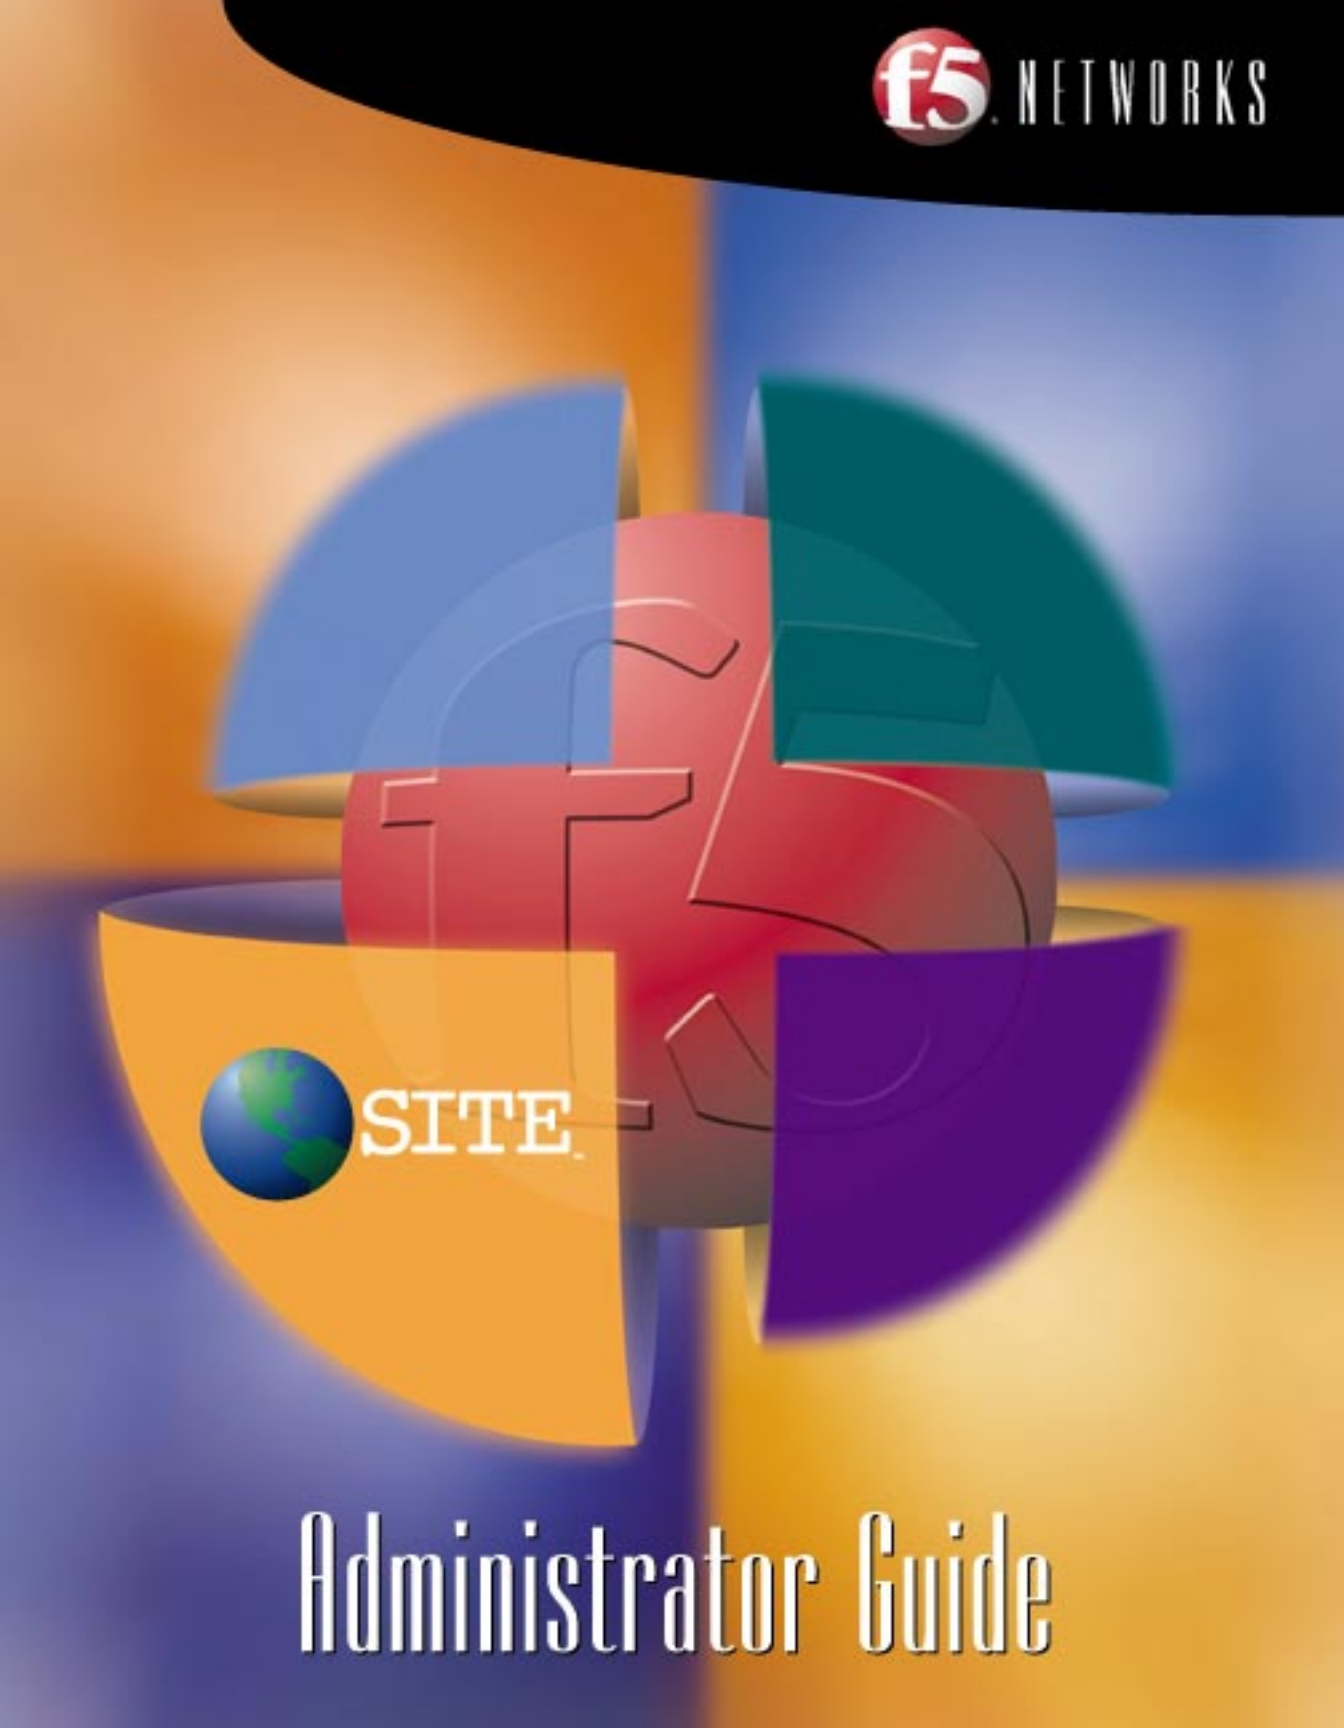
f5 NETWORKS
SITE
Administrator Guide

# Service and Support Information

## Product Version

This manual applies to version 1.0 of the global/SITE™ Controller.

## Obtaining Technical Support

| | |
|---|---|
| **Web** | tech.f5.com |
| **Phone** | (206) 505-0888 |
| **Fax** | (206) 505-0802 |
| **Email (support issues)** | support@f5.com |
| **Email (suggestions)** | feedback@f5.com |

## Contacting F5 Networks

| | |
|---|---|
| **Web** | www.f5.com |
| **Toll-free phone** | (888) 88BIG-IP |
| **Corporate phone** | (206) 505-0800 |
| **Fax** | (206) 505-0801 |
| **Email** | sales@f5.com |
| **Mailing Address** | 200 1st Avenue West<br>Suite 500<br>Seattle, Washington 98119 |

# Legal Notices

## Copyright



F100

## Trademarks

F5, BIG/ip, and 3DNS are registered trademarks of F5 Networks, Inc., and see/IT and global/SITE are trademarks of F5 Networks, Inc. Linux is a trademark of Linus Torvalds. Revision Control Engine is a registered trademark of DuraSoft GmbH.

Other product and company names are trademarks of their respective holders.

## Export Regulation Notice

The global/SITE™ Controller, if originally sold within the United States, may include cryptographic software. Under the Export Administration Act, the United States government may consider it a criminal offense to export this global/SITE™ Controller from the United States.

## FCC Compliance

This equipment has been tested and found to comply with the limits for a Class A digital device, pursuant to part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with this instruction manual, it may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference, in which case the user will be required to correct the interference at his own expense.

## Acknowledgments and Notices

3. All advertising materials mentioning features or use of this software must display the following acknowledgment:

"This product includes software developed by the Apache Group for use in the Apache HTTP server project (http://www.apache.org/)."

4. The names "Apache Server" and "Apache Group" must not be used to endorse or promote products derived from this software without prior written permission. For written permission, please contact apache@apache.org.

5. Products derived from this software may not be called "Apache" nor may "Apache" appear in their names without prior written permission of the Apache Group.

6. Redistributions of any form whatsoever must retain the following acknowledgment:

"This product includes software developed by the Apache Group for use in the Apache HTTP server project (http://www.apache.org/)."

THIS SOFTWARE IS PROVIDED BY THE APACHE GROUP ``AS IS'' AND ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE APACHE GROUP OR ITS CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

# F5 Networks Limited Warranty

This warranty will apply to any sale of goods or services or license of software (collectively, "Products") from F5 Networks, Inc. ("F5"). Any additional or different terms including terms in any purchase order or order confirmation will have no effect unless expressly agreed to in writing by F5. Any software provided to a Customer is subject to the terms of the End User License Agreement delivered with the Product.

## Limited Warranty

Software. F5 warrants that for a period of 90 days from the date of shipment: (a) the media on which the software is furnished will be free of defects in materials and workmanship under normal use; and (b) the software substantially conforms to its published specifications. Except for the foregoing, the software is provided AS IS.

In no event does F5 warrant that the Software is error free, that the Product will operate with any software or hardware other than that provided by F5 or specified in the documentation, or that the Product will satisfy Purchaser's own specific requirements.

Hardware. F5 warrants that the hardware component of any Product will, for a period of one year from the date of shipment from F5, be free from defects in material and workmanship under normal use.

Remedy. Purchaser's exclusive remedy and the entire liability of F5 under this limited warranty and any other guarantee made by F5 is, at F5's option, to repair or replace any Product or component that fails during the warranty period at no cost to Purchaser. Products returned to F5 must be pre-authorized by F5 with a Return Material Authorization (RMA) number marked on the outside of the package, and sent prepaid, insured, and packaged appropriately for safe shipment. The repaired or replaced item will be shipped to Purchaser, at F5's expense, no later than 7 days after receipt by F5. Title to any returned Products or components will transfer to F5

upon receipt. F5 will replace defective media or documentation or, at its option, undertake reasonable efforts to modify the software to correct any substantial non-conformance with the specifications.

Restrictions. The foregoing limited warranties extend only to the original Purchaser, and do not apply if a Product (a) has been altered, except by F5, (b) has not been installed, operated, repaired, or maintained in accordance with F5's instructions, (c) has been subjected to abnormal physical or electrical stress, misuse, negligence, or accident, or (d) has been operated outside of the environmental specifications for the Product. F5's limited software warranty does not apply to software corrections or upgrades.

