabled properties on the source server are not imported.

Before custom properties are added to the target server, **Import/Export Manager** will do a row length check to determine if the target database table allows new columns to be added. **Import/Export Manager** will not continue if this check fails.

Please also consider:

Property with the same internal name but different enumerations	If two enumerated (single/multi-select) custom fields have the same internal name and the same type, the target server configuration must have the enumerations found on the source as well as those found on the target. For example, if the source field has `100 red`, `200 blue`, and `300 white` and the target field has `100 blue` and `200 orange`, **Import/Export Manager** will add `red` and `white` to the list of enumerations for the target field. **Import/Export Manager** will see that the `200 blue` is equivalent to the `100 blue` and ignore `orange`, but it needs to find `red` and `white`, regardless of their code numbers. After the import, this field will have the following enumerations on target server configuration for this example field: `100 blue`, `200 orange`, `<some number> red`, `<some number> white`. **Import/Export Manager** preserves the values (visible to users) and does not worry about the code numbers. However, you may want to set up the extra enumerations needed on the target server manually before importing to control those codes. Disabled enumerations are still copied for enabled custom fields because some items may have those values. The enumerations will be disabled on the target unless the target field already has an equivalent enabled enumeration. The default values for an enumerated property are not merged.
Custom property being added to the target already exists with the same data type but different length	The target custom property's length will be increased to prevent data truncation issues for the imported data.
Custom property being added to the target already exists with a different data type	**Import/Export Manager** will create a new custom property with a system generated name.

Deleting projects from source server	**Import/Export Manager** does not delete the **Projects selected for export** from the source server. It is up to the user to decide based on their requirements.
Filters, Queries, and Audits	Filters, queries, and audits will not be imported by default. This can be changed in the user interface.
Hives	In addition to the database metadata, the **Import/Export Manager** also copies over vault contents to the target server. If the target server has multiple hives, the contents are added to all the hives in a round-robin fashion, similar to the way StarTeam Server processes content. The **Import/Export Manager** does an up front estimation of the

	total size of the vault content being exported from the source and checks to see if the target hive has the capacity to handle importing the contents. If there is not sufficient storage space, the **Import/Export Manager** will throw an error message and abort the operation before importing any data. Estimation of space on the target hives depends on various parameters like storage threshold limit of the hive, disk partition capacity, and also the **Allow new Archives** option. It is recommended that you ensure that there is ample disk space on the target machine and that the hives are configured appropriately before starting import. File contents, attachments, and contents from content properties are copied over to the target as is, without any processing. Change package attachments, however, need to be processed before being imported to the target server. Change package attachments have ids embedded in the content. They need to be re-mapped to the corresponding id in the target. The imported contents will therefore end up having a different md5 value compared to the one in source. This is essential so that the imported change packages can be opened and used.
Licenses	**Import/Export Manager** does not copy over the source server's licensing information. It is important for the user to ensure that the target server is licensed appropriately. For example, if the source server has enterprise advantage license, and the user is importing requirements data from this server to the target server, it is required that the target server has an enterprise advantage license to view the imported requirements.
Source and Target Servers	Source and Target StarTeam Servers must be running the same build number.
Tip Items	If a tip item is moved across project boundaries, the historical revisions of the item will be available in the target to support rolled back configurations. The tip item will not be available in the target since it does not belong to the imported project.

Best Practices

- Always backup the target database before running an import. The import process significantly changes the target database and if the import fails, the target database may be unusable.
- It is strongly recommended to do multiple test runs against a copy of the database and the vault before it is deployed in production.
- **Import/Export Manager** copies significant amount of data to the target database. It is important to ensure that target server machine and database machine have available free space to support this import.

Exporting

This section contains topics showing you how to export projects.

Exporting a Project

Use the following steps to export one or more projects from a Local server configuration into a file on your computer. After that, you can import the project into another StarTeam Server.

Note: If you are moving projects within your Local server configuration, you do not need to export them first. You can go directly to the import process.

1. Open the **Server Administration** tool.
2. Select a server configuration.
3. Stop the server.
4. Click **Tools** > **Import Export Manager** > **Export**. The **Export** window opens.

5. Select one or more projects to export in the **Select projects to export** list.
6. Click the **Select all projects** checkbox to select all of the projects in the **Select projects to export** list.
7. Click **Ignore filters and queries** to exclude filters and queries from the export file.
8. Under the **Choose a location for the export file** field, click **Browse** to select a location for the export file.

   The export file will be named the same name as the configuration and contain a `.stbk` extension.
9. Click **Export** to begin the export process. The **Import/Export Manager** begins the export process and displays progress in a window. The message, `The export operation completed successfully`, will be displayed when the export is done.

Export Window

Use the **Export Window** to export one or more projects from a Local server configuration into a file on your computer. To access the window, click **Tools** > **Import Export Manager** > **Export**.

Note: You must stop the server configuration before you can export a project.

Select projects to export	Select one or more projects to export.
Select all projects	Click to select all of the projects in the **Select projects to export** list.
Ignore filters and queries	Click to exclude filters and queries from the export file.
Choose a location for the export file	Click **Browse** to select a location for the export file.
Export	Click to begin the export process.

Importing

This section contains topics showing you how to import projects from a backup file.

Importing an Export File

Important: Make sure you read the *Business Rules* topic to understand how your data will be imported.

Use the following steps to import an export file into a local server configuration.

1. Open the **Server Administration** tool.
2. Select a server configuration.
3. Stop the server.
4. Click **Tools** > **Import Export Manager** > **Import from backup file**.

   Note: You can also import directly from a local server configuration. Click **Tools** > **Import Export Manager** > **Import from server configuration**. Select the server configuration and ignore the next step.
5. Next to the **Choose file on disk to import from** field, click **Browse** to select the backup file to use.
6. Check the **Merge Users and Groups** checkbox if you want to merge users and groups.
7. Check the **Rename projects upon conflict** checkbox if you want to have projects renamed upon conflict. If there is a conflict, you will be given a chance to rename the imported projects.
8. Check the **Ignore filters and queries** checkbox if you want to ignore filters or queries during the import.
9. Check the **Ignore Audits** checkbox to ignore audits during the import.
10. Check the **Ignore disabled properties in source server** checkbox if you want to ignore properties that have been disabled on the source server.

11. Check the **Intervene upon merge conflicts** checkbox if you want to show all conflicts in a window during import. From there, you can choose to either use the source or target values.
12. Click **Import**. If you selected **Intervene upon merge conflicts** and there are conflicts, the **Handle conflicts during import** dialog box displays. Select how you want conflicts to be handled. A message, `The Import operation completed successfully`, displays when the import is complete.

Import Window

Use the **Import** window to import a backup file into a Local server configuration. To access the window, click **Tools** > **Import Export Manager** > **Import from backup file**.

Note: You must stop the server configuration before you can import a project.

<table>
<tr><td><b>Choose file on disk to import from</b></td><td>Click <b>Browse</b> to select the backup file to use.</td></tr>
<tr><td><b>Handle Merge conflicts</b></td><td>Contains different options allowing you to control how conflicts are handled.
<table>
<tr><td><b>Merge Users and Groups</b></td><td>Select to merge users and groups.</td></tr>
<tr><td><b>Rename projects upon conflict</b></td><td>Select to have projects renamed upon conflict. If there is a conflict, you will be given a chance to rename the imported projects.</td></tr>
<tr><td><b>Ignore filters and queries</b></td><td>Select to ignore filters or queries during the import.</td></tr>
<tr><td><b>Ignore Audits</b></td><td>Select to ignore audits during the import.</td></tr>
<tr><td><b>Ignore disabled properties in source server</b></td><td>Select to ignore properties that have been disabled in the source server.</td></tr>
<tr><td><b>Intervene upon merge conflicts</b></td><td>Select this option to show all conflicts in a window during import. From there, you can choose to either use the source or target values.</td></tr>
</table>
</td></tr>
</table>

Handle Conflicts During Import Dialog Box

The **Handle conflicts during import** dialog box displays when:

- There are assets in the source and target that appear to be the same.
- You checked the **Intervene upon merge conflicts** checkbox on the **Import** window.

Review the conflicts in the list on the dialog box and select how you want the conflicts to be handled.

For users and groups, you will see:

Use Source	Selecting this option will override the values in the server configuration with the values in the backup file.
Use Target	Selecting this option will ignore the conflicting values in the backup file and use the values from the server configuration.
Do the same for all remaining conflicts	Selecting this saves whichever option you selected, **Use Source** or **Use Target**, and applies it to all merge conflicts.

For project name conflicts, you will see:

Default action is to rename project using suffix (`ConfigName.stbk`).	Select to let the application rename the project for you.
Rename project using alternate suffix instead.	Type in the name of the project in the field.

Maintenance Task Scheduler

The **Maintenance Task Scheduler** allows you to schedule scripts to run against local server configurations. The **Maintenance Task Scheduler** provides you with the ability to automate performance improvements on your database on a schedule that you choose. You can run scripts that will update query optimization statistics and online index rebuilding.

Here are some items to consider when using the **Maintenance Task Scheduler**:

- You cannot modify scripts or add to the list of scripts.
- If a script or folder containing the scripts is moved and there is a schedule for the script set, StarTeam Server will still attempt to run the script but it will log an error.
- You can provide a **Recurring** schedule for your scripts. This gives you the flexibility to run scripts at different times. You can also run scripts ad hoc.
- Script execution is logged in a file called `execScript.log`, located at the same location as the server log.
- It is not possible to schedule scripts remotely through the **Server Administration** tool. This only works for local configurations.
- You can define a new script schedule while the server configuration s running. The modification will be picked up after about 30 seconds.

Assigning a Schedule to a Script

1. On the **Server Administration** tool, select a local server configuration.
2. Click **Actions** > **Configure Maintenance Tasks**. The **Maintenance Task Scheduler** dialog box opens.
3. Select a script from the **Scripts** list.
4. From the **Frequency** list, select **Recurring**.

   Note: To remove a schedule from a script, select **Not Scheduled**.
5. Select the hours and minutes from the **Time** controls.
6. Select on which days you would like to run the script.
7. Click **Save**. Your schedule is saved and will execute at the specified date(s) and time(s).

Scripts

The following scripts are available for the **Maintenance Task Scheduler**.

Note: The location of the script log is a subfolder named `SchedulerLogs` under the server log path.

Microsoft SQL Server

Script Name	Description
`starteam_sqlserver_create_sp_update_stats.sql`	Updates query optimization statistics. Updating statistics ensures that queries compile with up-to-date statistics. Can be run online or offline.

Script Name	Description
`starteam_sqlserver_create_index_maintenance_script.sql`	Rebuilds indexes. Scheduling this script requires an edition of Microsoft SQL Server that supports online index rebuilding. If not supported but scheduled, there will be an exception in the log when running it. Clustered indexes are rebuilt using the online option. To rebuild clustered indexes, the offline version should be executed whenever significant amounts of data is in the server. It is always possible to run it offline by clicking **Execute**.

PostgreSQL

Script Name	Description
`starteam_postgres_create_compute_stats.sql`	Updates query optimization statistics. Updating statistics ensures that queries compile with up-to-date statistics. Can be run online or offline.
`starteam_postgres_create_index_maintenance_script.sql`	Rebuilds indexes. Can be run online or offline.

Oracle

Script Name	Description
`starteam_oracle_create_compute_stats.sql`	Updates query optimization statistics. Updating statistics ensures that queries compile with up-to-date statistics. Can be run online or offline.
`starteam_oracle_create_index_maintenance_script.sql`	Rebuilds indexes. Scheduling this script requires an edition of Oracle database that supports online index rebuilding. (Currently Oracle Enterprise Edition only.) If not supported but scheduled, there will be an exception in the log when running it. It is always possible to run it offline by clicking **Execute**.

Maintenance Task Scheduler Dialog Box

Use the **Maintenance Task Scheduler** dialog box to schedule scripts to run against your local server configurations.

Scripts Provides a list of scripts that you can schedule to run.

Frequency Select how often you'd like the script to run.

Not Scheduled Default value for scripts without a schedule. Selecting this value will ensure that the script does not run. Other timing controls will be disabled.

Recurring This will cause the server to execute the script at the first available opportunity on or after the selected day of week and time of day. Selecting **Recurring** allows any number of days of the week to be selected. After execution, the schedule will remain in the **Recurring** state.

Time	Select the hour and minutes from the lists.
Day of Week Check boxes	Depending on the frequency you've selected, select the day(s) of the week that the script should execute.

In addition to the controls to set the schedule, you can use the following options:

Save	Saves the schedule.
Execute	Runs the script immediately. The **Execute** button only works for offline execution and can only be done when the server configuration is not running. Otherwise the option will be disabled in the dialog box.

Managing Log and Initialization Files

Initialization Files

Initialization files have different locations on different Microsoft Windows platforms. On Microsoft Windows 2008+, the pathPrefix is `C:\Documents and Settings\`.

ClientLicenses.st	Located at `%PROGRAMDATA%\Borland\StarTeam\ClientLicenses.st`. This file is installed by StarTeam Runtime and the application clients. If the `ClientLicenses.st` file is missing, you are asked to register the product.
ConnectionManager.ini	Used for starting the application and is located at `%PROGRAMDATA%\Borland\StarTeam\ConnectionManager.ini`.
starteam-client-options.xml	Located at `%APPDATA%\Borland\StarTeam\starteam-client-options.xml`.
starteam-server-configs.xml	Used for server session information, is located at `Server Installation Folder\starteam-server-configs.xml`.
starteam-servers.xml	lists the server configurations for which you have created (or will create) server descriptions. These descriptions are used while opening or creating projects. This file is located at `%PROGRAMDATA%\Borland\StarTeam\ServerList` and is installed by StarTeam Runtime and the application clients.

ConnectionManager.ini

`ConnectionManager.ini` contains information that the client must be able to locate in order to run. It is created at the time that the application is installed. In actual `ConnectionManager.ini` files, the x's are replaced by hexadecimal numbers. If the `ConnectionManager.ini` file is missing or corrupted, the application asks if you want it recreated. Reinstallation can also recreate the missing `ConnectionManager.ini` file. Example:

```
[ConnectionManager]
WorkstationID=xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx
```

starteam-client-options.xml

This file is installed by StarTeam Runtime and the application clients. The `starteam-client-options.xml` file lists personal option settings and any alternate working folders you have set with the application (accessed by selecting **Tools** > **Personal Options**). You may want to back up this file or put it under version control. If the `starteam-client-options.xml` file is missing, the application automatically recreates it. However, the recreated file contains only the default settings for the personal options and no alternate working folder information. If the `starteam-client-options.xml` file is corrupted, you can delete it, but you may be able to edit it. You can tell that the `starteam-client-options.xml` file is missing or corrupted when your personal options are no longer correct, changes you made to personal options disappear when you restart the application, files do not change even though you have checked them out (because they have been copied to the wrong working folders), or the application reports that old files are missing and does not see new files, because it is looking for them in the wrong place.

The options that are check boxes in the dialog are set equal to 1 for selected or 0 for cleared. Intervals are set to a number of minutes or seconds depending on the option. Paths are in text. No quotation marks are used with the text.

For example, the *Project Component* information provides the paths to alternate working folders for projects accessed from your workstation. The entry for this component in the `starteam-client-options.xml` file includes the following parts.

- The phrase `Project Component`.
- `ViewWorkingFolderOverrides` (if the alternate working folder is for an entire view) or WorkingFolderOverrides (if the alternate working folder is for an individual folder).
- A hex identification of the project view and folder.
- The alternate working folder's path.

starteam-server-configs.xml

The `starteam-server-configs.xml` file contains session options for one or more server configurations. StarTeam Server session options specify the core information that the Server requires to start a server configuration. One `starteam-server-configs.xml` file exists per computer and is located in the same folder as the StarTeam Server application. On Microsoft Windows, this file is usually located in `C:\Program Files\Micro Focus\StarTeam Server <version #>`.

The session option information for each server configuration begins with the name of the configuration in brackets and is followed by a set of options and their settings. The StarTeam Server creates and maintains this file, which is created when the first server configuration is created. The file is updated whenever a server configuration is created, modified, deleted, started, or stopped. Do not edit this file directly. We recommend that you back up the `starteam-server-configs.xml` file or put it under version control.

Information in this file includes:

Server-Managed Settings

⚠ **Important:** Options in this section should only be modified using the Admin tool, if at all.

Configuration Creation Settings	**ServerGuid**	Value supplied by the server. This option is set by the server. Do not edit it.
	ComputerName	Identifies the computer on which Microsoft Windows resides. This option is set by the server. Do not edit it.
	CreatedByBuild	Indicates that a server configuration cannot have any Native-I archive files.
	DBCreated	Indicates whether the database tables used by the application are already created. Do not edit this option.
	UserName	The domain user name for the user who created the server configuration. This option is set by the server. Do not edit it.

	Initialized	Indicates whether this server configuration was initialized. This option is maintained by the server. Do not edit it.
Application Settings	**RepositoryPath**	The complete path to the repository folders. This information is set using the `-r` option with `starteamserver` command. The repository path can only be specified when creating a new server configuration. This information cannot be modified for existing server configurations.
	ListenIP	Binds a server configuration to a specific TCP/IP (sockets) address. For example, if the StarTeam Server has more than one IP address (more than one NIC card), you can configure it to listen to one specific port. When this option is set to `0` (default) the server configuration listens to all IP addresses on the specified port. The port is specified on the **Protocol** tab of the **Configure Server** tool in the Server Administration Tool.
	LogPath	Specifies the location of the server log file.
	EndPoint	Server endpoint.
	MessageBrokerType	This information is set using the `-mb` option with `starteamserver` command. The message broker type can be specified when creating a new server configuration. Message broker options available include: • `0 = None` • `1 = Legacy StarteamMPX` • `2 = ActiveMQ MPX`
Database Settings	**DBPort**	The port used by the database server.
	DBServerName	Name of the database server.
	DBServiceName	(Oracle only). The name of the service on the Oracle server. Either `-dbservicename` or `-dbsid` is required when creating an Oracle configuration from the Microsoft Windows command line.
	DBSID	(Oracle only). This is the SID on the Oracle server. Either the `-dbsid` or the `-dbservername` parameter is required when creating an Oracle configuration from the Microsoft Windows command line.
	DBType	This information is set using the `-t` option with `starteamserver` command. The database type can only be specified when creating a new server configuration. This information cannot be modified for existing server configurations. You can review the database type using the `-view` option from the command line or in the application on the **Database** tab of the dialog. • 2 = Microsoft SQL Server or SSE. • 3 = Oracle. • 7= PostgreSQL.

DBUserName Name by which the application accesses the server configuration database. This information is set using the `-u` option with the `starteamserver` command.

DBPassword Password for `DBUserName` by which the application accesses the server configuration database. This information is set using the `-p` option with the `starteamserver` command.

Process Execution Settings

PID The *Process ID* for the instance of the server configuration that is currently running. Otherwise, this option is set to `0`. This option is maintained by the StarTeam Server. Do not edit it. When this option is missing, `starteamserver` creates it.

Sample Indicates the sample server configuration. For internal use only. Do not edit this option.

ServiceMode For use on Windows NT systems only. Use `1` to run the server configuration as an NT Service. Use `0` to run the server configuration as an application.

Status Indicates whether the server configuration is `Ready`, `Starting`, `Running`, or `Stopping`. This option is maintained by the server. Do not edit it. When this option is missing, `starteamserver` creates it.

Performance Tuning

DBMinimumConnections Number of connections created at server startup. Default: `25`. Valid range: `10-100`.

DBMaximumConnections Maximum number of database connections that the server is allowed to have, connections added as needed. If max limit is reached, you receive a message: `"Failed to acquire a database connection. Please contact an administrator to increase number of database connections."`

Default: `25`. Valid range: `10-65,535`.

MinCommandThreads Not used.

MaxCommandThreads Specifies the maximum number of command threads that the server configuration can create. A setting of `0` for this option causes the server configuration to use the default value. Default value: `25`. Valid range: `16-120`.

ItemCachePriority_File Defines the relative preference of file data in the server data cache. Each component will occupy no more than it's percent share, which is calculated as shown in the following example: `FileShare% = (100 * ItemCachePriority_File ) / (ItemCachePriority_File + ItemCachePriority_Change + ItemCachePriority_Requirement + ItemCachePriority_Task + ItemCachePriority_Topic + ItemCachePriority_Folder )`.

Default value: `100`. Valid range: `0-100`.

ItemCachePriority_Folder Defines the relative preference of component data in the server data cache. Default value: `0`. Valid range: `0-100`.

ItemCachePriority_Changes Defines the relative preference of component data in the server data cache. Default value: `0`. Valid range: `0-100`.

ItemCachePriority_Topic Defines the relative preference of component data in the server data cache. Default value: `0`. Valid range: `0-100`.

ItemCachePriority_Task Defines the relative preference of component data in the server data cache. Default value: `0`. Valid range: `0-100`.

ItemCacheOnStartup_%s Cache on startup setting by class name (File, Changes, Topic, etc). Default value: 64-bit = `1`, 32-bit =0. Valid values: `0` (disabled) or `1` (enabled).

ItemCachePriority_Requirement Defines the relative preference of component data in the server data cache. Default value: `100`. Valid range: `0-100`.

ItemCacheLogRecovery Enables logging of item cache cleanup. Default value: `0`. Valid values: `0` (disabled) and `1` (enabled).

ItemCacheMemoryLimit Defines the maximum amount of memory that can be used for caching item data, in MB. The default value is `-1`, which means no limit is set and the server will use up to the maximum amount of memory available to the server process for caching the data. Example: Adding `<option name="ItemCacheMemoryLimit" value="100"/>` to the configuration file will set the cache limit to 100 MB. Valid range: `1000-7000`.

ItemCacheProcessMemoryLimit Maximum allowed process memory size, used as cap for Item Data Cache size. Default values:

Win32	1.5GB
Win32/PAE	2.3 GB
Wow64	2.8.GB
Win64	8.0 GB

SocketTimeoutMinutes Socket connection timeout. Default value: `0`. Valid range: `0-10`.

ObjectQueryChunkSize Data Cache queries tuning. Default value: `10000`.

MaxWorkstationConnections Limits the maximum number of connections per workstation. Default value: `25`. Valid range: `25-100`.

LinkValueCache Used to enable/disable LinkValue caching. Default value: `1`. Valid values: `0` (disabled) or `1` (enabled).

MinimumClientCommandAPI Minimum command API accepted by the server. Default value: `1.25`. Valid range: `1.25-1.xxx`.

LinuxThreadStackSize Used to set stack size in MB during thread creation on Linux. Use to override the default system size of 1 MB.

Functional Settings

InfoStreamCRStatus Association of status between completed user stories (created using Info Streams) and Original Change Request. When enabled, a Story is marked complete in agile and triggers the corresponding Change Request to be marked fixed. Default value: `1` . Valid values: `0` (disabled) or `1` (enabled).

DisableAdvancedViews

Disables creation of certain type of views. Default value: `1`. Valid values: `0` (enabled) or `1` (disabled).

NotificationLocale

Override the locale for notification messages. Default value: `empty`, using server locale. Valid values: `En, fr, de, pt, zh-cn, or ja`.

VaultCompressionCutOffThreshold

File content is compressed in the archives directory only if the compressed file is smaller than the threshold value (in percent). Default value: `10`. Valid range: `0-100`.

WorkspaceModifiedTimeFeature

Enable/disable View Workspace Modified Time feature. Default value: `1`. Valid values: `0` (disabled) or `1` (enabled).