Support, Upgrades. F5 provides software telephone support services at no charge for 90 days following the installation of any Product: Monday through Friday, from 6 a.m. to 6 p.m. Pacific time, excluding F5's holidays. Such support will consist of responding to trouble calls as reasonably required to make the Product perform as described in the Specifications. For advisory help requests, which are calls of a more consultative nature than a standard trouble call, F5 will provide up to two hours of telephone service at no charge. Additional service for advisory help requests may be purchased at F5 Networks' then-current standard service fee. During this initial 90 day period, Customer is entitled, at no charge, to updated versions of covered software such as bug fixes, and incremental enhancements as designated by minor revision increases (for example, BIG/ip V1.5 to BIG/ipV1.6). In addition, Customer will receive special pricing on upgraded versions of covered Products such as new clients, new modules, and major enhancements designated by major revision increases (for example, BIG/ip V1.x to BIG/ip V2.0). Customer may purchase a Maintenance Agreement for enhanced maintenance and support services.

**DISCLAIMER; LIMITATION OF REMEDY:** EXCEPT FOR THE WARRANTIES SPECIFICALLY DESCRIBED HEREIN, F5 DOES NOT MAKE ANY GUARANTEE OR WARRANTY, EXPRESS OR IMPLIED, WITH RESPECT TO PRODUCTS, SPECIFICATIONS, SUPPORT, SERVICE, OR ANYTHING ELSE. F5 HAS NOT AUTHORIZED ANYONE TO MAKE ANY REPRESENTATION OR WARRANTY OTHER THAN AS PROVIDED ABOVE. F5 DISCLAIMS ANY AND ALL

WARRANTIES AND GUARANTEES, EXPRESS, IMPLIED, OR OTHERWISE, ARISING WITH RESPECT TO THE PRODUCTS OR SERVICES DELIVERED HEREUNDER, INCLUDING BUT NOT LIMITED TO THE WARRANTY OF MERCHANTABILITY, THE WARRANTY OF FITNESS FOR A PARTICULAR PURPOSE, AND ANY WARRANTY OF NON-INFRINGEMENT OF THE INTELLECTUAL PROPERTY RIGHTS OF ANY THIRD PARTY. F5 WILL HAVE NO OBLIGATION OR LIABILITY, WHETHER ARISING IN CONTRACT (INCLUDING WARRANTY), TORT (INCLUDING ACTIVE, PASSIVE, OR IMPUTED NEGLIGENCE, STRICT LIABILITY, OR PRODUCT LIABILITY), OR OTHERWISE FOR ANY SPECIAL, INCIDENTAL, CONSEQUENTIAL, OR INDIRECT DAMAGES INCLUDING BUT NOT LIMITED TO LOSS OF USE, LOSS OF DATA, BUSINESS INTERRUPTION, LOSS OF REVENUE, LOSS OF BUSINESS, OR OTHER FINANCIAL LOSS ARISING OUT OF OR IN CONNECTION WITH ANY OF THE PRODUCTS OR OTHER GOODS OR SERVICES FURNISHED TO CUSTOMER BY F5, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

# End-user Software License

**IMPORTANT! READ BEFORE INSTALLING OR OPERATING THIS PRODUCT.**

**CAREFULLY READ THE TERMS AND CONDITIONS OF THIS LICENSE BEFORE INSTALLING OR OPERATING THIS PRODUCT: BY INSTALLING, OPERATING, OR KEEPING THIS PRODUCT FOR MORE THAN THIRTY DAYS AFTER DELIVERY, YOU INDICATE YOUR ACCEPTANCE OF THESE TERMS AND CONDITIONS. IF YOU DO NOT AGREE WITH THE TERMS AND CONDITIONS OF THIS AGREEMENT, PROMPTLY CONTACT F5 NETWORKS, INC. ("F5") TO ARRANGE FOR RETURN OF THE PRODUCT FOR A REFUND.**

1. Scope. This License applies to the software for the global/SITE Controller, whether such software is provided separately or as an integral part of a hardware product. As used herein, the term "Software" will refer to all such software, and the corrections, updates, new releases and new versions of such software. A product that consists of Software only will be referred to as a "Software Product" and a combination Software/Hardware product will be referred to as a "Combination Product." All Software is licensed, not sold, by F5. This License is a legal agreement between F5 and the single entity ("Licensee") that has acquired Software from F5 under applicable terms and conditions.

2. License Grant. Subject to the terms of this License, F5 grants to Licensee a non-exclusive, non-transferable license to use the Software in object code form solely on a single central processing unit owned or leased by Licensee. Other than as specifically described herein, no right or license is granted to Licensee to any of F5's trademarks, copyrights, or other intellectual property rights. Licensee may make one back-up copy of any Software Product, provided the back-up copy contains the same copyright and proprietary information notices as the original Software Product. Licensee is not authorized to copy the Software contained in

a Combination Product. The Software incorporates certain third party software which is used subject to licenses from the respective owners.

3. Restrictions. The Software, documentation, and the associated copyrights are owned by F5 or its licensors, and are protected by law and international treaties. Except as provided above, Licensee may not copy or reproduce the Software, and may not copy or translate the written materials without F5's prior, written consent. Licensee may not copy, modify, reverse compile, or reverse engineer the Software, or sell, sub-license, rent, or transfer the Software or any associated documentation to any third party.

4. Export Control. F5's standard Software incorporates cryptographic software. Licensee agrees to comply with the Export Administration Act, the Export Control Act, all regulations promulgated under such Acts, and all other laws and governmental regulations relating to the export of technical data, and equipment, and products produced therefrom, which are applicable to Licensee.

5. Limited Warranty.

   a) Warranty. F5 warrants that for a period of 90 days from the date of shipment: (i) the media on which the Software is furnished will be free of defects in materials and workmanship under normal use; and (ii) the Software substantially conforms to its published specifications. Except for the foregoing, the Software is provided AS IS. In no event does F5 warrant that the Software is error-free, that it will operate with any software or hardware other than that provided by F5 or specified in the documentation, or that the Software will satisfy Licensee's own specific requirements.

   b) Remedy. Licensee's exclusive remedy and the entire liability of F5 under this limited warranty and any other guarantee made by F5 is, at F5's option, to repair or replace any F5 product that fails during the warranty period at no cost to Licensee. Any products returned to F5 must be pre-authorized by F5 with a Return Material Authorization (RMA) number marked on the outside of

the package, and sent prepaid, insured, and packaged appropriately for safe shipment. The repaired or replaced item will be shipped to Licensee, at F5's expense, no later than 7 days after receipt by F5. Title to any returned product or components will transfer to F5 upon receipt. F5 will replace defective media or documentation or, at its option, undertake reasonable efforts to modify the Software to correct any substantial non-conformance with the specifications.

c) Restrictions. The foregoing limited warranties extend only to the original Licensee, and do not apply if a Software Product or Combination Product (i) has been altered, except by F5, (ii) has not been installed, operated, repaired, or maintained in accordance with F5's instructions, (iii) has been subjected to abnormal physical or electrical stress, misuse, negligence, or accident or (iv) has been operated outside of the environmental specifications for the product. F5's limited software warranty does not apply to software corrections or upgrades.

6. DISCLAIMER; LIMITATION OF REMEDY. EXCEPT FOR THE WARRANTIES SPECIFICALLY DESCRIBED HEREIN, F5 DOES NOT MAKE ANY GUARANTEE OR WARRANTY, EXPRESS OR IMPLIED, WITH RESPECT TO THE SOFTWARE, SPECIFICATIONS, SUPPORT, SERVICE OR ANYTHING ELSE. F5 HAS NOT AUTHORIZED ANYONE TO MAKE ANY REPRESENTATION OR WARRANTY OTHER THAN AS PROVIDED ABOVE. F5 DISCLAIMS ANY AND ALL WARRANTIES AND GUARANTEES, EXPRESS, IMPLIED OR OTHERWISE, ARISING WITH RESPECT TO THE SOFTWARE OR SERVICES DELIVERED HEREUNDER, INCLUDING BUT NOT LIMITED TO THE WARRANTY OF MERCHANTABILITY, THE WARRANTY OF FITNESS FOR A PARTICULAR PURPOSE, AND ANY WARRANTY OF NON-INFRINGEMENT OF THE INTELLECTUAL PROPERTY RIGHTS OF ANY THIRD PARTY. F5 WILL HAVE NO OBLIGATION OR LIABILITY, WHETHER ARISING IN

CONTRACT (INCLUDING WARRANTY), TORT (INCLUDING ACTIVE, PASSIVE, OR IMPUTED NEGLIGENCE, STRICT LIABILITY OR PRODUCT LIABILITY), OR OTHERWISE, FOR ANY SPECIAL, INCIDENTAL, CONSEQUENTIAL, OR INDIRECT DAMAGES INCLUDING BUT NOT LIMITED TO LOSS OF USE, LOSS OF DATA, BUSINESS INTERRUPTION, LOSS OF REVENUE, LOSS OF BUSINESS, OR OTHER FINANCIAL LOSS ARISING OUT OF OR IN CONNECTION WITH THE SOFTWARE OR OTHER GOODS OR SERVICES FURNISHED TO LICENSEE BY F5, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

7. Termination. This License is effective until terminated, and will automatically terminate if Licensee fails to comply with any of its provisions. Upon termination of this License, the Licensee will destroy the Software and documentation and all copies or portions thereof.

8. Miscellaneous. This Agreement will be governed by the laws of the State of Washington, USA without regard to its choice of law rules. The provisions of the U.N. Convention for the International Sale of Goods will not apply. Any provisions found to be unenforceable will not affect the enforceability of the other provisions contained herein, but will instead be replaced with a provision as similar in meaning to the original as possible. This Agreement constitutes the entire agreement between the parties with regard to its subject matter. No modification will be binding unless in writing and signed by the parties.

# Table of Contents

## *Chapter 1*



## *Chapter 2*

## *Appendix A*



## *Appendix B*


# 1

# Introduction to the global/SITE Controller

- **Welcome to the global/SITE Controller**
- **global/SITE Controller features**
- **Finding help and other technical support resources**

# Welcome to the global/SITE Controller

The global/SITE Controller is a network appliance that manages publishing, distribution, and synchronization of file-based content and applications to local and geographically distributed Internet sites. The global/SITE Controller uses a publishing metaphor, tracking your content as *publications*. When you log on to the global/SITE Controller, you can view the current state and status of your publications at a glance. You can see specific information about each publication, view or modify sections and versions, and update or publish any publication as easily as using the mouse to point and click.

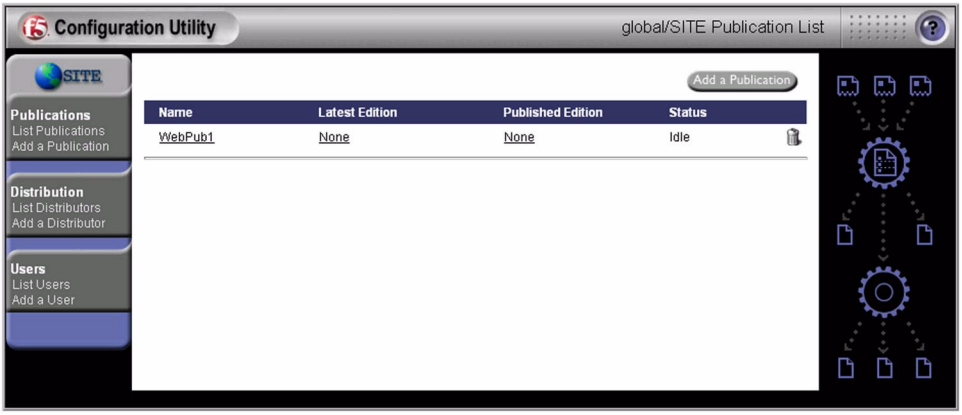


***Figure 1.1*** *The global/SITE Controller Publication List screen*

## Application tools

The global/SITE Controller supports the F5 Configuration utility, which is a browser-based application, and also supports an administrative command line utility. The browser application provides access to all global/SITE Controller features and functions. Figure 1.1 shows the Publication List screen, the application screen that you see when you first connect your browser

to the Configuration utility from the global/SITE Controller home page. Figure 1.2 shows the F5 Networks Configuration Utility home page for global/SITE Controller.

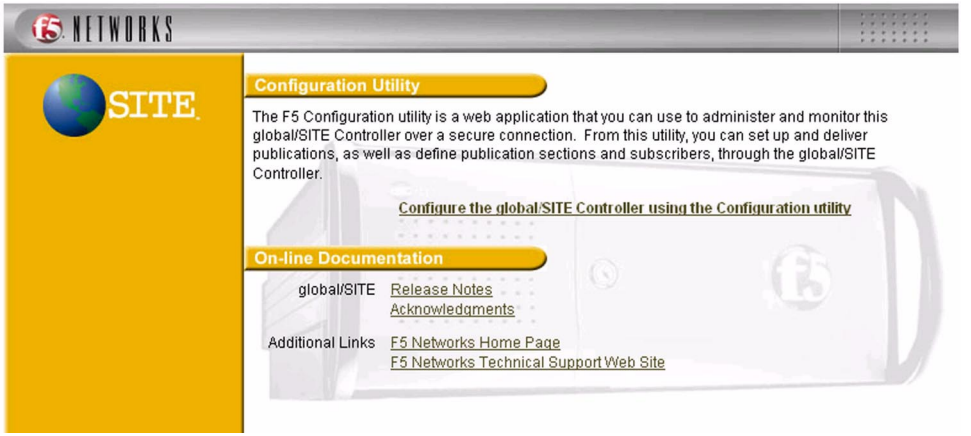


***Figure 1.2*** *F5 Networks Configuration Utility home page*

## The browser application

The Publication List screen (shown in Figure 1.1) gives you a quick overview of your publications, and status of each, and a quick way to delete any publication. In addition, it provides a link to the Publication Editions screen where you can view further details about each publication. Figure 1.3 shows the Publication Editions screen, which shows details for the editions of each publication, and provides access to the sections, section versions, subscribers, and publication options associated with each publication that you manage.

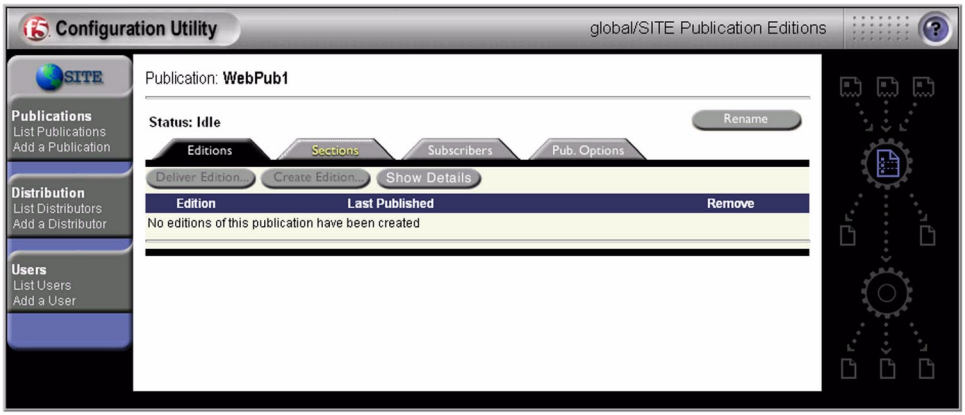


***Figure 1.3*** *The global/SITE Controller Publication Editions screen*

The Publication Editions screen, shown in Figure 1.3, is one of a series of publication screens that display the basic building blocks of your publications, provide status of each, and lead you through the publication process. There are tabs for each of the major pieces of a typical publication: Editions, Sections, Subscribers, and Publication Options. Each tab displays a screen that provides an overview of the element, and access to status and details for that element. You can typically create, modify, or delete each element, and add or remove that element from a publication. You also use the publication screens to start and monitor the publication process.

The User and Distribution screens assist you in managing the administrative functions of the global/SITE Controller. Use them for adding or removing user accounts, and viewing or adding global/SITE Controllers that function as distributors.

The ***navigation pane***, displayed in the left frame, provides links that you can use for administrative tasks, such as managing or adding new user accounts or distributors. Click **Publications**, **Distribution**, or **Users** to see summary help for that group. Click any of the other options on the navigation pane to access an application screen for use.

The global/SITE Controller screens have a standard application title bar across the top, which includes both the screen title and a help button. You can click the help button to open a page of help pertinent to the screen you are viewing when you click the button.

The global/SITE Controller screens also have a ***locator*** on the right side of the screen. This representative graphic provides a reminder of which part of the publication you are currently working with.

### The UNIX command line utility

The supplemental UNIX command line utility supports important maintenance commands that allow you to perform troubleshooting tasks. The command line utility also supports basic administrative commands, such as creating and updating the global/SITE Controller. For more information on the command line, refer to Appendix A and Appendix B.

## Specifications

The global/SITE Controller requires no additional hardware or software from the user. Simply plug it in, configure it, and it is ready to go with your existing server and publishing systems.

### Operating systems

The global/SITE Controller runs on the Redhat$_{TM}$ Linux version 5.2.1.

The global/SITE Controller supports a variety of source and target server operating systems:

- Windows NT®
- Windows® 95
- UNIX platforms
- Mac® OS

### Web browsers

You can connect to the global/SITE Controller using either of the following popular web browsers:

- Microsoft® Internet Explorer, version 4.0 and later
- Netscape® Navigator, version 4.0 and later

# global/SITE Controller features

The global/SITE Controller offers a variety of management and control features for your content synchronization and distribution needs.

## Application and content synchronization

The global/SITE Controller provides total control of all content publishing processes. The browser-based application makes it easy for you to pull data from multiple staging servers, and then synchronize final delivery across all servers simultaneously, at one or more sites, regardless of location.

The global/SITE Controller allows web publishers to intelligently deploy files and directories to web servers. Web publishers can define standard file distribution profiles, or set up one-time distribution rules. In addition, when doing updates between publisher and distributor controllers, the global/SITE Controller transmits only updates, reducing traffic on network resources.

## Version control

The global/SITE Controller is designed to track and save all changes while providing continual access to all prior versions. It synchronizes the project data accordingly. Your web publishers are no longer limited to publishing only the most recent content. The complete history of each project is now accessible.

## Site recovery and one-button rollback

The global/SITE version-control feature provides instantly available application and content rollback. This feature is critical for effective management of your applications and content at Internet speed. In case of problems with a site or with content, web publishers can immediately select any previous edition to be republished to production web servers and they can accomplish this with the click of a button.

## Content and administration security

Built-in security mechanisms throughout the global/SITE Controller protect Internet content and applications for transfer between the publisher and the distributor, then ensure proper delivery of content and applications to the production server. The browser-based administrative tool uses SSL to encrypt information passed between the browser and the global/SITE administrator web server.

## Visibility into publishing process

The design of the browser interface makes viewing the publication organization, and monitoring the publishing process easy. You start at the overview and work to more specific details of each publication. At any level, you are one click away from a different perspective of your publication.

With the Section Progress and Publish Progress screens, web publishers can monitor the process as they update sections, or deliver content to subscribers. The Section File Listing screen shows the name, size, and revision date for each file in a section version.

Throughout the browser interface, screens provide status of the publication, edition, or section. If you choose the optional edition delivery pause options on the Pub. Options screen, your publishing process will stop at selected times so that you can confirm status before continuing with the delivery of the edition.

## Adaptable open systems and web publishing support

The global/SITE Controller supports all source and target servers running UNIX, NT, and Macintosh operating systems. You manage all processes through a browser-based interface, or the command line utility.

The global/SITE Controller also works with all existing web publishing tools, including those with staging servers. Some of the file-based content and applications that the global/SITE Controller supports are:

- HTML
- Java Scripts
- Active Server Pages
- Adobe® PDF
- Microsoft Windows applications

The global/SITE Controller adapts to your organization's existing IT investment by synchronizing content, applications, and configurations across arrays of heterogeneous servers. You do not even need to contemplate the tedious and expensive task of installing additional server software.

## Secure, succinct network traffic

Pairing a global/SITE Controller distributor with a global/SITE Controller publisher provides increased scope for delivering content while maintaining both security and cost effectiveness. This publisher/distributor configuration allows you to publish securely behind firewalls, and take advantage of less expensive communication options.

When you use global/SITE publisher to global/SITE distributor publishing, this is what happens:

- The information is encrypted by the global/SITE publisher.
- The publisher sends to the distributor only the differences (file changes) between the current edition on the distributor and the new edition.

- The information is transferred across the WAN only once, from the publisher to the global/SITE distributor.
- The distributor then transfers the whole edition of changed files to all of the subscribers.

# Finding help and other technical support resources

In addition to this administrator guide, you can find help and technical documentation about the global/SITE Controller in the following locations.

## Release notes

You can access the release notes for the current version of the global/SITE Controller by clicking Release Notes on the global/SITE Controller home page. The release note contains the latest information for the current version, including a list of new features and enhancements, a list of fixes, and, in some cases, a list of known issues. You can always view the newest release notes on the F5 Technical Support page.

## Online help

You can find different types of help online in the following locations:

- **PDF version of the *global/SITE Controller Administrator Guide***
  A PDF version of this administrator guide is available from the global/SITE Controller technical support page at **http://tech.F5.com**. This site always has the most recent version of the Administrator Guide. Note that if you upgrade your global/SITE Controller at a later time, the upgrade may include an updated PDF version of the guide.

- **Getting started help**
  Each heading in the navigation pane, including **Publications**, **Distribution**, and **Users**, provides a basic summary of the related features.
- **Screen help**
  Each screen in the global/SITE Controller has an online help page that you can access by clicking the help button at the upper right of the screen. The help page describes the screen you are currently working in. The Tasks tab gives brief procedures for how to use the screen, and the Controls tab lists each control on the screen and provides an explanation of what it does.

## Technical support

The F5 Networks Technical Support web site, **http://tech.F5.com**, provides the latest technical notes, answers to frequently asked questions, and updates for administrator guides (in PDF format). To access this site, you need to obtain a customer ID and a password from your F5 service engineer.

You can also obtain technical support from the following sources:

| | |
|---|---|
| **Web** | tech.f5.com |
| **Phone** | (206) 505-0888 |
| **Fax** | (206) 505-0802 |
| **email (support issues)** | support@f5.com |
| **email (suggestions)** | feedback@f5.com |


# 2

# Publishing Documents

- **Why global/SITE Controller?**
- **Planning to configure the global/SITE Controller**
- **Using the publishing cycle**
- **Understanding the global/SITE phases**
- **Making ongoing configuration changes**

# Why global/SITE Controller?

The global/SITE Controller is used to manage the process of accumulating a logically consistent set of data, and replicating it to a set of systems in a controlled way.

The basic process involves:

- Capturing files from one or more sources and storing them on the global/SITE Controller
- Distributing those files to a set of servers

In order to define related sets of files, including where and how they are distributed to different servers, you establish relatively static relationships among a set of objects.

Those relationships are the key to achieving your goal in updating your file sets on the set of systems. The global/SITE Controller helps you manage the various steps in the process of moving the files through the pre-defined states of publishing, whether it is updating or delivering the files.