AtomicBehavior

Enable/disable Atomic transactions for VCM. Default value: `0`. Valid values: `0` (enabled) and `1` (disabled).

VerboseLevel

Level of logging detail. Default value: `0`. Valid range: `0-6`.

MixedEOLAsBinary

Treat files containing a mix of EOF characters as binary. Default value: `0`. Valid values: `0` (no) or `1` (yes).

UpgradeMissingAttachments

Upgrade attachments even if upgrade form 12.0 was previously run. Used to force re-migration of attachments. Default value: `0` 0. Valid values: `0` (disabled) or `1` (enabled).

ViewsCachedAtServerStartup

IDs of views to be cached on startup. Example: `22, 57, 1547`.

CheckinChangePackages

Create checkin change packages during file checkin. Default value: `1`. Valid values: `0` (disabled) or `1` (enabled).

ForceCustomTools

Enable/Disable custom tools. Default value: `0`. Valid values: `0` (disabled) or `1` (enabled).

CheckoutUseRevisionFileName

Checkout files using revision's file name. Default value: `0`. Valid values: `0` (disabled) or `1` (enabled).

AllowExternalTraces

Enable/disable external traces. Default value: `1`. Valid values: `0` (disabled) or `1` (enabled).

`DefaultBinaryExtensions`

You can specify a semi-colon separated list of extensions which will automatically be treated as binary files on check-in. For example:

```
DefaultBinaryExtensions=”.pdf;.bin;.s19;.doc”
```

InfoStreamCRStatus

Association of status between completed user stories (created using Info Streams) and Original Change Request. When enabled, a Story is marked complete in agile triggers the corresponding CR to be marked fixed.

Default value: `0`. Valid values: `0` (disabled) or `1` (enabled).

SDKWorkflow

Used to enable/disable workflow in the command line SDK.

Default value: `0`. Valid values: `0` (disabled) or `1` (enabled).

DefaultBinaryExtensions

List of file extensions to treat always as binary. Use quotes ("") to denote the extensions. For example, `"pdf;.doc;.xls;"`.

URL Settings

StarTeamAgileWebAddress	Web address of StarTeam Agile used by the Cross-Platform Client to access StarTeam Agile. Default value: "". Example: `http://agile:8080/topaz-web/login.jsp`?
ALMServiceURL	URL of REST service, used to generate notification URL. Default value: `<ServerHostnameUrl>:80`. Example: `"hostname:portnumber".`
RESTServicePort	Rest service port over-ride. This value must match the Apache Tomcat Port Number. Default value: `9090`. Valid range: `0-10000`.

Purge Settings

OnlinePurgeThreadPriority	Online purge thread priority. Default value: `Background`. Valid values: `Background` or `Foreground`.
OnlinePurgeExecTime	Online purge execution time, in seconds. Default value: `60`. Valid range: `10-400`.
OnlinePurgeLoggingType	Online purge log type. Default value: `0`. Valid values: `0` (normal) or `1` (verbose).

Recommended Configuration Options

StarTeam has a wide variety of customers ranging from small development shops with a few developers to large enterprises with tens of thousands of concurrent users. It is practically impossible to have *one size fits all* in terms of hardware recommendations and configuration options.

There are no hard rules about what makes a configuration *small*, *medium*, or *large*. However, for our purposes, we'll use these definitions based on concurrent users:

Small configuration	Supports no more than 50 concurrent users.
medium configuration	Supports no more than 100 concurrent users.
large configuration	Supports 100+ concurrent users.

This does not factor data volume or the type of users: on-line users or bulk applications. This is because, in our experience, the amount of data managed by a configuration (particularly items) tends to grow proportionally with the number of projects and views, which grows in proportion to the team size. Moreover, the ratio of on-line users to bulk applications tends to be roughly the same across organization sizes. The concurrent user count seems to be the best metric for judging configuration size for purposes of deployment planning.

So how big can a configuration get? To date, we've seen single instances with over 500 concurrent users, over 10,000 total define" users, over 4,000 views, tens of millions of items, and up to a terabyte of vault data. With continuous hardware advances and software improvements, these limits get pushed every year.

Setting Name	Small	Medium	Large
`DBMinimumConnections`	10	10	70
`DBMaximumConnections`	25	50	80
`MinCommandThreads`	16	40	80
`MaxCommandThreads`	25	50	80
`ItemCachePriority_File*`	100	100	100
`ItemCachePriority_Folder`	0	0	0
`ItemCachePriority_Changes*`	100	100	100

Setting Name	Small	Medium	Large
`ItemCachePriority_Topic`	0	0	0
`ItemCachePriority_Task`	25	25	25
`ItemCachePriority_Requirement`	0	0	0
`ItemCacheLogRecovery`	1	1	1
`ItemCacheMemoryLimit` This setting depends on the physical memory. These values assume that the machine is exclusively used for StarTeam Server. When the physical memory is not constrained on 64-bit machines , this can be set to `-1`. -1 for 64-bit translates to 8GB of IDS cache.	• Win32: 300 • Win32/PAE: 500 • WoW64: 700 • Win64: 60% of physical memory when physical memory is 8GB or more (In MB) For lowerend machines use recommended values for Wow64.	• Win32: 500 • Win32/PAE: 600 • WoW64: 800 • Win64: 60% of physical memory when physical memory is > 8GB (in MB). For lower-end machines use recommended values for Wow64	• Win32 : 700 • Win32/PAE : 600 • WoW64: 1000 • Win64: 60% of physical memory when physical memory is > 8GB (in MB) For lower-end machines use recommended values for Wow64.

*The recommended values for `CachePriority` for all the artifacts depends on the specific use case. For customers who have Change-Request only configurations, it would be desirable to set `ItemCachePriority_Changes` to `100` and reset all the others to `0`. For files-only configurations, it is recommended to set `ItemCachePriority_File` to `100` and reset the other options to `0`.

StarTeam.Log

The `StarTeam.Log` file records the operations performed on your client workstation during a session. It helps you troubleshoot and document errors or operations between the server and your workstation that failed during server configuration sessions. The file may contain commands sent by your workstation to a server configuration when you open and work with a project, commands performed locally on your workstation, error messages generated while using the application, or events performed by StarTeamMPX.

Every time you start your client, the system creates a `StarTeam.Log` file in the folder location specified in your personal options. On most systems, the default location for the log file is `%PROGRAMDATA%\Borland\StarTeam`. If there is a log file already in this folder, The application renames the existing file to include the date and time at which it was renamed. For example, if you create a `StarTeam.Log` file on November 9, 2011 at 10:35 A.M., the old log file is renamed `StarTeam-09-Nov-11-10-35-18.Log`, and a new log file is created.

Because the application creates a new `StarTeam.Log` file every time you start the client, the log folder can fill up quickly. To control the number of log files in the folder, you may want to periodically delete old log files from the output folder or disable the log option. To disable the option, clear the **Log Errors** and the **Log Operations** check boxes on the **Workspace** tab of the **Personal Options** dialog. To display the `StarTeam.Log` file, select **Tools** > **StarTeam Log File** from the menu bar. You can also import and view the data from a log file using any application that supports tab-delimited fields. For example, if you save the file with a `.csv` extension, the file can be opened in Microsoft Excel.

The **Workspace** tab on the **Personal Options** dialog enables you to specify the location and the type of data recorded in the `StarTeam.Log` file

Displaying and Customizing StarTeam.Log

The `StarTeam.Log` file records the operations performed on your client workstation during a session. It helps you troubleshoot and document errors or operations between the server and your workstation that failed during server configuration sessions.

Because the application creates a new file every time you start the client, the log folder can fill up quickly. To control the number of log files in the folder, you may want to periodically delete old log files from the output folder or disable the `StarTeam.Log` option. To disable the option, clear the **Log Errors** and the **Log Operations** check boxes on the Workspace tab of the **Personal Options** dialog. To display the `StarTeam.Log` file, select **Tools** > **StarTeam Log File** from the menu bar. You can also import and view the data using any application that supports tab-delimited fields. For example, if you save the file with a `CSV` extension, the file can be opened in Microsoft Excel.

The **Workspace** tab on the **Personal Options** dialog enables you to specify the location and the type of data recorded in the `StarTeam.Log` file.

1. From a client, select **Tools** > **Personal Options** . The **Personal Options** dialog box appears.
2. On the **Workspace** tab, enter or browse for the location of the `StarTeam.Log` file in the **Log Output Path** text box. The default is the location in which the application is installed, for example, `%PROGRAMDATA%\Borland\StarTeam`. The current log file is always named `StarTeam.Log`. Log files from earlier sessions of the application include the date and time the file was last modified.

   Note: `StarTeam.Log` contains data about operations sent from your workstation to one or more servers, depending on what project views you have open. This data includes the name of the project so that you can isolate data for a particular server, when necessary.
3. Select the types of data you want to include in `StarTeam.Log`.

   Log errors
   Set by default. Records errors that occur while you are using the client. The errors log lists the date and time you started your server configuration and any errors or failed operations between the server and client. The application identifies each failed operation by an internal ID and provides an explanation. For example, you might see:`...Operation 40956 failed: TCP/IP Socket Error 10054:...`

   Log StarTeamMPX events
   Selecting this option records information about StarTeamMPX events for this client. The log identifies the time and date on which a StarTeamMPX event (an automatic refresh or file status update) took place. The log prefaces a StarTeamMPX event as “Statistics for Events” and uses internal IDs and brief explanations to identify the server event. The following example describes a status change for a file:`...Statistics for Events /1b21dd1-e208-51ea-01b2-1dd1e20851ea/Object/File/ Modify` You can log StarTeamMPX events only if you check the “Enable StarTeamMPX” checkbox on the StarTeamMPX tab. For StarTeamMPXrelated operations, any changes you make on the Workspace tab do not apply to projects already open. However, the application will log StarTeamMPX events for any projects you open from this point forward.

   Log operations that take at least ___ milliseconds
   Select this option to record operations that take longer than a specific number of milliseconds. (An operation is a command that results from user actions. Operations can be executed on either the Server or the client.) The milliseconds time setting stops the log from filling up with operations of little importance. The default is 10 milliseconds. This option records information on the date, time, and UI Operation number for each command executed by your workstation. You can record either Summary or Detail information.

Summary Records the time spent for the total operation, client execution time, and server execution time. The application identifies each operation by an internal ID, such as Statistics for Operation 40001.

Detail Records a detailed breakdown of all server commands performed for each operation. The log identifies the server address, project, and component (File, Change Request, Requirement, Task, or Topic) for each server command. The application identifies each server command by an internal ID, such as Public Server Command 10.

4. Click **OK**.

Audit Logs

By default, the Server is automatically configured to generate audit logs. With this option activated, the Server logs audit events for projects in the server configuration database. For example, the log records when change requests are created, and when a file is added. The audit log entries can be viewed from a client by selecting the Audit tab in the upper pane. This operation can be performed only on a server configuration that is running.

A chronological record, the Audit log accumulates data about the actions performed on folders, files, requirements, change requests, tasks, and topics. Each log entry shows the user who carried out the action, the date and time the action was performed, the class name (type of item), the event (type of action), the view name, and the project name. By using filters or queries, you can locate all the entries for a particular item.

For most items, events may be added, branched, commented, created, deleted, modified, moved from, moved to, and shared. For files, events may also include converted, edited, item overwritten, locked, lock broken, and unlocked. Log entries themselves cannot be moved, shared, modified, or branched. If the Audit tab of the main window displays no entries, your administrator has probably disabled the Audit log function.

Enabling and Purging the Audit Log

When you select the Enable Audit Generation option, the Server logs audit events for projects in the server configuration database. For example, the log records when change requests are created, and when a file is added. The audit log entries can be viewed from a client by selecting the Audit tab in the upper pane. This operation can be performed only on a server that is running.

Note: If setting the option to purge logs on server configuration startup, you need to restart your server configurations fairly regularly to avoid startup problems.