## Naming global/SITE objects

In order to discuss the basic global/SITE Controller processes, it is important to understand the terminology. These terms define the fundamental global/SITE elements, and correspond to the issues you need to consider when planning the organization of your publishing systems, deciding what you want to do in configuring the global/SITE Controller, and managing and delivering your content.

- A ***section*** contains information that tells the global/SITE Controller how to get a set of files that it collects from a source. Editions contain ***versions*** of sections, which reference the changed files.
- A ***publication*** is one or more sections to be published together to one or more ***subscribers***. An ***edition*** of a publication is a set of section/version pairs.

- A ***distributor*** is a remote global/SITE Controller that provides access to additional systems. A ***publisher*** can deliver publications directly to subscribers, or deliver to a distributor which then delivers the publication to the subscribers.

# Planning to configure the global/SITE Controller

Before you begin configuring the global/SITE Controller, it is a good idea to consider what you need, and plan what you are going to do. This helps save time since you can set things up initially in the way that will be most efficient for you.

You need to define:

1. Where you are going to look for new content
2. Where you are going to put it

You can use any file as content if you can specify a path to a server, and directory name, and method of access for the file. This is where you tell the global/SITE Controller to look for the content. The global/SITE Controller collects the content from the source, and stores it as a version of a section. The global/SITE Controller saves sections as a combination of the content to store, and the directions for collecting the content.

Once you have defined the content (that the global/SITE Controller stores as a version of a section), you can later distribute the content to one or more subscribers. The subscribers are where you are going to put the content. However, before you can publish the content, you need to know who your subscribers are, and provide a path for delivering the content to them. This process of taking content from one place that you define, and putting it somewhere else that you define, using a prescribed set of steps, is referred to as the publishing cycle.

In the terminology of the global/SITE Controller, a publication delivers one or multiple sections to one or more publication subscribers. The publication forms a single logical unit that the global/SITE Controller uses to move all of the individual parts through the publishing cycle in a controlled fashion.

The key to configuring a global/SITE Controller is:

❖ Determine what the logical publications are.

❖ Determine which sections each publication includes.

For example, one of your publications may be a quarterly report. This publication may include sections such as: Word from the CEO, Company Mission Statement and Goals, Investors, Customer Listing, Investments this year, Profits, Projections, and Summary. Another publication might be an monthly marketing update or a weekly recruiting listing. You may find that some sections, such as Company Mission and Goals, are useful in more than one publication.

Although you decide about the publication first, when you are using the global/SITE Controller, it makes sense to define the sections first, and then define the publication. Before proceeding further, you need to plan the overall organization of the publications, and the sections per publication.

## Defining Publications

Each publication defines two major relationships:

1. The sections that are published together
2. The set of targets or subscribers that receives the content from a given section

The publication, organizes the sections you have defined. It is this publication that is propagated, or delivered, to the target servers, or subscribers. For example, a publication can be defined as your web site (**www.yoursite.com**) which is composed of many separate sections (**www.yoursite.com/cgi-bin**, **www.yoursite.com/marketing**, and so on).

## Creating the publication

You need to assign a unique name to each publication. This name is an identifier, and shows up in almost all screens to identify publications in almost everything else you do from this point on.

For our example, assume you are creating a publication called **WebPub1**. You use the New Publication screen. (From the navigation pane on the left side of the global/SITE graphical interface, click **Add a Publication**.)

**◆ Note**

*In the global/SITE Controller, the screen name appears in the application title bar, across the top of the browser interface, and is not included on most screens shown as figures in this chapter. The screen names are shown in the figure **caption** for each screen.*

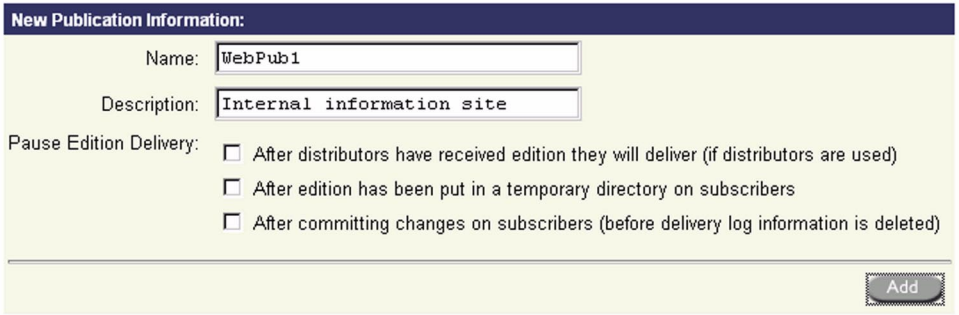


***Figure 2.1** The New Publication screen*

There is a field for the name of the publication, and an optional description. You can set or change the **Pause Edition Delivery** options now, or later. They specify whether the publication delivery process pauses automatically at certain points in the process.

### To add a publication

1. In the **Name** box, type the name of the publication you are adding.

2. In the **Description** box, type a description. This is optional, for your own use.
3. If you want the publication to pause during the publishing process, under **Pause Edition Delivery**, click any check box that applies. You can check one, two, or all three.

   When you publish this publication, the process will pause as you have directed it to.

   Note that you do not have to decide on the publication options now. You can set them at a later time with the Pub. Options screen. For more information about delivery pauses, and other publication options, refer to the help on the Publication Options screen.
4. Click the **Add** button to complete the addition. The global/SITE Publication List screen displays, with your new publication listed at the bottom.

   Once you have added the publication, your newly created publication appears in the list on the Publications List screen.

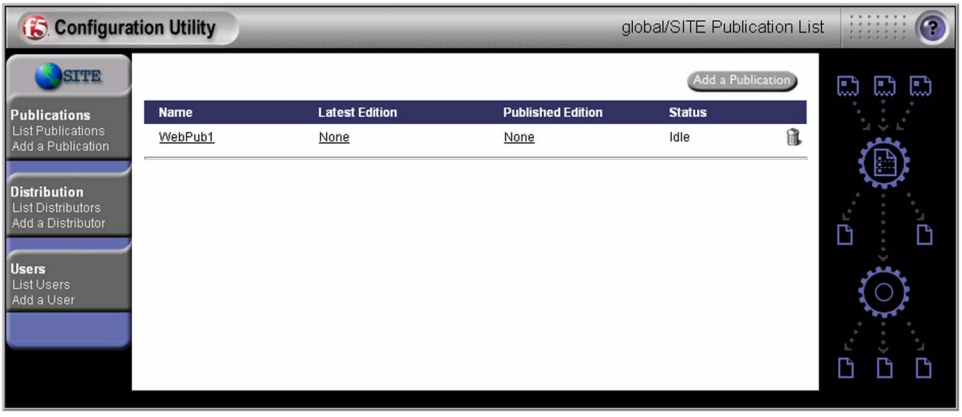


***Figure 2.2*** *The Publications List screen.*

5. Click the publication name (**WebPub1** in this example) in this list to view the Publication Editions screen for details of your publication.

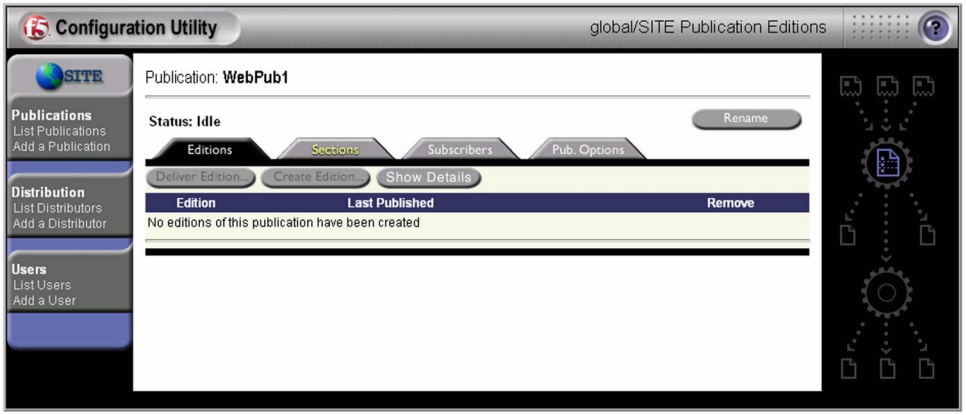


***Figure 2.3*** *The Publication Editions screen*

## Using the Publication screens

The Publication screens are the center of operations for nearly all configuration activities and content delivery for a publication. The Publication Editions screen is the first you will see. It has four major tabs that correspond to the four main areas of information: Editions, Sections, Subscribers, and Pub. Options. Click any tab to display a detail screen of the same name for dealing with that aspect of your publication. Each detail screen shows the same four tabs. You can click any one, to view a different area of your publication.

- The Publication Editions screen shows details of your editions and is used to deliver content from sections to subscribers.
- The Publication Sections screen shows section details for a specific publication, and is used to manage sections.
- The Publication Subscribers screen lists the subscribers for the publication, and is used to define where the section content will be delivered.

- The Publication Options screen is where you set options for the delivery process, including edition delivery pauses.

Before delivering content, you need to inform the global/SITE Controller where to get it. You do that by clicking the Sections tab, and using the Publication Sections screen.


***Figure 2.4*** *The Publication Sections screen*

## Defining sections

A global/SITE Controller section includes both data and a path to it.

Before you begin to set up the global/SITE Controller, you should answer some basic questions. Start by working backwards:

1. You have a server (subscriber); what data needs to be managed?
2. How many sections do you need?

   Remember that different roots (**/path1/part1**, **/path2/part2**) must have different sections.

   A section contains all the files and directories at a given path. A section can stop at any specified subdirectories, known as ***exceptions***. These exceptions exclude the listed directories or files from a particular section.
3. What will you name each section?

Remember that in order to create a section, you must have a unique name for it. This name is unique among all sections, within all publications.

4. What is the path to each section?

   You must be able to identify the specific directory path on the server that is the source of the section's content. Remember that, in all cases, the directory path for all section sources must already exist on the specified servers.

5. Do you have authority to access the information?
   You need to provide authorization access by specifying the user ID and password for access to the server where the content is stored.

## Defining a simple section

The first step, once you have a publication, is creating sections. Updating these sections is one of the most frequent tasks that you will perform. To create a section, you:

- Give the section a name.
- Specify the server where it can be found.
- Provide access settings; the user ID and password to access the server.
- Define the location of the content with a path.

Once you have created the section, you have several options, including testing the connection, creating exceptions to the section, or creating an initial version of the section on the global/SITE Controller. The following text in this section of the guide introduces all of these tasks.

### Example

This example shows how to create a new section called **WebContent1**, which gets its content from the directory **/home/webcontent** on the system **python** using the FTP account **website**.

1. Start by selecting the Sections tab for your new publication. The Publication Sections screen displays.


***Figure 2.5*** *The Publication Sections screen*

2. Since there are no sections yet, you need to create one: click the **Create New Section** button.
   The Create a New Section screen displays.


***Figure 2.6*** *The Create a New Section screen*

### To create a section

1. In the **Section Name** box, type a unique name for the section that you want to add. The name cannot contain double quotes, slashes, or angle brackets. You cannot change the name of a section later. For our example, type **WebContent1**.
2. In the **Staging Serve**r box, type the name of the server (node name or IP address) that is the source of the content for this section that you want to add. This is also referred to as the source or content server.
3. In the **Description** box, type a description for the section. This is optional, to provide mnemonic assistance.
4. To include this section in editions of the publication, leave the **Enabled** check box checked (the default). If you want to create the section now, but disable it for the publication, clear the check box. You can change this later.
5. Leave the **Transfer Method** set at FTP.
6. In the **User ID** box, type the User ID (FTP user name) you use to access the server you specified in step 2.
7. In the **Password** box, type the user password (FTP password) that matches the User ID you specified in step 6, and allows access to the server.
8. In the **Path** box, type the specific directory path for this section on this staging server. The section includes all the files and directories at the given path. If you want to exclude a subdirectory, you must specify it as an exception, using the Section Details screen.
9. To complete the addition of this section, and add it to the publication, click **Create**.

   If there are no errors, the Section Detail screen displays, showing the source detail for this section.

## Testing your section

Once the section is created, you may want to try a test connection on the newly defined section to confirm whether basic FTP access can be established. Test Connection verifies that the user ID, password, and path information are valid, by connecting to the specified system.

### To test your section connection

1. On the Section Detail screen, click the **Test Connection** button.
   The Connection Test screen displays.

The Connection Test screen displays the test results generated by clicking the **Test Connection** button. The Test Results table lists the name for the source server and the path for the section, and the success or failure of the test connection. The left side shows what is being tested, the right side shows the result.


***Figure 2.7*** *The Connection Test screen*

## Creating a section version

Now that you have created a section, you may also want to try creating an initial version of the section on the global/SITE Controller.

### To create a version

On the Section Detail screen:

1. Click the **Get New Version** button.

The global/SITE Controller goes through the process of gathering data from the specified location and storing an initial version of it.

For more information, refer to the later section, ***Getting new content***.

## Modifying your section configuration, later

Once you have created a section, you can return to the Section Detail screen later to change parameters if necessary.

### To change your section parameters

1. On the Publication Editions screen, click the Sections tab. The Publication Sections screen displays.

   The Publication Sections screen lists all sections for a publication.

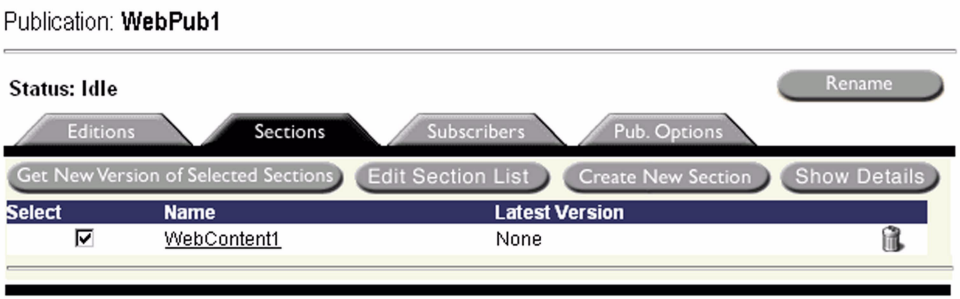


***Figure 2.8*** *The Publication Sections screen with one section listed*

2. Then click the name of the section you want to work on.

   This displays the Section Details screen where you can make changes.

***Figure 2.9*** *The Section Details screen, Source tab*

3. Press the **Update** button to save any changes you have made to the settings on this screen.
   Use the **Revert** button to ignore any changes made since the screen was brought up.

## Defining exceptions

Once you have created a section, you may want to create an ***exception*** to the section. Exceptions instruct the global/SITE Controller, when the specified path is encountered, to disregard any subdirectories in that directory. An exception is a way of excluding a subdirectory from being published as part of a section.

To create an exception path, you use the Section Detail screen after the section is created. You can use the **New Exception** box on the Section Detail screen to enter a directory path that will be treated as an exception.

### Example

In our example, the source directory **/home/webcontent** contains a subdirectory **/home/webcontent/images** that we want to manage with a separate section.

Assuming you have already created your initial section, start at the Section Detail screen for it (in our example, it is **WebContent1**).

***Figure 2.10*** *The Section Detail screen, add a new exception*

### To create an exception

Start at the Section Detail screen.

1. In the New Exception box, type the name of the directory you want to exclude. For this example, it is **/home/webcontent/images**.
2. Click the **Update** button.
3. Finally, create a new section, such as **WebImages**, which specifies **/home/webcontent/images** as its source.

**◆ Tip**

*In all cases, the directory path for all section sources and subscribers must already exist on the specified servers. The global/SITE Controller does not create them. It can create any subdirectories, but it cannot create the initial root path.*

### A little more about exceptions

You may have questions about directory paths when you are setting up multiple sections that use exception paths. The typical situation is that one section's exception is another section's root. In that context, it may seem like the global/SITE Controller should create the exception path, but it does not. That path must already exist, or you must first create it manually on the subscriber.

**Example**

Say that **section1** goes to **/root/section1**. It seems reasonable that **/root/section1** must already exist. You cannot assume you have permission to make it, especially on NT servers where it might be virtual.

## Loading initial content into sections

At this point, you can update the content for a section from its staging source, or you can create an edition.

### To update the content of a section

1. Start at the Publication Sections screen, where you see your sections listed.

***Figure 2.11*** *Publication Sections screen, showing two sections*

2. Click the **Get New Version of Selected Sections** button to begin the process of gathering new files from those sections.
   The Section Progress Display screen displays, where you can monitor the progress of that process.

This is discussed in more detail in the section ***Displaying status***.

Remember, it is still necessary to specify the destination for section in order to actually publish any content.

***Figure 2.12*** *The Section Progress Display screen*

## Specifying the subscribers for section content

Now we reach the heart of the matter: configuring subscribers. In the example, you have defined your web site as a group of sections, and created a single publication to conveniently manage them as the logical group they are.

You also have one or more web servers that need to receive that content. You need to specify the exact relationship between each section and each web server. Again, the publication provides that connection.

A ***subscriber*** is the server that receives your publication. To the global/SITE Controller, a subscriber describes a single system that is the target for delivering the contents of one or more of the sections belonging to a publication.

## Specifying a subscriber

1. From any of the Publication detail screens, click the Subscribers tab.
   The Publication Subscribers screen displays. In our example, there are no subscribers.

***Figure 2.13*** *The Publication Subscribers screen*

2. Click the **Add a Subscriber** button.
   The Add Subscriber screen displays.

***Figure 2.14*** *The Add Subscriber screen*

### To add a subscriber

Using the Add Subscriber screen:

1. In the **Subscriber Name** box, type the name of the subscriber that you want to add. This information is optional; the data you enter is displayed in various screens that list subscribers for the publication. If it is not filled in, the server name/user ID pair is used instead.
2. In the **Server** box, type the name of the server for the subscriber. This is the node name (or IP address) for the server.
3. If you will publish to this subscriber directly, leave the **global/SITE Distributor** box set at **None**, which is the default. If you intend to publish content to this subscriber through a secondary global/SITE Controller, or distributor,

select that distributor from the **global/SITE Distributor** box. Note that to select a distributor, you must have already added it. For more information on setting up a global/SITE distributor, see the following section, ***Defining a distributor***.

4. Leave the **Transfer Method** set at FTP, the default.
5. In the **User ID** box, type the user ID that you use to access the server specified in step 2.
6. In the **Password** box, type the user password that matches the User ID specified in step 5, and that allows access to the server in step 2.
7. Below the password is a list of sections associated with this publication. A path must be specified for those included. In our example, we have created two sections, and they are both listed.

   **Under Include Section**

   - Leave the check box checked (the default), for each section that you want included in the edition published to this subscriber.
   - In the **Destination Path** box, type in a specific destination path for each section on the subscriber.
   - For any sections that you do not want to go to this subscriber, clear the check box.

8. When you have specified whether each section will be included for this subscriber, click **Add**. This completes the addition of this subscriber, and adds it to the publication.

Once the subscriber is created, you can use the **Test Connection** button to verify that the global/SITE Controller can successfully log into the specified server using the user account, and locate the specified path(s). This is a good thing to do at this point, and at any time that you make changes to the subscriber definition.

## Defining a distributor

If you will be using a secondary global/SITE Controller to distribute to remote subscribers, you will need to set up distributors. You use the Add a Distributor screen to set up and add a new global/SITE distributor. (From the navigation pane, click **Add a Distributor**.) Our sample shows the information that you will add.

***Figure 2.15*** *The Add a Distributor screen, with information entered*

### To add a distributor

1. In the **Host Name** box, type the unique identifier host name or IP address of the global/SITE Controller that you want to function as a distributor.
2. In the **Description** box, type in an optional description. This is makes it easy for you to identify the distributor.
3. In the **global/SITE Identifier** box, type the unique, pre-assigned identifier for the global/SITE Controller that you are adding as a distributor.
4. In the **Key Phrase** box, type in the key phrase that is shared by the pair of publisher and distributor. Both the publisher and the distributor must have the same Key Phrase.
5. Click **Add** to complete the addition.
   The global/SITE Distribution List screen displays, with the new distributor added at the bottom of the list.

***Figure 2.16*** *The List Distributors screen shows the new distributor*

Once you have your distributors set up, you can easily specify a distributor for any subscriber that you add. Use the Add Subscriber screen.

***Figure 2.17*** *The Add Subscriber screen, adding a distributor*

### To add a distributor with a subscriber

1. Add the subscriber as you did earlier in ***To add a subscriber***.
2. At the **global/SITE Distributor** box, select the distributor that you want to use for this subscriber.
   Simply click the down arrow at the right of the box, and choose the subscriber you want from the available subscribers listed.
3. Complete the rest of the screen, and click **Add**. This completes the addition of this subscriber, with a distributor specified.

Once you have added the subscriber and distributor, you may want to try a test connection to verify that the global/SITE Controller can successfully log into the remote server. Test Connection verifies that FTP access can be established with the distributor using the user ID, password, and path information that you have provided.

### To test your remote connection

1. On the Subscriber Detail screen, click the **Test Connection** button.
   The Connection Test screen displays.

***Figure 2.18*** *The Connection Test screen for the distributor*

The Connection Test screen displays the test results generated by clicking the **Test Connection** button. The Test Results table lists the paths for the subscriber's website and the source server and the success or failure of the test connection. The left side shows what is being tested, the right side shows the result.

# Using the publishing cycle

So far, we have discussed how content is organized into sections, how sections are organized into publications, and where those publications go. Now we need to look at the process of capturing a consistent set of content to be published to one or more servers (subscribers).

Within the publication cycle, there are two independent activities:,

1. Updating the contents of sections from their staging sources
2. Publishing the contents of some or all the sections for a publication

These two activities are typically executed sequentially, but they need not be.

## Getting new content

To gather new content, the global/SITE Controller makes a copy of the content on the specified source server. The copy is stored as a version of the section. For ease in tracking, versions are numbered, starting with **1** the first time a new section is updated.

There are four ways to update a section's content at any time; both display the Section Progress Display screen.

- The first is from the Publication Sections screen, by clicking the **Get New Version of Selected Sections** button.
- The second is from the Section Detail screen, where you can click the **Get New Version** button.

- The third is to click **Create Edition** from the Publication Editions screen, and select the **Get new version** option.
- The fourth is to click the **Deliver Edition** button from the Publication Editions screen, and select the **Create new Edition** option.

***Figure 2.19*** *The Section Progress Display screen*

In all of these cases, the global/SITE Controller examines the directory tree specified by the combination of its staging server name, FTP user ID, and path. The global/SITE process ignores any exception directories on the path, and compares the content to the highest version number currently in the global/SITE Controller's versioned section. This process gathers any new files, or any files with different (not necessarily newer, just different) modified times, or a different size, and stores them in the global/SITE Controller. The number of new or changed files is what is displayed on the Update Progress screen.

If the global/SITE Controller determines that the set of files is identical to the latest version in the global/SITE Controller's versioned section, then no new version will be created.

The Section History tab on the Section Detail screen lists all versions currently stored in the section.

***Figure 2.20*** *The Section Detail screen, Section History tab*

In the example, we have already stored an initial version, **1**. You can view the contents of that version by clicking the number under the **Version** heading.