To enable the audit log

1. Open the Server Administration tool. If you are using the Server Administration tool installed with the client, you can administer remote servers only.
2. From the list of servers, select the server configuration that you want to change. If you have not yet logged on, you will be asked to do so.
3. Click the **Configure Server** shortcut or **Tools** > **Administration** > **Configure Server**. The **Configure Server** dialog box appears.
4. Select the **Audits** tab.
5. Select the **Enable Audit Generation** check box.
6. Optionally, to automatically delete entries after a specified length of time, select the **Purge Audit Entries Older Than** check box. (Clearing this check box keeps the entries indefinitely.) Type a number of days in the Days text box. The range is from 7 to 1000 days. For example, to delete entries when they become approximately one month old, type 30 days in the Day text box. When the server configuration starts, entries that exceed this purge limit are deleted.

7. Click **OK**.

Security Logs

The application's clients and servers generate a number of log files. These logs enable an administrator to evaluate the performance of the system and potentially troubleshoot problems. Each server configuration has its own server log and security log. Each client creates its own log file, which records activity between that client and the server configurations it is connected to.

Users must have the appropriate security access rights in order to view a log file. These access rights can be set using the **Tools** > **Accounts** > **Access Rights** menu option in the **Server Administration** tool.

Server Log Files

The server log file (`Server.locale.Log`) records the activity on a server configuration. Each time you start a server configuration, the Server renames the existing log file and creates a new log file for the current server configuration session. The log file from the previous startup is renamed to include the date and time at which it was renamed (`Server.locale.date.Log`). For example, if you start a server configuration on November 9, 2011 at 5:22 P.M., the old `Server.locale.Log` file is renamed `Server.en-US.2011-11-09-17-22-59.Log` and a new `Server.locale.Log` file is created whose time stamp might be `11/9/2011 17:23:03`.

If the locale specified for the operating system on which your server runs is not US English, you will have two server log files: one for US English and one for your locale. For example, you might have both `Server.en-US.Log` and `Server.fr-FR.Log`. The first log is for support purposes, and the second log is for your use.

You can view the contents of the server log file at any time, even while the server configuration is running by choosing **Tools** > **Administration** > **Server Log** .

Security Log Files

A security log records all security-related events for a server configuration. For each secured event (such as logging on or off), the security log records the date and time it occurred, the user performing the operation, the workstation from which the operation was performed, the item acted upon, and whether the operation failed.

Depending upon the number of users and the amount of activity on a server configuration, the security log may grow rapidly. To keep the log to a reasonable size, you can have the server delete old entries. First, decide how long you want to have security events available, then configure the server configuration to purge entries that are older than this time period. See "Working with the Security Log" topic in Related Information for how to purge security log entries.

If you have access rights to a server configuration, you can view its security log at any time the server is running. The security log is not a typical log file, as its data is stored in the application database. The security log is available by choosing **Tools** > **Accounts** > **Security Log** .

StarTeam.Log File

The `StarTeam.Log` file records the operations performed on your client workstation during a session. It helps you troubleshoot and document errors or operations between the server and your workstation that failed during server configuration sessions.

The `StarTeam.Log` file may contain the following types of information:

- Commands sent by your workstation to a server configuration when you open and work with a project. If you work with projects on several different server configurations, you can configure the `StarTeam.Log` file to include the server configuration name with the command information it records.
- Commands performed locally on your workstation, such as setting personal options.

- Error messages generated while using the application.
- Events performed by StarTeamMPX.

StarTeam.Log File Creation

Every time you start your client, the system creates a `StarTeam.Log` file in the folder location specified in your personal options.

StarTeam.Log File Location

On most systems, the default location for the `StarTeam.Log` file is `C:\Program Files\Borland\StarTeam x.x`. If there is a `StarTeam.Log` file already in this folder, the application renames the existing file to include the date and time at which it was renamed. For example, if you create a `StarTeam.Log` file on November 9, 2011 at 10:35 A.M., the old `StarTeam.Log` file is renamed `StarTeam-09-Nov-11-10-35-18.Log`, and a new `StarTeam.Log` file is created.

Tip: Tip Because the application creates a new `StarTeam.Log` file every time you start the client, the log folder can fill up quickly. To control the number of log files in the folder, you may want to periodically delete old log files from the output folder or disable the `StarTeam.Log` option. To disable the option, clear the **Log Errors** and the **Log Operations** check boxes on the **Workspace tab** of the **Personal Options** dialog.

Working with the Security Event Log

If you have access rights to a server configuration, you can view its security event log at any time. The security event log is not a typical .Log file, as its data is stored in the application database. This operation can be performed only when the server is running.

To view the security event log

1. Open the Server Administration tool. If you are using the Server Administration tool installed with the client, you can administer remote servers only.
2. Select the appropriate server configuration from the list of servers. If you have not yet logged on, you will be asked to do so.
3. Select **Tools** > **Accounts** > **Security Log** . These actions display the Security Log contents. This log lists each secured event (such as logging on or off), the date and time it occurred, the user performing the operation, the workstation from which the operation was performed, the item acted upon, and whether the operation failed.
4. Use the **Security Event Type** drop-down list box to view all events of a particular type.
5. To reload the security event log and review the most recent entries, click **Reload** from the **Security Event Log** dialog box.
6. To print the data selected from the log, click **Print Selection** from the **Security Event Log** dialog box.

Depending upon the number of users and the amount of activity on a server configuration, the security event log may grow rapidly. To keep the log to a reasonable size, you can have the Server delete old entries. First, decide how long you want to have security events available, then configure the server configuration to purge entries that are older than this time period. This operation can be performed only when the server is running.

To set the interval for purging the security event log

1. Open the Server Administration tool. If you are using the Server Administration tool installed with the client, you can administer remote servers only.
2. Select the appropriate server configuration from the list of servers. If you have not yet logged on, you will be asked to do so.
3. Click **Tools** > **Accounts** > **System Policy** from the menu. The **System Policy** dialog box appears.

4. Select the **Security Events** tab.
5. Select the **Purge Security Event Entries Older Than __ Days** check box. (Clearing this check box keeps the entries indefinitely.)
6. Type the number of days in the text box. The range is 30 to 1000. The default is 180. The next time the server configuration starts, entries that exceed the purge limit are deleted.
7. Click **OK**.
8. Restart the server configuration for the purge interval to take effect.

Security Event Types

If you have access rights to a server configuration, you can view its security event log at any time. The security event log is not a typical .Log file, as its data is stored in the application database. This operation can be performed only when the server is running.

Add item owner	Indicates that a user created a folder or an item.
Add user/group	Indicates that a user or group was added to the server configuration.
Add/Edit container access rights	Indicates that access rights were added or changed for a group of objects contained in another object. For example, if you select **Project** > **Access Rights** and change rights for all change requests in the project, that event fits into this category.
Add/Edit item access rights	Indicates that access rights were added or changed for a specific object. For example, if you change access rights for a project, that event fits into this category.
Change user	Indicates that someone changed user names as part of a replication process. This event can occur when special clients, such as Notification Agent, perform operations.
Delete container access rights	Indicates that access rights were deleted at the container level.
Delete item access rights	Indicates that access rights were deleted at the item level.
Delete user/group	Indicates that a user or group was deleted.
Edit user/group	Indicates that the properties for a user or group were changed in some way.
Force user logoff	Indicates that a user was forced to log off the server configuration.
Item access check	Indicates that access rights were checked to see if the user could access a specific item.
Logoff	Indicates that a user logged off the server configuration.
Logon	Indicates that a user logged on to the server configuration.
Logon attempt: Account lockout	Indicates that a user attempted to log on and the account was locked.
Logon attempt: Expired password	Indicates that a user attempted to log on and the password had expired.
Logon attempt: No such user name	Indicates that a user attempted to log on with a non-existent user name.
Logon attempt: Restricted access time	Indicates that a user attempted to log on at a time when he or she was not allowed access.
Logon attempt: Suspended account	Indicates that a user attempted to log on and the account was suspended.
Logon failure	Indicates that an incorrect password was used during the logon process.

Policy change	Indicates that a system policy has changed.
User status change	Indicates that an administrator suspended, reactivated, locked, unlocked, or required a password change on a user’s account.

Working with the Server Log

You can view the contents of the server log file at any time, even while the server configuration is running. To see the entire file, use Notepad, WordPad, or another text editor to display it. To determine the location of a server log file from the Server Administration tool:

1. Open the Server Administration tool. If you are using the Server Administration tool installed with the client, you can administer remote servers only.
2. Select the appropriate server configuration from the list of servers. If you have not yet logged on, you will be asked to do so.
3. Click the **Configure Server shortcut** or **Tools** > **Administration** > **Configure Server** . These actions display the **Configure Server** dialog box.
4. Look at the top of the **General** tab to find the location of the log file.

To review the contents of a server log file

1. Open the Server Administration tool. If you are using the Server Administration tool installed with the client, you can administer remote servers only.
2. Select the appropriate server configuration from the list of servers. If you have not yet logged on, you will be asked to do so.
3. Click the **Server Log** shortcut or **Tools Administration Server Log**. These actions display the Server Log contents.

   The format provides a line number, code, date and time, and message. The code numbers are arbitrarily assigned and not necessarily in order of severity. They represent the following:

   ```
   00000001 Message
   00000002 Warning
   00000004 Error
   00000008 Unexpected Condition
   ```
4. To display only the errors in the log, select the **Errors Only** check box.
5. To review the most recent entries, click **Reload**.

On Windows systems, you can copy data from the log window to the Windows clipboard. From the clipboard, you can paste the data into other applications, such as Microsoft Word or Notepad.

To copy data from the server log

1. From the Server Administration tool, click the **Accounts** bar in the lower left pane and select the **Server Log** Shortcut.
2. Select the data that you want to copy.
3. Press **Ctrl + C** .
4. Click **OK** to exit the dialog.
5. Using **Ctrl + V** , paste the information into a text editor or word processing application.
6. Click **File** > **Print** from the menu to print the data.

Any users in the System Managers group, a default, will receive email when an error is logged in the server log. This group is initially empty.

Server Log

The server log file (`Server.locale.Log`) specifies the DBMS version and build number and records the activity on a server configuration. Each time you start a server configuration, the Server renames the existing log file and creates a new log file for the current server configuration session. The log file from the previous startup is renamed to include the date and time at which it was renamed (`Server.locale.date.Log`). For example, if you start a server configuration on November 9, 2011 at 5:22 P.M., the old `Server.locale.Log` file is renamed `Server.en-US.2011-11-09-17-22-59.Log`, and a new `Server.locale.Log` file is created whose time stamp might be 11/9/2011 17:23:03.

If the locale specified for the operating system on which your server runs is not US English, you will have two server log files: one for US English and one for your locale. For example, you might have both `Server.en-US.Log` and `Server.fr-FR.Log`. The first log is for support purposes, and the second log is for your use.

Formatting of the log gives a line number, code, date and time, and message. The code numbers are not in any order of severity.

Server Log Error Codes

Formatting of the log gives a line number, code, date and time, and message. The code numbers are not in any order of severity.

```
Line #        Code          Date          Time           Error Message
-------------------------------------------------------------------------
45     00000001     2011-05-19       05:05:08          Message
46        00000002     2011-05-19       05:05:10          Warning
47        00000004     2011-05-19       05:05:12          Error
48        00000008     2011-05-19       05:05:14          Unexpected Condition
```

DbConvert.<local>.log

This log records the progress of database migration. Migration is a process of creating a new database (destination database) and copying data from an existing database (source database) into it.

DbConvert.log consists of two parts (when starting a new migration):

- Server log from starting a new configuration with a destination database. Server log part is no different than the log from a normal server startup/shutdown for a new configuration.
- Migration log from the migration of the source database to the destination database. The Migration log itself starts after the line`***** Server shutdown complete *****` It logs the source configuration name and the destination configuration name `Started database conversion: configuration "Test" to configuration "TestMigrated"..."` It lists all the tables as their data is copied from the source database to the destination database: "`<tablename> migrated successfully”` At the end, the source configuration is disabled and the destination configuration is enabled. `Source config successfully disabled Target config successfully enabled` It ends with a statement: `Migration completed successfully`.

If migration stopped and then re-started at a later time, the server log part might be missing in the DbConvert.log, if the destination database was already created by the previous run of migration. A sample of this log is displayed below:

```
1         00000001  2009-03-29 11:08:28  Microsoft Windows Server 2003 family
Service Pack 2 (Build 3790)

3         00000001  2009-03-29 11:08:29  Started database conversion:
configuration "newsql" to configuration "NewORA10R2_Hamachi"...
4         00000001  2009-03-29 11:08:30  Catalog Tables converted successfully
5         00000001  2009-03-29 11:08:38  Catalog Fields converted successfully
6         00000001  2009-03-29 11:08:54  Catalog table converted successfully
7         00000001  2009-03-29 11:08:54  Catalog table converted successfully
8         00000001  2009-03-29 11:08:54  Catalog table converted successfully
9         00000001  2009-03-29 11:08:56  Microsoft Windows Server 2003 family
Service Pack 2 (Build 3790)

11        00000001  2009-03-29 11:08:57  ServerSettings migrated successfully
```

```
12        00000001  2009-03-29 11:08:57  CommProtocol migrated successfully
13        00000001  2009-03-29 11:08:57  IPRangeObject migrated successfully
14        00000001  2009-03-29 11:08:57  User migrated successfully
15        00000001  2009-03-29 11:08:57  Group migrated successfully
16        00000001  2009-03-29 11:08:58  GroupMembers migrated successfully
17        00000001  2009-03-29 11:08:58  AccessControlData migrated
successfully
18        00000001  2009-03-29 11:08:58  SystemPolicyObject2 migrated
successfully
19        00000001  2009-03-29 11:08:58  ObjectSecurityLog migrated
successfully
20        00000001  2009-03-29 11:08:58  ProfileType migrated successfully
21        00000001  2009-03-29 11:08:58  ProfileData migrated successfully
22        00000001  2009-03-29 11:08:58  Merge migrated successfully
23        00000001  2009-03-29 11:08:58  Project migrated successfully
24        00000001  2009-03-29 11:08:58  View migrated successfully
25        00000001  2009-03-29 11:08:58  Folder migrated successfully
26        00000001  2009-03-29 11:08:58  Folder_QNodes migrated successfully
27        00000001  2009-03-29 11:08:58  Folder_QParts migrated successfully
28        00000001  2009-03-29 11:08:58  Folder_Queries2 migrated successfully
29        00000001  2009-03-29 11:08:58  Folder_Filters2 migrated successfully
30        00000001  2009-03-29 11:08:59  Folder_FColumns migrated successfully
31        00000001  2009-03-29 11:08:59  ViewMember migrated successfully
32        00000001  2009-03-29 11:08:59  ConfigLabel migrated successfully
33        00000001  2009-03-29 11:08:59  ConfigLabelEntry migrated
successfully
34        00000001  2009-03-29 11:09:00  Link migrated successfully
35        00000001  2009-03-29 11:09:00  LinkPin migrated successfully
36        00000001  2009-03-29 11:09:01  PromotionDefinition migrated
successfully
37        00000001  2009-03-29 11:09:01  PromotionModel migrated successfully
38        00000001  2009-03-29 11:09:01  PromotionState migrated successfully
39        00000001  2009-03-29 11:09:01  PromotionStatus migrated successfully
40        00000001  2009-03-29 11:09:01  ItemLock migrated successfully
41        00000001  2009-03-29 11:09:01  File migrated successfully
42        00000001  2009-03-29 11:09:01  Files_BookmarkObjects migrated
successfully
43        00000001  2009-03-29 11:09:01  Files_QNodes migrated successfully
44        00000001  2009-03-29 11:09:01  Files_QParts migrated successfully
45        00000001  2009-03-29 11:09:02  Files_Queries2 migrated successfully
46        00000001  2009-03-29 11:09:02  Files_Filters2 migrated successfully
47        00000001  2009-03-29 11:09:02  Files_FColumns migrated successfully
48        00000001  2009-03-29 11:09:02  Change migrated successfully
49        00000001  2009-03-29 11:09:02  Changes_UnreadObjects migrated
successfully
50        00000001  2009-03-29 11:09:02  Changes_BookmarkObjects migrated
successfully
51        00000001  2009-03-29 11:09:02  Changes_Attachments migrated
successfully
52        00000001  2009-03-29 11:09:03  Changes_QNodes migrated successfully
53        00000001  2009-03-29 11:09:03  Changes_QParts migrated successfully
54        00000001  2009-03-29 11:09:04  Changes_Queries2 migrated
successfully
55        00000001  2009-03-29 11:09:04  Changes_Filters2 migrated
successfully
56        00000001  2009-03-29 11:09:07  Changes_FColumns migrated
successfully
57        00000001  2009-03-29 11:09:07  Requirement migrated successfully
58        00000001  2009-03-29 11:09:07  Requirements_Attachments migrated
successfully
59        00000001  2009-03-29 11:09:07  Requirements_UnreadObjects migrated
successfully
60        00000001  2009-03-29 11:09:07  Requiremen_BookmarkObjects migrated
successfully
```

```
61        00000001  2009-03-29 11:09:07  Requirements_QNodes migrated
successfully
62        00000001  2009-03-29 11:09:07  Requirements_QParts migrated
successfully
63        00000001  2009-03-29 11:09:07  Requirements_Queries2 migrated
successfully
64        00000001  2009-03-29 11:09:08  Requirements_Filters2 migrated
successfully
65        00000001  2009-03-29 11:09:08  Requirements_FColumns migrated
successfully
66        00000001  2009-03-29 11:09:08  Task migrated successfully
67        00000001  2009-03-29 11:09:08  WorkRecord migrated successfully
68        00000001  2009-03-29 11:09:08  Dependencies migrated successfully
69        00000001  2009-03-29 11:09:08  Tasks_UnreadObjects migrated
successfully
70        00000001  2009-03-29 11:09:08  Tasks_BookmarkObjects migrated
successfully
71        00000001  2009-03-29 11:09:08  Tasks_Attachments migrated
successfully
72        00000001  2009-03-29 11:09:09  Tasks_QNodes migrated successfully
73        00000001  2009-03-29 11:09:09  Tasks_QParts migrated successfully
74        00000001  2009-03-29 11:09:09  Tasks_Queries2 migrated successfully
75        00000001  2009-03-29 11:09:09  Tasks_Filters2 migrated successfully
76        00000001  2009-03-29 11:09:09  Tasks_FColumns migrated successfully
77        00000001  2009-03-29 11:09:10  Topic migrated successfully
78        00000001  2009-03-29 11:09:10  Topics_Attachments migrated
successfully
79        00000001  2009-03-29 11:09:10  Topics_UnreadObjects migrated
successfully
80        00000001  2009-03-29 11:09:10  Topics_BookmarkObjects migrated
successfully
81        00000001  2009-03-29 11:09:10  Topics_QNodes migrated successfully
82        00000001  2009-03-29 11:09:10  Topics_QParts migrated successfully
83        00000001  2009-03-29 11:09:10  Topics_Queries2 migrated successfully
84        00000001  2009-03-29 11:09:10  Topics_Filters2 migrated successfully
85        00000001  2009-03-29 11:09:11  Topics_FColumns migrated successfully
86        00000001  2009-03-29 11:09:12  Trace migrated successfully
87        00000001  2009-03-29 11:09:12  Traces_LinkValues migrated
successfully
88        00000001  2009-03-29 11:09:12  Traces_UnreadObjects migrated
successfully
89        00000001  2009-03-29 11:09:12  Traces_BookmarkObjects migrated
successfully
90        00000001  2009-03-29 11:09:12  Traces_Attachments migrated
successfully
91        00000001  2009-03-29 11:09:12  Traces_QNodes migrated successfully
92        00000001  2009-03-29 11:09:12  Traces_QParts migrated successfully
93        00000001  2009-03-29 11:09:12  Traces_Queries2 migrated successfully
94        00000001  2009-03-29 11:09:12  Traces_Filters2 migrated successfully
95        00000001  2009-03-29 11:09:12  Traces_FColumns migrated successfully
96        00000001  2009-03-29 11:09:13  Audits migrated successfully
97        00000001  2009-03-29 11:09:13  Audits_QNodes migrated successfully
98        00000001  2009-03-29 11:09:13  Audits_QParts migrated successfully
99        00000001  2009-03-29 11:09:13  Audits_Queries2 migrated successfully
100       00000001  2009-03-29 11:09:13  Audits_Filters2 migrated successfully
101       00000001  2009-03-29 11:09:13  Audits_FColumns migrated successfully
102       00000001  2009-03-29 11:09:14  ChangePackage migrated successfully
103       00000001  2009-03-29 11:09:14  ChangePackageChange migrated
successfully
104       00000001  2009-03-29 11:09:14  ChangeReference migrated successfully
105       00000001  2009-03-29 11:09:14  ChangePackag_UnreadObjects migrated
successfully
106       00000001  2009-03-29 11:09:14  ChangePack_BookmarkObjects migrated
successfully
```

```
107       00000001  2009-03-29 11:09:14  ChangePackages_Attachments migrated
successfully
108       00000001  2009-03-29 11:09:14  ChangePackages_QNodes migrated
successfully
109       00000001  2009-03-29 11:09:14  ChangePackages_QParts migrated
successfully
110       00000001  2009-03-29 11:09:14  ChangePackages_Queries2 migrated
successfully
111       00000001  2009-03-29 11:09:14  ChangePackages_Filters2 migrated
successfully
112       00000001  2009-03-29 11:09:15  ChangePackages_FColumns migrated
successfully
113       00000001  2009-03-29 11:09:15  Workstation migrated successfully
114       00000001  2009-03-29 11:09:15  Source config successfully disabled
115       00000001  2009-03-29 11:09:15  Target config successfully enabled
116       00000001  2009-03-29 11:09:15  Migration completed successfully.
```

Analyze Server Log and Analyze Archived Server Logs

The Analyze Server Log and the Analyze Archived Server Logs analyze the server log(s) and report the connections usage over time.

Choose **Tools** > **Administration** > **Server Administration** to access these logs.

Analyze Server Log Connect to the current server, obtains most recent server log and runs analysis, displaying a graph of Connections Count over time, and a report with multiple tabs:

Tab Name	Description
Users	List of users that connected, total connection time cumulated over all connections, and license type used.
Workstations	List of workstations used to connect, including total connection time.
Clients	List of client applications used for connecting, including total connection count.
SDK	List of StarTeam SDK versions used for connecting, including total connection count.
Java	List of Java versions used to connect, including total connection count.
Operating System	List of operating system versions used to connect, including total connection count.

Analyze Archived Server Logs Selects a number of StarTeam Server log files. The application runs analysis over the combined logs, reporting the same as above but cumulative for all the server logs selected.

Tracing Data from Check-out Operations

Tracing Data from Check-out Operations with the Check-out Trace Utility

The StarTeam check-out Trace utility generates a `*.csv` file that provides data about check-out operations for the server configuration for which tracing is enabled. Before you run the utility, you must enable tracing for the server configuration in the `starteam-server-configs.xml` file. With tracing enabled, the server generates a trace record for each checked out file and saves the information in a trace file (`check-`

out.cotrc). The utility uses the trace file as input and outputs a *.csv file containing data about the check-out operations. You can import the output from the *.csv file into Datamart or an Excel spreadsheet.

The *.csv file contains the following information for each check-out:

Note: Checkout data will not be included in the generated .cotrc file if a Cache Agent performed the checkout. Data will only be included in the .cotrc file if the check-out operation was performed by the Server.