***Figure 2.21*** *The Section File Listing screen*

Keep in mind that multiple publications can share a single section. For instance, the **WebImages** section used in the preceding example was intended for use by numerous publications.

Creating the very first version of a section may take some time, as all of the data from the source must be copied to the global/SITE Controller and converted into its stored form in the section. However, subsequent updates should be faster, as they first compare the date and time of all the files on the staging source with those stored in the latest version in the section. Only new files, and those that have different timestamps or sizes, are copied into the section.

A final way to update section content is as the first part of the edition delivery process. (That is discussed later in this chapter.)

The next section reviews the publication process step-by-step so that you can better understand how the global/SITE Controller works.

**◆ Tip**

*An important reason not to update section content automatically as part of the content delivery process is so that the content being delivered is not inadvertently different from what is expected. When you update sections prior to delivery, the exact contents of that version are frozen, and remain independent of any possible changes made to the files on the staging server for that section.*

## Publishing new content

The publication process includes a number of distinct phases. The overall flow of the content distribution phase of the publishing cycle can be seen like this:

- **Create the edition**
  Define an edition and establish the publication settings that control this specific distribution process.
- **Deliver the edition**
  This runs the global/SITE Controller through a series of phases that result in the specified edition being published to all subscribers.

There are two methods of creating a new edition. You can create a new edition automatically as part of a publication delivery, after you first check all section staging servers for new content.

Alternately, you can perform the section update and edition creation process separately. This allows you to have more precise control over exactly what content is delivered.

## Create the edition

The preparation phase involves defining exactly which version of each section in the publication will be published. This set of section versions defines an edition.

Once an edition has been established, the global/SITE Controller creates a snapshot of all the parameters necessary to complete the rest of the process.

### To manually create a new edition

1. From the Publication Editions screen, click the **Create New Edition** button.
   The Create New Editions screen displays.

***Figure 2.22*** *The Create New Editions screen*

You can use the **remove version** option from the **Version** box to have the publication delivery process remove from the subscribers all the files from the last published edition.

The Create New Editions screen displays all the sections in the publication, with a box listing the existing version. You can click the arrow at the right of the box to view additional versions, or to create new versions, or remove existing versions. The displayed version is always the latest

existing version (with the highest version number). You can accept the default version of all sections by clicking the **Create Edition** button.

2. Select the version you want for each section, and click the **Create Edition** button.

   The Publication Editions screen displays, where you see the newly created edition added to the list. The new edition in turn lists each section it contains, and which version of that section will be delivered.

The global/SITE Controller automatically assigns both version and edition number, and simply increments each to the next higher value whenever a new one is created. Most listings also display the date and time they were created.

From the Publication Editions screen, you can click the section name in the edition listing to view the Section Detail screen for that section. Click the version number and date for a section to view the Section File Listing screen for that version of that section.

## Displaying status

At this point, we are ready to proceed with the publishing process. Some of the publishing phases can be lengthy, and are therefore designed as asynchronous processes.

In order that you can follow what is going on, the overall publication status is always displayed wherever the publication is listed. On the Publication Detail screen, for example, you can see that it is currently **Idle**. You can also see that in the overall publication list screen. Click the publication status in either of those screens to display the delivery details screen.

***Figure 2.23*** *The Publication Editions screen showing* ***Idle*** *status*

During the **Idle** state, the screen refreshes itself infrequently. Once the delivery process has started, however, it automatically updates more frequently, and more detailed delivery status is available.

Of course, you can use the **Refresh** or **Reload** button on your browser at any time to update it on demand.

1. Click the **Deliver Edition** button.
   The Deliver Edition screen displays.

***Figure 2.24*** *The Deliver Edition screen*

The Deliver Edition screen shows that neither path has anything published to it yet. The **Edition** field defaults to the latest edition, in this case, **edition 1**.

2. Click **Deliver Edition**. This starts the delivery process. The Publish Progress Display screen displays.

***Figure 2.25*** *Publication Progress Display, status:* ***WANSync***

On the Publish Progress Display screen you can observe the overall publication status change through multiple states, reflecting the individual subscriber path states in the detail list.

Note that Figure 2.25 shows details from an early phase in the process, and the details are about the sections. Figure 2.26 shows details from a later phase of the process, and those details are about the subscriber path.

***Figure 2.26*** *Publication Progress Display; Status:* ***CommitPhaseDone***

## The Cancel button

Sometimes either the section update process (getting new content) or publication delivery process need to be stopped before they complete. Perhaps a section source or subscriber becomes unavailable during the process, and there is no point in proceeding further until the problem is corrected.

Whenever the publication state is not **Idle**, three control buttons are displayed on the detail status screens.

The **Cancel** button  stops all processing activity, and places the publication into a **Cancelled** state. If the publication has not yet reached the **CopyPhase**, this resets the publication to **Idle**.

Canceling the copy phase puts the publication into a **CopyPhaseCancelled** state which allows two options:

- Proceed with any subscribers that have successfully completed the copy phase for all paths.

- Reset the entire publication to **Idle** without completing any subscribers.

Any subscribers that are not in the **gsCopyDone** state will not be committed. ***Committed*** means that the files are moved from the temporary subdirectory to a permanent directory on the subscriber.

During transitions of the publishing cycle, you may need to know what is going on, what the status of the publication is, or how to tell if something has stalled or if it is still progressing. Table 2.1, *The global/SITE Controller phases of publication*, lists the various activity or transition states, gives information on what should be happening, and includes suggestions on how to proceed if it is stalled, and when to start over.

# Understanding the global/SITE phases

During various processing activities, the overall state of the global/SITE controller is displayed at two levels of detail:

- The overall state of the publication, (which is often the synthesis of the section and subscriber path states).
- The detailed section or subscriber path state, whichever object is being processed

The overall state determines what other actions the user can perform, including the ability to proceed or stop when error conditions arise.

During a publication delivery with no errors, and no pauses specified on the Publication Options screen, the normal sequence of states follows this progression:

- **Idle**
  The publication is in a stable state, either ready for updating or publishing.
- **UpdatingSections**
  When this has been specified as part of the delivery. If an edition was specified with pre-existing section versions, this is skipped.

- **PublicationPreparation**
  Brief state which prepares some data structures for the rest of the process.
- **WANSync**
  A very brief state, unless distributors are involved. The process may be paused here. If there are distributors, they are updated with changes prior to publishing. If there are no distributors, there is nothing to change, and the state is extremely brief.
- **CreateChangeTrees**
  Files are prepared for quick and efficient transmittal to subscribers.
- **CopyPhase**
  Changed files are copied to temporary sub-directories for the subscribers. The process may be paused here after the copy phase.
- **CommitPhase**
  Changed files are moved from the temporary sub-directories to their final location at the path you specified, and obsolete files are removed from subscribers.

You can stop the entire process at any point by clicking the **Cancel Delivery Control** button. If you do this during the **CopyPhase**, you can then complete the publication cycle--for any subscribers that have already had all section files copied to them--by clicking the **Continue** or **Complete Delivery Control** button.

A **Cancel** used in any other phase just allows the publication to be reset to the **Idle** phase where it can be restarted from the beginning.

Table 2.1, *The global/SITE Controller phases of publication,* which follows, describes what is happening to the process during several publication state and detail state combinations.

Keep in mind that in all stages up to the **copy** stage, the status is reporting on sections for each distributor. After that, the **copy** and **commit** phases deal with the subscriber path.

| Overall Publication State | Section or Path State | Description |
|---|---|---|
| **Idle** | N/A | All configuration actions are possible. Must be in this state to initiate section update or content delivery process. |
| **CreatingEdition, EditionReady** | N/A | Only in these states briefly during the creation of a new edition. |
| **UpdatingSections** | **gsReposStageUpdate** | In this state when performing the potentially lengthy process of determining which files to fetch from section source, and fetching them. |
| Idle | **gsReposStageUpdateDone** | If section update is done separately from delivery, success returns to **Idle.** |
| **UpdatingSectionsIn-complete** | mix of **gsReposStageUpdateDone**, **gsReposStageUpdateFailed** and **Cancelled** | When multiple sections are being updated, not all completed. |
| **UpdatingSectionsFailed** | **gsReposStageUpdateFailed** | All section updates failed. |
| **PublicationPreparation** | N/A | A brief phase which constructs a work list for each global/SITE publisher and distributor from the configuration database. |
| **PubPrepDone** | **gsStable** | Successful completion of publication preparation. |
| **PubPrepFailed**, **PubPrepSuspended** | N/A | An error message is displayed, typically a distributor was not reachable. The delivery cannot proceed. Cancel resets the pub state to **Idle** so it can be retried. |
| **WANSync** | gsWanCDir | The changes necessary to bring each distributor up-to-date are created on the publisher global/SITE Controller. |
| **WANSync** | **gsWanCDirDone** | The set of changes is transmitted to each distributor. |

| Overall Publication State | Section or Path State | Description |
|---|---|---|
| **WANSync** | **gsWanSync** | Each distributor is updating its section(s) with the received changes. |
| **WANSyncDone** | **gsWanSyncDone** | Ready to proceed to **CreateChangeTrees.** |
| **WANSyncSuspended** | **gsWanSyncDone** | Will proceed to **CreateChangeTrees** after **Continue** or **Complete** is clicked. Only stops here when the corresponding Pub. Option is set. |
| **WANSyncCanceled** | some canceled | User clicked **Cancel** during phase. Cannot proceed with delivery. Reset publication to **Idle** and restart. |
| **WANSyncFailed** | **gsWanCDirFailed** | An error is displayed for the specific cause, typically lack of disk space on the publisher to store the temporary change file. |
| **WANSyncFailed** | **gsWanSyncFailed** | An error occurred while the distributor was updating its sections, typically lack of disk space on the distributor. |
| **CreateChangeTrees** | **gsCreateCDir** | The set of files necessary to update a set of subscribers from the current edition to the new edition is extracted from the section and stored on each global/SITE Controller. |
| **ChangeTreesDone** | **gsCreateCDirDone** | Ready to proceed to **CopyPhase.** |
| **ChangeTreesCanceled** | **Canceled** | The process was canceled by the user. The publication cannot proceed further. It must be reset to **Idle** and started over. |
| **ChangeTreesFailed** | **gsCreateCDirFailed** | An error, typically lack of disk space, will be displayed. The delivery process cannot proceed. |
| **CopyPhase** | **gsDataCopy** | Each global/SITE Controller copies new and changed files to temporary subdirectories on each subscriber |

| Overall Publication State | Section or Path State | Description |
|---|---|---|
| **CopyPhaseDone** | **gsDataCopyDone** | Ready to move to **CommitPhase** |
| **CopyPhaseSuspended** | **gsDataCopyDone** | Will proceed to **CommitPhase** after clicking **Continue** or **Complete**. User selected the Publication Option to stop at this point. |
| **CopyPhaseCanceled** | **canceled** plus various | User stopped subscriber path copies in progress, as well as those not started yet, by clicking the **Cancel** control. Only subscribers that successfully completed the copy phase for all sections can proceed to **CommitPhase** if the **Continue** or **Complete** control is clicked. Clicking **Reset** sets the publication state back to **Idle** without completing any deliveries. |
| **CopyPhaseFailed** | **gsDataCopyFailed** | An error is displayed for the specific cause, typically communication failure or lack of disk space on the subscriber. |
| **CommitPhase** | **gsDataCommit** | The copied new and changed files are being renamed from their temporary subdirectories to their final location. |
| **CommitPhaseDone** | **gsDataCommitDone** | The publication proceeds to update the configuration database and set its state back to **Idle**. |
| **CommitPhaseCanceled** | **canceled** plus others | The **Cancel** control was clicked, which stopped all commits in progress, as well as any not yet started. Avoide this if possible, as it leaves any incomplete subscribers in an indeterminate state. |
| **CommitPhaseFailed** | **gsDataCommitFailed** | Error message indicates the problem, usually a communication failure. |

***Table 2.1*** *The global/SITE Controller phases of publication*

# Making ongoing configuration changes

After you initially define a publication, various situations can arise that require changes to the configuration. This section provides overviews of a few of these situations:

- One or more subscribers may be unavailable during one or more publication cycles.
- A subscriber path has been removed permanently from a publication.
- A section has been removed from a publication.
- Other configuration changes may be required.

## Unavailable subscribers

For various reasons, you may have a given subscriber offline during a publishing cycle. In order to avoid having those subscribers fail during the process, you can mark them offline, or unavailable.

### Marking a path offline, and back online

You can use the Subscriber Detail screen to disable paths for individual sections.

***Figure 2.27*** *The Subscriber Detail screen*

### To disable a section for a particular subscriber

For each subscriber that is off line, you need to exclude all sections. On the Subscriber Detail screen:

1. Clear the check box next to each section.
2. When you have cleared the check boxes for each section, click the **Save Changes** button.

Any sections you have marked this way are ignored by the attempt to publish. And each must be returned to **Include** status before you can publish to it again.

### To re-enable a section for a particular subscriber

For each subscriber that was off line, when it returns online, you need to include all sections in order to publish to it. On the Subscriber Detail screen:

1. Click a check into the checkbox next to each section.
2. When you have checked the check boxes for each section that you want to include, click the **Save Changes** button. All sections with a check mark will be included in the next published edition.

You can also use the Deliver Edition screen to disable the path for a section. Clear the **Include** check box for each subscriber on a path to keep that subscriber from receiving the edition.

***Figure 2.28*** *The Deliver Edition screen*

### To disable a subscriber for a particular edition delivery

For each subscriber that is off line, you need to exclude the subscriber from the publication, for each section, before attempting to deliver the edition. On the Deliver Edition screen:

1. Clear the checkbox next to each offline subscriber, for each section listed.
2. When you click the **Deliver Edition** button, that subscriber will not receive the edition.

Any subscribers/sections that you have marked this way will be ignored by the attempt to publish this edition. Each must be returned to **Include** status before you can publish to it.

### To enable a subscriber for a publication

For each subscriber that was off line, when it returns online you need to include each subscriber for each section before you can publish to it. On the Deliver Edition screen:

1. Click a check into the checkbox next to each subscriber, for each section.
2. When you have checked the check boxes for each section that you want to include, click the **Deliver Edition** button.

All subscriber sections with a check mark will be included in the next published edition.

## To synchronize subscribers

If you run one or more publishing cycles while a subscriber is offline, the offline subscriber almost certainly has a stale version of the content. This can cause problems with future edition delivery, if you are not careful.

In a single publishing cycle, each section can change by only a single version. For example, if section **A** previously published version 5, and the new version is 7, that is fine, but that means all subscribers must contain either version 5, or version 7 (that is, up-to-date already, and not needing to change). If another subscriber for section **A** contains version 4, when you click the **Deliver Edition** button on the Deliver Edition screen, you will get an error message.

If you have out-of-date subscribers, you must bring them up-to-date before publishing a new edition. To do this, you run a publishing cycle to bring the stale subscriber paths up-to-date. The simplest way to do this is to re-publish the currently published edition to the entire publication. Even though all subscribers are listed in the Delivery Detail screen, no action will be taken because they are already up-to-date.

**◆ Tip**

*It is perfectly fine to skip multiple versions in a given publication. For example, the currently published version for section A in the publication is 7, and a subscriber currently contains version 2. This*

*can be published in a single cycle, going directly from version 2 to version 7, as opposed to going from version 2 to 3, then 3 to 4, and so on.*

### To check the current version for each subscriber

If you are concerned about out-of-date subscribers, you may want to check on your sections for subscribers to a particular publication before publishing a new version.

Start with the Publication Subscribers screen (the Subscribers tab).

***Figure 2.29*** *The Publication Subscribers screen*

The Publication Subscribers screen displays the currently published version for each subscriber if all paths are up-to-date.

Or you can use the Subscriber Detail screen, which shows the edition published to each path (as does the initial Edition Delivery screen).

***Figure 2.30*** *Subscriber Detail screen*

## Freeing up disk space by deleting editions and unused versions

Depending on the amount of change between versions, sections will occasionally need to be purged of old versions to recover disk space. This is a two step process:

1. You can delete editions using the delete button on the Publication Editions screen (Editions tab) of the publication detail screen.
   - You cannot remove the last published edition, and its listing does not show a delete button.
   - You should not remove any other editions that you may still want to rollback to at some point, or to publish to other subscribers.

2. Once you have deleted all unnecessary editions, go to the Section Detail screen and click the **Section History** tab, then click the **Remove Unused Versions** button.

***Figure 2.31*** *The Section Detail screen, History tab*

It may take a while for this process to complete if you have large sections. This process removes any versions that are no longer referenced by the remaining editions. Because distributors store versions of the publication's sections, this function also attempts to contact any distributors that contain versions of the section and remove the unneeded versions there, too.

## Removing a subscriber from a publication

The publication retains information about the last published section version on a subscriber until the subscriber is deleted, for example, if version 3 section **A** was last published to the subscriber. If that information becomes invalid, it may be necessary to remove the subscriber.

### To remove a subscriber from a publication

On the Publication Subscribers screen:

1. In the Subscriber list, locate the publication you want to remove.
2. Click the delete button  to its right.

***Figure 2.32*** *The Publication Subscribers screen*

This removes all knowledge of that subscriber from the publication. It does not remove any files from the subscriber itself.


# A

# Administrative Functions

# Using global/SITE administrative commands

This appendix contains:

- Backup strategy
- Command line utility terminology

# Backup strategy for global/SITE

Backup and restoration of the global/SITE Controller is currently performed using the command utility. This section provides procedures for backup and restoration under four potential situations.

## Backing up the global/SITE Controller

We recommend that you back up only the configuration database **(/gSITE/Configdb)**. Use the backup procedure that you prefer and are familiar with, and work with your F5 field service engineer to set it up.

## Restoring the global/SITE Controller

We provide procedures for restoration from four types of situations.

- One repository (section) on the primary/publisher is corrupted.
- There is a complete system failure or database corruption on the primary (publisher).
- A single repository/section on a distributor is corrupted.
- There is a complete system failure or database corruption problem on a distributor.

## Scenario one

One repository (*section*) on the primary (*publisher*) is corrupted.

1. To clean out the repository/section, run this command on the primary/publisher to remove this directory and all subdirectories:

```
rm -rf /gSITE/Data/GSID<local IID>/<group name>/Repos/<section
   name>
```

2. For each publication that uses this section, create an edition with version **-1** for that repository/section **(gsCreateEdition)**. All other sections must have their currently published version in this edition.
3. Use **gsSetPublished** to set all paths to the new edition in each applicable publication.
4. Delete all previous editions in each applicable publication, as they now refer to versions for this section that do not exist, but possibly will soon exist again with different content.
5. Do a section update for that repository/section. This will recreate and repopulate the directory that was removed above.

   If the stage is not in sync with the target/subscriber paths (perhaps untested and/or unpublished modifications have been made), then one of the target /subscriber paths must be used temporarily as the source path for this section. If you use a temporary source path, once the update is finished, remember to set the source path back to the original path.
6. If there are any distributors, you must copy the directory mentioned in step 1 (and its subdirectories) to each distributor, putting the distributors into sync with the primary's/publisher's repository/section. This may require the field engineer to do some setting up.
7. For each applicable publication, create an edition with version **1** for that repository/section (just updated with version **1**) and with the published version for all other sections.

8. Set all paths as published **(gsSetPublished)** with this new edition in each applicable publication. (This step assumes that all paths for this section were published with this version of this section. If not, you must do an actual publish.)

The effect of this restoration is that all previous versions of this repository/section have been lost.

## Scenario two

There is a complete system failure or a database corruption problem on the primary (publisher).

1. Restore the configuration database.
   This may require that **/gSITE/Configdb** be writable via NFS, and may require assistance from the field engineer.

   Update any changes made since the last backup of the database that are pertinent or needed immediately.

2. Perform the following steps using the procedures from ***Scenario one***:

   a) Perform step 1 of ***Scenario one*** for each section.

   b) Create an edition with version **-1** for all sections in all publications.

   c) Perform steps 3 and 4 of ***Scenario one***.

   d) Perform a section update for all sections using the instructions in step 5 of ***Scenario one***.

   e) Perform step 6 of ***Scenario one*** for each section.

   f) For each publication, create an edition with version **1** for each applicable section (just updated with version **1**).

   g) The same as step 8 of ***Scenario one***.

## Scenario three

On a distributor, a single repository (*section*) is corrupted.

1. On that distributor, to clean out the repository (*section*), run this command:

```
rm -rf /gSITE/Data/GSID<primary IID>/<group name>/Repos/<section
   name>
```

2. Copy that same directory and all of its subdirectories from the primary/publisher to the distributor.

## Scenario four

On a distributor, there is a complete system failure or database corruption problem.

1. Restore the configuration database and/or rebuild by hand (via the command utility).
2. For each section that this distributor distributes, go through the steps in ***Scenario three***.

**◆ Note**

*All of the above scenarios are complex and challenging. We recommend that you call F5 Tech Support for assistance from a field engineer. (206) 505-0888*

# Command line utility terminology

This section lists terms used primarily in the command utility. For a complete list of terms used in the browser-based application, check the ***Glossary*** section.

**content server**
Target servers for sets of files. The ***subscriber***, where content is available to the client's customers.

**FTPaccess**
An FTPaccess object name is a specific set (consisting of a node, user name, and password) used for FTP access to a node.

**FTPpath**
An FTPpath is a particular node, user name, and path. This is the directory path of either the source or target for publishing with the global/SITE Controller.

**Node**
The name of a source or target for an operation in the global/SITE Controller publication process. The node can be the IP address or the host name.

**primary**
In the browser-based application, referred to as a ***publisher.***

**pub group**
In the browser-based application, referred to as ***publication.***

**publication target**
In the browser-based application, referred to as ***subscriber.***

**repeater**
In the browser-based application, referred to as ***distributor.***

**repository**
In the browser-based application, referred to as ***section.***

**staging server**
Source servers that house files that will be included in a site.

**target server**
Content servers for sets of files. The server where global/SITE Controller delivers content. The ***subscriber***, where content is available to the client's customers.


# B

# Administrative Commands

# Using global/SITE commands

This appendix contains an alphabetic listing of the various global/SITE Controller commands, and a syntax sample. Most entries include descriptions of their typical usage, and some contain additional information about the command.

Available options include:

**-i** interactive, prompts you for next parameter as you enter each one. Some commands that require only two or three arguments, and/or very repetitive by nature, do not have this option.

Options specific to only a few commands are documented with the command.

Arguments enclosed in brackets [ ] are optional.

Note that there are some differences in terminology between the browser interface and the command line utility. Following is a quick reference chart. For a listing of command utility terminology refer to ***Appendix A***; for a complete browser interface glossary, refer to the ***Glossary***.

| command utility | browser interface |
|---|---|
| edition | edition |
| primary | publisher |
| pub group | publication |
| publication target | subscriber |
| repeater | distributor |
| repository | section |

# gsActivatePubGroupRepos

**gsActivatePubGroupRepos <pubgroup name> <repository name>**

**gsActivatePubGroupRepos [-i]**

## Description

This command activates a section, <**repository name**>, attached to publication, <**pubgroup name**>, for use by that publication. By default, a section is active, unless specifically declared inactive (see command gsInactivate**PubGroupRepos**). Active sections can be included when creating new editions. When updating sections, only active sections are to be updated, unless a specific list of sections to be updated is given (see **gsUpdateReposFromStage**).