- user ID
- user name
- time stamp (date/time of check-out)
- project ID
- project
- view
- view ID
- folder ID
- folder path
- file ID
- filename
- file revision number

Note: To optimize performance, StarTeam does not immediately update trace files. StarTeam buffers the information for the trace file in memory and writes it to the trace file during idle time processing.

You can find the check-out Trace utility in the StarTeam Server root installation folder (CheckoutTraceDump.exe). For information about using the utility, refer to the links at the bottom of this topic.

Enabling Tracing for Server Configurations Manually

When you enable tracing for a server configuration, the server saves a trace record for each file that is checked out in a trace file (check-out.cotrc). You then use the trace file as input for the Check-out Trace utility to generate a *.csv file containing information about check-out operations.

To manually enable tracing for a server configuration

1. Open the starteam-server-configs.xml file. You can locate the file under the StarTeam Server root installation folder.
2. Update the following elements with a value of "1" for each server configuration that you wish to enable tracing on:

   ```
   <option name="FileAllowCheckoutTrace" value="1"/>
   ```

   ```
   <option name="FileEnableCheckoutTrace" value="1"/>
   ```

   The first option activates the tracing code. While the second option turns tracing on or off. A value of "1" represents true or on. A value of "0" represents false or off. To enable tracing, both values must be "1".
3. Save the changes to starteam-server-configs.xml.
4. Shut down and restart the server configuration so that it can detect changes from starteam-server-configs.xml.

Note: To enable tracing using the Server Administration Tool, refer to *Activating Diagnostic Tests*.

When you set both options in starteam-server-configs.xml to "1", the server configuration creates Checkout.cotrc files in the *Trace* folder (a subfolder of the repository folder Checkout.cotrc). When the size of the current trace file reaches 128 MB, the server saves the current trace file and creates a new trace file. The name of the older trace file becomes the name plus a time-stamp, similar to the time-stamp naming convention found in StarTeam `server.log` files. When you shutdown the server configuration, the server saves the trace file with a time-stamp appended to the filename. When you restart the server configuration, the server creates a new trace file.

Note: To optimize performance, StarTeam does not immediately update trace files. StarTeam buffers the information for the trace file in memory and writes it to the trace file during idle time processing.

Generating .CSV Files About Check-out Operations

Before you run the utility, you must enable tracing for the server configuration. With tracing enabled, the server generates a trace record for each checked out file and saves the information in a trace file (`Checkout.cotrc`). The utility uses the trace file as input and outputs a `*.csv` file containing data about the check-out operations.

The Check-out Trace utility takes one or more check-out trace (`*.cotrc`) files as input and outputs one `*.csv` text file containing check-out trace data as comma-delimited values. The default filename for the `.csv` file is identical to the name of the trace file with the extension `.csv`. For example, when the trace filename is `Checkout.cotrc`, then the `csv` output filename is `Checkout.cotrc.csv`.

To run the Check-out Trace utility

1. At the command prompt, navigate to the `CheckoutTraceDump.exe` file in the StarTeam Server root installation folder.
2. The `-go` option signals the utility to run with default options. You can set many parameters for the utility. For a list of all of the available options, review the command line operations for the utility at the link referenced at the bottom of this topic.

   Note: By default, the server saves the trace files in the Trace folder (a subfolder of the repository folder `Checkout.cotrc`). You cannot run the utility against the current trace file, but you can copy the trace file and run the utility against the copy.

Tip: If you want to run the utility from a workstation rather than on the server, you can copy `CheckoutTraceDump.exe` and `OSSup.dll` to the alternate location. Be cautious not to *move* `OSSup.dll` to the new location because the server configuration also relies on it. Additionally, the utility depends on the Microsoft .NET Runtime, so it must be available on the alternate workstation.

Data Storage Options

Data Storage Overview

As part of creating a new server configuration, StarTeam Server creates a number of folders for storing log files, attachments, archive files, and so on. This topic explains the location and purposes of the files and folders contained in the Native-II vault.

Native-II Vaults

All server configurations created using StarTeam Server use Native-II vaults to store new archive files. The Native-II vault improves StarTeam performance and allows much larger files to be stored than in earlier StarTeam releases. For StarTeam, your server configuration will have only Native-II archive files, and this means that backups can be done without shutting down the server.

Caution: You should never delete or modify repository files other than through StarTeam Server.

Understanding Repositories

Consider the following server configuration whose repository path starts with a drive letter (not shown) and ends with the folder name `MyConfiguration`. As shown in the figure below, the repository contains `Attachments`, `DefaultHive`, `HiveIndex`, `Notifications`, and `Trace` subfolders. The `DefaultHive` folder contains `Archives` and `Cache` subfolders.

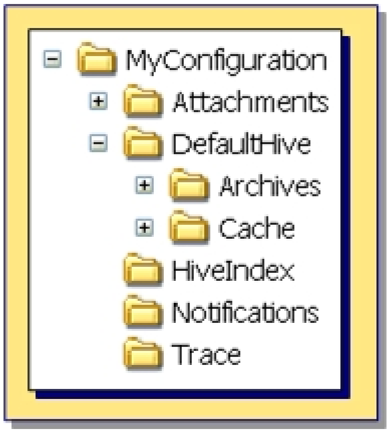

The name of the server configuration may also be `MyConfiguration`. The repository path is a general location for initial storage of a variety of objects, most of which can be moved to new locations later, independent of one another.

Log Files and Folders

The repository path folder, such as the `MyConfiguration` folder in the above example, becomes the home of the following related objects.

The server log files The server creates a new server log file each time you start the server configuration.

.dmp files The server creates .dmp files when you use server diagnostics to log errors and unexpected conditions it encounters. Usually, you have no .dmp files or trace files, discussed below as the contents the *Trace* subfolder, unless a technical support representative is working with you on a problem.

The Trace subfolder The Trace subfolder stores the files that are created when and if you use server diagnostics to trace server commands.

These objects do not have to remain in the repository path. You can change the path to all of the above by changing the **Log Path** using the Server Administration Tool.

Tip: These folders do not have to be included in a backup.

Native-II Vault Folders

For server configurations the repository path is also the initial home of several folders used by the Native-II vault to store archive files and information about them. The `DefaultHive` folder contains two subfolders, `Archives` and `Cache`. These folders are described below.

HiveIndex The `HiveIndex` folder stores the `hive-index.xml` file, which contains the properties for each hive used by the server configuration.

You can change the path to the `HiveIndex` folder by changing the repository path in the `starteam-server-configs.xml` file. You would make this change only when necessary, for example, because of a drive failure.

Tip: The `HiveIndex` folder must be included in a backup.

DefaultHive

If you accepted all the defaults when you created the server configuration or if you started an upgraded server configuration without first creating a hive, StarTeam Server automatically creates the folder `DefaultHive`. It is a subfolder of the repository path and is created when you start the server configuration for the first time.

Whether the initial hive is called `DefaultHive` or not, you will have at least one hive for each server configuration. You may have several hives. Each hive has an archive and cache path. An easy, but not mandatory, naming convention is using `DefaultHive`. The name of the hive becomes the name of a folder with two subfolders: `Archives` and `Cache`. However, you can place these paths anywhere. They do not need to be on the same drive or volume.

Archives subfolder

This folder stores file revisions in archive files, which may be compressed.

Cache subfolder

This folder stores uncompressed versions of archive files. It has two subfolders `Temp` and `Deltas`. `Temp` is used for files that are being added to StarTeam and for new file revisions that are being checked in. `Deltas` stores the differences between working files and tip revisions when a user asks that transmissions over slow connections be optimized —an option found in the client on the **File** tab of the **Personal Options** dialog.

Native-II Vaults

The Native-II vault improves StarTeam performance (as compared to the old vault structure referred to as Native-I) and allows you to store much larger files than in earlier releases of StarTeam. Additionally, server configurations using Native-II archive files enable you to perform backups without shutting down the server. StarTeam server configurations support Native-II vaults only.

Native-II Vault Performance

The sections below explain how StarTeam handles add, check-in, and check-out operations.

Add Operations

To add a file to the Native-II vault, StarTeam Server stores the revision in a temporary folder, computes the MD5 value of its contents, and checks how well it compresses. If the compression is 10% or greater, the Server moves the compressed version to the archive for the hive and its uncompressed version to the cache for the hive. If the revision does not compress well, the Server moves the uncompressed version to the archive for the hive.

StarTeam converts the MD5 value to a hex string and uses it as the name for the archive file. StarTeam uses the .gz extension when it compresses the file archive. If an archive file already exists with that name, StarTeam does not create a new archive file—although the StarTeam properties for that file are set to identify the hive in which the revision is stored, the use of compression, and the name of the archive file.

Check-in Operations

To check in a file revision to a Native-II vault, StarTeam Server stores the revision in a temporary folder in the next hive in the hive rotation. Then the server computes the MD5 value of its contents. If an archive file with the correct name already exists in the hive, StarTeam does not create a new archive file, although StarTeam updates the revision properties for the file. Otherwise, StarTeam creates a new file archive. Notice that no two files that are identical in content are ever stored in a given hive.

If the StarTeam file was initially identified as one that compresses well, StarTeam compresses the file revision and places it in the hive's archive with a .gz extension. Its uncompressed version is moved to the hive's cache. Otherwise, the uncompressed version is moved to the hive's archive.

Check-out Operations

To check out a file revision from a Native-II vault, StarTeam Server checks the hive ID of the revision and archive filename. Then the server retrieves the file revision from the specified hive's cache or archive, and it sends the archive file directly. These clients know how to decompress the archive file when necessary.

Archives and Cache Structure

Every archive path and cache path for a hive has the same structure. This structure is similar to that used by StarTeam clients to store file status records.

StarTeam organizes the files located in the Archives and Cache folders into subfolders. This makes browsing and managing the files easier. The name of the subfolders in which StarTeam stores a file revision is based on the initial characters in the name of the archive.

For example, suppose the contents of a file revision has an MD5 value of `01fc3c4ac5e0e92cc707f30fb73a0726`. Assuming the user specified an Archives path of `C:\DefaultHive\Archives`, the Archives path for this revision would be one of the following, depending on whether or not StarTeam compresses the archive file:

```
C:\DefaultHive\Archives\01\f\01fc3c4ac5e0e92cc707f30fb73a0726
```

```
C:\DefaultHive\Archives\01\f\01fc3c4ac5e0e92cc707f30fb73a0726.gz
```

Note: You must include the Archives path for each hive (for example `C:\DefaultHive\Archives`) in a backup.

Delta Storage

StarTeam uses deltas to optimize for slow connections. To use this feature, users set the personal option named **Optimize for slow connections** found in the client on the **File** tab of the **Personal Options** dialog. Then when a user checks out a new revision of a file that is already in his or her working folder, the server recognizes the revision number for the working file and sends only the difference between that revision and the revision that is being checked out.

StarTeam Server stores each delta for later use in the *Deltas* folder, a subfolder of the *Cache* folder found in each hive. The file containing the delta is given a name that combines the names of the two archive files used to generate the data. For example, if the file revision for the client on disk has an MD5 value of:

```
7f46c2bb9602fe972d952f4988ab85cd
```

and the requested revision has an MD5 value of:

```
7f46c2bb9602fe972d952f4982ab35aa
```

then the server generates a delta between these two revisions and names it:

```
7f46c2bb9602fe972d952f4988ab85cd.7f46c2bb9602fe972d952f4982ab35aa
```

Hives

A hive is a computer location where StarTeam Server stores archive files and a cache. These items are contained in the `Archives` and `Cache` folders. For example, if you created a server configuration named `MyConfiguration` and located it on the root of your `C:\` drive, by default, StarTeam Server generates a folder under `C:\MyConfiguration` named `DefaultHive` containing `Archives` and `Cache` subfolders. The `DefaultHive` folder and its subfolders represent the hive.

StarTeam Server Native-II vaults can have any number of hives, each of which has its own archives and cache. If one hive fills up, you can add another without having to change any data locations or move any archive files. Companies with large files or large numbers of files can start off with more than one hive in the first place. They can even put the archives and cache on different drives or volumes (this is recommended).

Native-II vaults store each file revision in its entirety (even though the archive file may be compressed). This means that the Native-II vault eventually takes more space; however, you can spread the revisions over many drives or volumes by the use of hives for storage. This flexibility in using storage space becomes a greater benefit over time as hives become full.

When a server configuration has multiple hives, the Server adds files to each hive in turn before reusing the archive path of the first hive. If you are running a StarTeam client against a StarTeam Server and if a StarTeam Server configuration has more than one hive, then the Server does a round-robin as it stores files but it checks first to see that no hive already has this file before the client attempts to stream the file to the server.

When you create a server configuration, it automatically has at least one hive (either the default or a custom hive). To increase the amount of available space for this server configuration, you can add one or more new hives with the **Hive Manager** dialog. When remotely accessing a server configuration, you can create hives while the server configuration is running, because the configuration already has an initial path, if only to a `DefaultHive` in the repository path.

You can also use the **Hive Manager** dialog to change an individual archive path and/or cache path for a hive. Such changes should be done only when that hive must be moved. For example, you might move a hive as a result of a drive failure. You would also need to copy the contents of the archive path for that hive to the new location.

We recommend that the `Archives` and `Cache` volumes from one server configuration should not be mapped to `Archives` and `Cache` volumes from any other server configuration. Otherwise, the threshold settings on the `Archives` folders will not be calculated as accurately. This is because the server checks the available disk space when starting a server configuration, and it caches the value returned. As files are added or removed, for cache cleanup, the server adjusts the available space and determines if the threshold has been exceeded. Because of performance concerns, the threshold value is cached instead of checked every time a file is added or removed. The threshold value is a guideline used by the server to determine when to no longer place files in a particular hive. It is not meant to be an absolute cutoff.

Also, if the server believes that the threshold may have been crossed based on the cached available space, it will do one more check and query the file system as it does at startup to make sure current available disk space is correct before pulling the trigger on the hive. Accordingly, since the threshold value is tracked on a per-server-configuration-basis, in order for threshold calculation to be as accurate as possible (thus reducing the number of times the file system is checked for available space), we recommend that each server configuration point its hives to independent volumes.

Creating New Hives

You can use the Hive Manager dialog for creating new hives to increase the amount of available space, or for viewing and updating the properties of an existing hive.

If accessing a remote server configuration or if a local server configuration has been added as a remote server, you can create new hives while that server configuration is running. If accessing a local server configuration locally, you must first shut down the server configuration before creating a new hive.

1. Open the Server Administration tool and select the server configuration from the **Server Pane**.

   If you are not logged on, the Server Administration tool requires you to do so before continuing.

   Note: If accessing a local server configuration locally, you must shut down the server configuration before proceeding to the next step.
2. Do one of the following:

   - Click the Hive Manager shortcut button in the shortcut pane.
   - Select **Tools** > **Administration** > **Hive Manager** from the main menu.

   The **Hive Manager** dialog opens.
3. Click **New** in the **Hive Manager** dialog box.

   This opens the **New Hive** dialog box.

Note: The location of the `hive-index.html` file, which contains the properties for each hive used by the server configuration, displays at the top of the dialog.

4. Type information about the new hive in the following fields:

   - **Name**: Unique name for the hive. *DefaultHive* is the default.
   - **Archive path**: Path to the *Archives* folder for the new hive. The default path is `<repository path>\DefaultHive\Archives`.
   - **Cache path**: Path to the *Cache* folder for the new hive. The default path is `<repository path>\DefaultHive\Cache`.
   - **Maximum cache size**: Maximum number of megabytes of hard disk space that the *Cache* can use. The default is 20% of the disk space available. In the Server Administration tool, you can calculate the correct default maximum size for the cache. However, if you are using the Server Administration tool and it is not running on the same computer as the Server, you cannot calculate the maximum size. In this situation, type 100MB, as a default size.
   - **Cache cleanup interval**: Seconds between cache cleanup/refresh operations. The default value is 600. The range is 60 (1 minute) to 3153600 (1 year).
   - **Storage limit threshold**: Percentage of total disk space allowed for hive. When this percentage has been reached, StarTeam does not add any more archives to the hive. The default is 95% of total disk space.

   Tip: You can use UNC paths for the Archives and Cache paths.

5. Select or clear the option to **Allow new archives**. The default is selected. If no other hives exist for the server configuration, this check box must be selected.

   Note: If you are adding a hive because the original hive was low on space, you should also use the **Hive Manager** dialog to display the properties of that hive and clear the **Allow new archives** check box. This action allows the original hive to remain a check-out location, but keeps it from acquiring any new files. Files that are added go to the new hive.

6. Fill the **Root Cache Agent archive path** text box if you are using Cache Agent, and the Root Cache Agent is not on the same computer as the Server. Provide the path to the cache from the point of view of the Cache Agent.

   For example, suppose you create a new hive whose archive path is `C:\ProdServer\Hives\NewHive\Archives`, but the Root Cache Agent runs on a computer that has `H:\` mapped to `C:\ProdServer\Hives` on the StarTeam Server computer. The Root Cache Agent would see the new hive archive path as `H:\NewHive\Archives`, so in this situation, you would type `H:\NewHive\Archives` in the **Root Cache Agent archive path** text box.

7. Click **OK** to confirm your choices. This action returns you to the **Hive Manager** dialog.
8. Click **OK** to return to the main window of the Server Administration tool.

   Note: If accessing a local server configuration locally, you can now restart the server configuration.

Customizing the Archives Path

Changing the *Archives* path for a hive is generally done because of a serious problem, such as a drive failure. It must also be done with caution, or the results can be unexpected.

You must restart the server configuration for the new *Archives* path to take effect. The Server Administration tool saves the new path to the `hive-index.xml` file immediately; however, the changes take effect only after you restart the server configuration.

If accessing a remote server configuration or if a local server configuration has been added as a remote server, you can update the Archives path while that server configuration is running. If accessing a local server configuration locally, you must first shut down the server configuration before updating the Archives path.

To change the Archives path for a server configuration that is shut down

1. Open the Server Administration tool and shut down the server configuration for which you want to modify the Archives path.
2. Copy the Archives folder to its new location.
3. Open the **Hive Manager** dialog box in the Server Administration tool by doing one of the following:
   - Click the Hive Manager shortcut button in the shortcut pane,
   - Select **Tools** > **Administration** > **Hive Manager** from the main menu.
4. Update the **Archive path** field pointing to your new *Archives* path location.
5. Click **OK** to confirm your choices. This action returns you to the **Hive Manager** dialog.
6. Click **OK** to return to the main window of the Server Administration tool.
7. Restart your server configuration.

Note: If you already have more than one hive for your server configuration, and you cannot quickly move the *Archives* folder to its new location, then you can disable any new archives from being added to the problematic Archives path by clearing the option to **Allow new archives** in the **Hive Properties** dialog. With this option cleared, StarTeam does not add any new archives to the designated *Archives* folder for the specified hive.

To change the Archives path for a server configuration that is running

1. Open the Server Administration tool, and select the running server configuration containing the Archives path that you wish to update.

   Note: This must be a remote server configuration or a local server configuration that has been added as a remote server. You cannot access hive properties for local server configurations running locally.
2. Open the **Hive Manager** dialog box in the Server Administration tool by doing one of the following:
   - Click the Hive Manager shortcut button in the shortcut pane.
   - Select **Tools** > **Administration** > **Hive Manager** from the main menu.
3. Select the applicable hive in the **Hive Manager** dialog box, and click **Properties**.

   This opens the **Hive Properties** dialog box.
4. Clear the **Allow new archives** check box in the **Hive Properties** dialog if at least one other hive exists for the server configuration.

   The files that are added or checked in will be sent to the other hive.
5. Restart the server configuration.
6. At an appropriate time, do the following:
   a. Shut down the server configuration.
   b. Copy the archive files to their new location.
   c. Change the **Archive path** field in the **Hive Properties** dialog to the new location, and check the option to **Allow new archives**.
   d. Restart the server configuration.

Viewing and Customizing Hive Properties

Sometimes you may want to view the properties for a specific hive or change its settings. For example, you may want to move its *Archives* or *Cache* folders to an alternate location. In that situation, you must use the **Hive Manager** dialog to display the properties for the hive, and then change them.

If accessing a remote server configuration or if a local server configuration has been added as a remote server, you can view and update hive properties while that server configuration is running. If accessing a local server configuration locally, you must first shut down the server configuration before viewing or updating hive properties.

1. Open the Server Administration tool and select the server configuration from the **Server Pane**.

   If you are not logged on, the Server Administration tool requires you to do so before continuing.

   Note: If accessing a local server configuration locally, you must shut down the server configuration before proceeding to the next step.
2. Do one of the following:
   - Click the Hive Manager shortcut button in the shortcut pane.
   - Select **Tools** > **Administration** > **Hive Manager** from the main menu.

   The **Hive Manager** dialog opens.
3. Select the applicable hive in the **Hive Manager** dialog, and click **Properties**. The **Hive Properties** dialog opens.
4. Review and, if desired, change the information in this dialog.

   Tip: With the exception of the **Name** field, you can edit all of the fields in the **Hive Properties** dialog. For the default and possible settings for these fields, refer to the link “Creating New Hives” at the bottom of this topic.
5. Click **OK** to confirm your choices when satisfied with your changes.

   This action returns you to the **Hive Manager** dialog box.
6. Click **OK** to return to the main window of the Server Administration tool.

   Note: If accessing a local server configuration locally, you can now restart the server configuration.

Migrating Servers

Moving Server Configurations Overview

Below is an overview for moving server configurations and some common assumptions that lead to errors when moving server configurations. You should backup each of the StarTeam components before attempting to migrate or move a server configuration:

- Database
- Repository
- `starteam-server-configs.xml` file

Overview for Moving a Server Configuration

The following provides a database-independent overview for moving a server configuration:

1. Shut down the server configuration. Because of Native-II vaults, you do not *have* to do this, but it is still a good idea.
2. Create a database backup.
3. Verify the location of the files that you need to move. These are the database backup, `starteam-server-configs.xml` file, repository including Native-II archives if not located within the repository.
4. Copy the files from the source to the target location.
5. Copy the **entry** for the source server configuration into the target `starteam-server-configs-xml` file. If this file does not exist on the target machine, then you can copy the entire file and delete any entries from the file for server configurations that do not exist on the target machine. If this file *does* exist on the target machine, then take care to copy only the section needed for the server configuration that you are moving to the target machine.
6. Correct the repository/log path in `starteam-server-configs.xml` for the specified server configuration.

7. Restore the database from backup.
8. Configure the database connection.
9. If needed, start the server configuration with the **Start with Override** option.

Incorrect Assumptions about Moving Server Configurations

The following are some incorrect assumptions about moving server configurations that lead to errors, so do not try any of the following methods:

- Just move the database and create a new server configuration pointing to it.
- The server configuration consists only of a database and repository. Note that `starteam-server-configs.xml` file is also required when you move a server configuration.
- Just use the **Migrate Database** toolbar button. Note that this option migrates database types only.

Note: If you need help migrating a server configuration, contact the Micro Focus SupportLine at *http://supportline.microfocus.com*.

Migrating Server Configurations

Use the Server Administration tool to migrate from any database to any other database supported by the StarTeam Server. For example, you can use the migrate feature to migrate an Oracle or PostgreSQL database to a Microsoft SQL Server database. The migrate operation adds information about the new server configuration to the `starteam-server-configs.xml` file.

Note: You can perform this operation only if the server configuration is not running.

Preparing for Database Migration

1. Prior to migrating to another database, ensure that you have allowed enough table space for the database. Twenty data files of 1 GB each is a good foundation. If the tablespace size is too small, the database continues to extend the tablespace, reducing the performance.
2. Create a backup of the database you plan to migrate. Also ensure that you have backups of the files and folders in the server configuration repository.
3. Ensure that the database to be migrated is still in good repair by doing one or both of the following:
    - Run Vault Verify.
    - Run any verification tools provided by your database vendor.
4. Do one of the following:
    - Manually create a database or schema user as the recipient of the migrated data. See the *StarTeam Installation Guide* for details. Make sure that you note the names provided for the server name, and the user name and password for the database or schema user. At a minimum, this user must have permission to create tables and stored procedures (if the database supports stored procedures).
    - Use the Server Administration tool to automatically create a database or schema user as the recipient of the migrated data.
5. Plan the database migration for a time at which it will inconvenience few users. A server configuration cannot be running while the database is being migrated. Advise team members ahead of time that you plan to make the transition, let them know the time at which it will take place, and request that they check in their files.

Migrating a Database

1. Open the Server Administration tool, and select the desired server configuration from the server pane.

    Note: The server configuration cannot be running during a database migration.

2. Click the **Migrate Database** button, or choose **Actions** > **Migrate** from the main menu. A message warns you that you cannot migrate a server configuration if the server is not registered.
3. If your server is registered, click **Yes**. The **Create a New Target Configuration for Migration** wizard opens.
4. In the first page of the wizard, **Select Target Configuration for Migration**, do the following:
   a) Type the name for the new server configuration in the **Target Configuration name** field.
   b) Click **Next**. The second page of the wizard, **Type New Configuration Data**, opens.
5. Indicate the type of database in the **Database type** list.
6. Do one of the following:
   - (Default and recommended action): Check **Create new StarTeam database**.
   - Clear the check box if you have already manually created a database or schema user for this migration.
7. Click **Next** when this information is complete.
8. From this point on, the dialog boxes are the same as those that display when you create a server configuration.

Completing the Migration

1. Disable the prior server configuration. This action prevents the server configuration from being started and accessed accidentally.
   a) In the Server Administration tool, select the prior server configuration.
   b) Click the **Disable Server** button or choose **Actions** > **Enable Server** . The status icon to the left of the server configuration name changes to a disabled icon.

   ⚠ **Caution:** Both the old and the new server configurations access the same vault, cache, and attachments folders. However, they do not access the same database. Continuing to use the prior server configuration will lead to vault verification errors and must be avoided.
2. Empty the `Cache` folder for the hive before starting the new server configuration. By default, the `Cache` folder is a child folder of the hive folder, under the repository root folder.
3. After verifying that the new configuration works correctly, delete the:
   - Prior server configuration
   - The database that it used
4. (Optional). The `Z99` table is a temporary table that records the progress of the database migration. If the migration process stops before completing, it uses the `Z99` table to determine the point at which it should resume the migration when you restart the process. If your migrate process did not complete properly, you can review the following columns to determine how far the migration process has progressed.
   - Column 1 contains the source table name.
   - Column 2 contains the ID of the last record copies.
   - Column 4 contains either a Y or N, indicating whether the table copy is complete.

Moving Server Configurations to a New Server

To prepare for migrating a server configuration

1. Shut down the server configuration.
2. Create a database backup.
3. Verify the location of the files that you need to move. These file are:
   - The database backup.
   - `starteam-server-configs.xml`.

- The repository including its Native-II archives if they are not located under the repository folder.

To migrate a server configuration

1. Copy the files from the source to the target location.
2. Open `starteam-server-configs-xml` in a text editor.
3. Copy the entire entry for the source server configuration into the target `starteam-server-configs-xml` file.

   Note: If `starteam-server-configs-xml` does not exist on the target machine, then you can copy the entire file and delete any entries from it for any server configurations that do not exist on the target machine. However, if this file does exist on the target machine, then take care to copy only the section needed for the server configuration that you are moving to the target machine.
4. Open the `starteam-server-configs-xml` file.
5. Type the correct values in the `RepositoryPath` and `LogPath` options for your migrated server configuration so that it points to the new migrated locations for the specified server configuration.

   For example, you would update the following values for these options:

   - `<option name="RepositoryPath" value="C:\Program Files\Micro Focus\StarTeam Server <version>\QA Team Repository\" />`
   - `<option name="LogPath" value="C:\Program Files\Micro Focus\StarTeam Server <version>\QA Team Repository\" />`

To complete a server configuration migration

1. Restore the database from backup.
2. Configure the database for the migrated server configuration.
3. Start the migrated server configuration. Depending on whether you have other server configurations running on the same machine, you may need to start the migrated server configuration on a different port. If you need to do this, do one of the following:

   - Click the **Start with Override** toolbar button; or
   - Choose **Actions** > **Start With Override** from the main menu.

Guidelines for Data Files and Translation Logs

Guidelines for Microsoft SQL Server/Microsoft SQL Server Express Data Files and Transaction Logs

Based on the number of users, we suggest the following guidelines for data files and transaction logs. Your needs may be different from those shown in the table below.

Number of Users	Number of Data Files	Size of Each Data File	Number of Log Files	Size of Each Log File
Up to 15	3	50MB	3	50MB
Between 15 and 50	3	300MB	3	300MB
Between 51 and 100	5	300MB	5	300MB
Between 101 and 300	7	500MB	5	500MB
>300	7	800MB	6	500MB

Note: The transaction log file sizes are relevant only if the Transaction log backup is performed regularly.

Transaction log backups are essential. After a transaction is backed up, Microsoft SQL Server and Microsoft SQL Server Express databases automatically truncate the inactive portion of the transaction log.

This inactive portion contains completed transactions and is no longer used during the recovery process. The basic advantage comes with the fact that Microsoft SQL Server reuses this truncated, inactive space in the transaction log instead of allowing the transaction log to continue to grow and use more space. This is a huge plus from a performance standpoint.

Allowing files to grow automatically can cause fragmentation of those files if a large number of files share the same disk. Therefore, it is recommended that files or file groups be created on as many different available local physical disks as possible. Place objects that compete heavily for space in different file groups.

Guidelines for Oracle Schema User Data Files

Based on the number of users, we suggest the following guidelines for data files. Your needs may be different from those shown in the table below.

Number of Users	Number of Data Files	Size of Each Data File
Up to 15	3	50MB
Between 15 and 50	3	300MB
Between 51 and 100	5	300MB
Between 101 and 300	7	500MB
>300	7	800MB

Splitting Server Configurations

Splitting server configurations is not generally recommended. However, it may be appropriate to split a server configuration if, for example, its size or number of active users has outgrown its hardware or OS platform, or if your company process dictates that data must be moved from production systems to archival storage.

Please visit the Micro Focus SupportLine Web site at *http://supportline.microfocus.com* regarding any performance or scalability concerns before making a decision to split your StarTeam server configuration. In many cases, problems can be resolved without splitting the server configuration.

Before splitting a server configuration, consider the following implications:

Irreversibility Once split server configurations have begun to evolve independently, there is no way to merge them back together.

Data Traceability Shares, links, and floating item configurations between the moved and unmoved projects will be lost.

Administration Initially, a new server configuration will have the same set of configuration settings, users, and groups as the original configuration. Going forward, you must manage each server configuration individually, as changes will no longer be propagated between them.

Licensing Please contact your sales representative to discuss potential licensing issues that may arise from splitting a server configuration. To ensure compliance with the license agreement, you should use licenses managed in a license server, preferably FLEXlm, rather than native licenses.

1. Copy the original server configuration (`Server 1`) to a separate machine.
2. Remove the unwanted projects from the original server configuration.
3. Remap the Microsoft SQL Server logins for the new server configuration (`Server 2`).
4. Change the server GUID on the new server configuration (`Server 2`).

5. Remove the unwanted projects from the new server configuration (`Server 2`).

   Note: In this example, we will assume that the original server configuration (`Server 1`) has three projects named `Project A`, `Project B`, and `Project C`. The plan is to split the server configuration so that `Project A` and `Project B` will remain on `Server 1`, and `Project C` will reside on the new server configuration (`Server 2`).

6. Copy the `Server 1` configuration to a separate machine:

   a) Make a full database and Vault backups of the `Server 1` configuration.

   b) Restore the database and Vault on a secondary system.

      Caution: The database backup must be restored as a different database. Do not reuse the database location, Microsoft SQL Server database user, or Oracle schema user of the original server configuration.

   Once you complete the copy process, you should have two identical copies of your original server configuration running on two sets of hardware (server and database).

7. Remove the unwanted projects from the `Server 1` configuration:

   a) Make full database and Vault backups of the `Server 1` configuration.

   b) Start the `Server 1` configuration.

   c) Use the Cross-Platform Client to connect to `Server 1` and delete `Project C`.

   d) Shut down the `Server 1` configuration.

   e) Run **Purge** on the `Server 1` configuration to physically remove the deleted data.

   f) Use the **Vault Verify** utility to verify the integrity of the configuration data.

   g) If necessary, make full database and Vault backups of the `Server 1` configuration.

   h) Restore the database and Vault backups of the `Server 1` configuration from step 1 on `Server 2`.

   i) If using a different Vault location, configure the files `hive-index.xml` and `starteam-server-configs.xml` to point to the new location.

8. Remap the Microsoft SQL Server logins for the `Server 2` configuration:

   a) Connect to the database using sa or windows authentication and change the database context to the restored database.

   b) Run command `sp_change_users_login 'REPORT'`. This command will print the orphaned user name.

      Note: The following steps assume that the orphaned user is `starteam`. Use the appropriate orphaned user as reported by the command `sp_change_users_login 'REPORT'`.

   c) Run the following commands in SQL Query Analyzer:

```
sp_addlogin starteam
EXEC sp_change_users_login 'Update_One', 'starteam', 'starteam'
```

   d) Copy the contents of the script `set-owner-to-dbo.sql` and run it against the database.

      Note: This script can be found in the `DBScripts` folder under the StarTeam Server installation location.

   e) As sa user, execute the script by running the command `exec change_db_ownership 'starteam'`.

      Note: Warnings generated from this command are safe to ignore.

   f) Go to SQL Enterprise Manager or Microsoft SQL Server Management Studio for Microsoft SQL Server and delete user `starteam` from the database. Select **Yes** to also delete the schema.

      Note: This action deletes the database user `starteam`, not the Microsoft SQL Server Login `starteam`. Deleting the schema deletes all the database objects owned by this database user, which is required in order to delete a database user. This step is essential because, while there can be many users with dbo privileges, there can be only one database owner. StarTeam Server must be run by the database owner.

   g) Run the command `sp_changedbowner starteam`.
   h) Log into the database as user starteam (the password is blank by default) and run the SQL statement `select * from s0`.
   i) Ensure that one row is returned.

9. Change the server GUID on the new server configuration (`Server 2`):
   a) In the file `starteam-server-configs.xml`, update the option `ServerGuid` with a different GUID value for the new configuration name. For example: `<option name="ServerGuid" value="n"/>`, where `n` is the new server GUID value.
   b) Depending on whether it's a Microsoft SQL Server or Oracle database, perform one of the following steps to update the Server Settings table with the new server GUID value.

      - For Microsoft SQL Server, open a database connection using Microsoft SQL Server Management Studio or Studio Express, change the database context to the new database, and run the SQL statement `update s0 set f3 = n`, where `n` is the new server GUID value.
      - For Oracle, open a database connection to the new schema using SQL*Worksheet/SQL*Plus and run the following SQL statement, where `n` is the new server GUID value.

      ```
      update  s0 set f3 = n;
      commit;
      ```

10. Remove the unwanted projects from the new server configuration (`Server 2`):
    a) Configure a database connection for the restored server configuration.
    b) Start the `Server 2` configuration.
    c) Use the Cross-Platform Client to connect to `Server 2` and delete `Project A` and `Project B`.
    d) Shut down the `Server 2` configuration.
    e) Run **Purge** on the `Server 2` configuration to physically remove the deleted data.
    f) Use the **Vault Verify** utility to verify the integrity of the configuration data.
    g) If necessary, make full database and Vault backups of the `Server 2` configuration.

Index