The second form is the interactive form and is either without any arguments or simply with the **-i** option.

# gsCancelCommitPhase

**gsCancelCommitPhase [-v] <pub group name>**

## Description

This command cancels an active publication in any of the commit states:

- **CommitPhase**
- **CommitPhaseSuspended**
- **CommitPhaseCanceled**
- **CommitPhaseFailed**
- **CommitPhaseDone**

If successful, the publication returns to **Idle** state.

The optional **-v** switch can be included to give additional verbose printout.

# gsCancelCopyPhase

**gsCancelCopyPhase [-C][-v] <pub group name>**

## Description

This command cancels an active publication in any of the change tree or copy states:

- **CreateChangeTree**
- **ChangeTreesSusupended**
- **ChangeTreesCanceled**
- **ChangeTreesFailed**
- **ChangeTreesDone**
- **CopyPhase**
- **CopyPhaseSuspended**
- **CopyPhaseCanceled**
- **CopyPhaseFailed**
- **CopyPhaseDone**

The **-C** option sets the overall pubgroup state to **Idle**, allowing the publishing process to be restarted from the beginning. Without the **-C** option, any completed or failed operation will not be canceled. Only operations in process, as well as those that have not started yet will be canceled.

The optional **-v** switch can be included to give additional verbose printout.

## gsCancelEdition

**gsCancelEdition <pub group name>**

### Description

This command cancels an active publication in any of the **edition** states:

- **CreatingEdition**
- **EditionReadySuspended**
- **EditionReady**
- **EditionFailed**

The publication is set to **Idle** state.

## gsCancelUpdateFromStage

**gsCancelUpdateFromStage [-C] <pub group name>**

### Description

This command will stop updates in process, as well as cancel those that have not started yet.

Using this command with **-C** option sets the publication back to **Idle** (in the event that the **gsUpdateReposFromStage** command was interrupted before all updates were complete).

## gsCancelWanSync

**gsCancelWanSync [-C] [-v] <pub group name>**

### Description

This command cancels a publication in the following states:

- **WANSync**
- **WANSyncSuspended**
- **WANSyncCanceled**
- **WANSyncDone**
- **WANSyncFailed**

The **-C** option cancels and sets the publication back to **Idle** state. Without the **-C** option, any completed or failed operation will not be canceled. Only operations in process, as well as those that have not started yet, will be canceled.

The optional **-v** switch can be included to give additional verbose printout.

## gsChangeFtpPath

**gsChangeFtpPath <node name> <user name> <old path> <new path>**

### Description

Allows the user to modify the FTP path name.

The **<node name>** variable is node name to which FTP access entry applies. The **<user name>** variable is the current user name. The **<old path>** variable is the old full path, for example, **/usr/websrc/seahawks**. The **<new path>** variable is the new full path, for example, **/usr/websrc/seahawks**.

# gsCompletePublish

**gsCompletePublish <pub group name>**

## Description

This command completes the publication process once the publication is in either **CommitSuspended**, **CommitDone**, or **CommitCanceled** state.

Paths for which the **Commit** phase completed successfully will be updated with the appropriate edition, version(s), section(s), and a state of **Idle**.

**◆ Note**

*Paths for which the **Commit** phase failed will be set to a state of **CommitIncomplete**. Edition, version(s), section(s) information for that path will **NOT** change.*

The publication is set to **Idle** state, and if any path at all was successful, its published edition is set to the edition just published.

# gsConnectFtpPathRepeater

**usage: gsConnectFtpPathRepeater name=<nodename> user=<username> path=<directory path> repeater=<nodename>**

**usage: gsConnectFtpPathRepeater [-i]**

## Description

This command associates a particular path with a specific distributor. Both must already exist (see **gsCreateRepeater** and **gsCreateFtpPath**).

The second form is the interactive form and is either without any arguments, or simply with the **-i** option.

## gsConnectPubGroupPath

**gsConnectPubGroupPath name=<pub group>, node=<node name> user=<user name> [accesstype=<type>] path=<directory path>**

**gsConnectPubGroupPath [-i]**

### Description

This command assigns a particular path to a specific publication as a potential target for publication. Both must already exist (see **gsCreatePubGroup** and **gsCreateFtpPath**).

The second form is the interactive form and is either without any arguments or simply with the **-i** option.

## gsConnectPubGroupRepos

**gsConnectPubGroupRepos <pubgroup name> <repository name>**

**gsConnectPubGroupRepos [-i]**

### Description

This command assigns a specific section to be used as a source of content for a specific publication. Both must already exist (see **gsCreatePubGroup** and **gsCreateRepos**).

The second form is the interactive form and is either without any arguments or simply with the **-i** option.

## gsConnectReposSource

```
gsConnectReposSource name=<reposname> node=<nodename>
   [accesstype=<access type string>]  access=<username>
   path=<access path>


gsConnectReposSource [-i]
```

### Description

This command assigns a particular path as a staging source for this section. Both must already exist (see **gsCreateRepos** and **gsCreateFtpPath**).

The second form is the interactive form and is either without any arguments or simply with the **-i** option.

## gsConnectSourceToTarget

```
gsConnectSourceToTarget name=<pub group>, repository=<repository>
   node=<nodename> user=<username> [accesstype=<type>]
   path=<directory path>


gsConnectSourceToTarget [-i]
```

### Description

For the specified publication, this command specifies a section to be used as a source for content for the specified path. The publication, section, and path must already exist (see **gsCreatePubGroup**, **gsCreateRepos**, and **gsCreateFtpPath**). If the section and (or) the path have not already been associated with this specified publication, an attempt to do so will be made by this command. (For example, it will effectively attempt a **gsCommitPubGroupRepos** and/or **gsCommitPubGroupPath**.)

The second form is the interactive form and is either without any arguments or simply with the **-i** option.

# gsCreateEdition

**gsCreateEdition pubgroup=<pubgroup name> version=<vers. number>:<repos. name> ...**

**gsCreateEdition [-i]**

## Description

This command creates a new edition for the publication. If the non-interactive form of this command is used, the version/section pair for each section to be published as a part of this edition must be specified. If the interactive **-i** form of this command is used, the user will be prompted to enter the version for each active section.

The second form is the interactive form and is either without any arguments or simply with the **-i** option.

# gsCreateFtpAccess

**gsCreateFtpAccess name=<nodename> [description=<string>]**
**user=<username>**

**password=<password>**

**gsCreateFtpAccess [-i]**

## Description

Creates FTP access to the created primary, repeater, staging, or target servers.

The **<name>** variable is the name of a node on which you want FTP access. The **<description>** variable is the string describing the access entry. The **<user>** variable is the user name for FTP logon. The **<password>** is the clear text password for this user name.

The second form is the interactive form and is either without any arguments or simply with the **-i** option.

# gsCreateFtpPath

**gsCreateFtpPath name=<nodename> user=<username> path=<directory path> [repeater=<nodename>]**

**gsCreateFtpPath [-i]**

## Description

This creates a path associated with an access entry (and its node).

The **<name>** variable is a node name to which FTP access entry applies. The **<user>** variable is the user name. The **<path>** variable is the full path, that is, **/usr/websrc/seahawks**. The **<repeater>** variable specifies the repeater name, if applicable and preferred at this point.

Both the node name, user name, and repeater must already be known to the global/SITE Controller at this point (see **gsCreateNode**, **gsCreateFtpAccess**, and **gsCreateRepeater**).

# gsCreateNode

**gsCreateNode name=<nodename> [description=<string>] [cpu=<string>] [opsys=<string>]**

**gsCreateNode [-i]**

## Description

Create global/SITE Controller nodes for primary, distributor FTP accesses. Note that the local node is automatically created.

The primary is the main device performing most of the processing of content replication. Configuration and operational elements are stored in the configuration database resident on this device.

The **<name>** variable is a name by which the node can be accessed.

Note that if the node is to be used as a distributor, **gsCreateRepeater** will change the node name to a fully-qualified node name, if it is not already. For example, if **gsCreateNode** creates a node named **dist1**, **gsCreateRepeater** will change that node name to a fully-qualified node name. For example, assume **gsCreateNode** is used to create a node named **dist1**. If **gsCreateRepeater** is given the same name, the node and the repeater will be changed to a fully-qualified name; for example, **dist1.dev.net**.

The **<description>** variable is any additional description other than CPU type and operating system, that is, Western Region Production Machine.

The **<cpu>** variable is a string describing the CPU, that is, Intel PII or Alpha 8400.

The **<opsys>** variable is a string describing the operating system, that is, Linux or Windows NT.

The second form is the interactive form and is either without any arguments or simply with the **-i** option.

# gsCreatePubGroup

```
gsCreatePubGroup name=<pubgroup name> [pubstop=<pubstop type...>]
   [puberror=<puberror option>][description=<string>]

gsCreatePubGroup [-i]
```

## Description

This command creates a new publication. The **pubstop=** argument can be used as many times as necessary to specify the desired pub stops (**None**, **AfterEditionCreation**, **AfterPubPrep**, **AfterWANSync**, **AfterChangeTrees**, **AfterCopy**, **AfterCommit**).

Note that only **AfterWANSync**, **AfterCopy**, and **AfterCommit** are supported via the browser-based interface.

The **puberror=** option specifies action to be taken when a publication phase detects an error. Currently, the browser interface supports only one action, **Cancel**, but the command utility supports three: **Cancel**, **Pause**, **LogAndContinue**.

The second form is the interactive form and is either without any arguments or simply with the **-i** option.

# gsCreateRepeater

```
usage: gsCreateRepeater name=<nodename> [description=<string>]
   [identifier=<installation ID>] [secret=<encryption string>]
        where installation ID is either decimal or hexadecimal
   (0xffff)
usage: gsCreateRepeater [-i]
```

## Description

This command creates a repeater (*distributor*) that distributes the primary's (*publisher*) content to the target servers. Repeaters are remotely located and serve remote content servers (*subscribers*).

The **<name>** variable is the node name (already created by **gsCreateNode**) of the repeater/distributor.

The **<identifier>** variable is the installation ID (GSID) assigned when the global/SITE Controller was shipped.

The **<description>** variable is any description string.

The **<secret>** variable is the string used to create the 56-bit encryption key used to communicate with the distributor.

The second form is the interactive form and is either without any arguments or simply with the **-i** option.

The repeater name will be used to look for a node (see **gsCreateNode**) as shown in the following table:

| **gsCreateRepeater** has fully-qualified name | **gsCreateNode** has already created node with fully-qualified name | **Result** |
|---|---|---|
| No | No | Renames node to fully-qualified name. Creates repeater with fully-qualified name. |
| No | Yes | Creates repeater with fully-qualified name. |
| Yes | No | Creation fails because it either has no node created or it uses a short name. |
| Yes | Yes | Creates repeater with fully-qualified name. |

# gsCreateRepos

**gsCreateRepos name=<reposname> [type=FTPStage]**
**[description=<string>] [node=<nodename> [accesstype=<access type string>] access=<username> path=<accesspath>]**

**gsCreateRepos [-i]**

## Description

This command creates a repository. A repository (*section*) is a collection of related directories and files stored at the primary server (*publisher*). The data was extracted from staging servers.

Note that the FTP path must be created before creating a repository.

The **<name>** variable is a name given to the repository that is useful to the user of the global/SITE Controller.

The **<type>** variable is the type of source for this repository, currently only FTPStage is available.

The **<description>** variable is any descriptive string.

The **<node>** variable is the source node (stage) for FTP access.

The **<accesstype>** variable is currently unnecessary as there is only one type of source access, FTP.

The **<access>** variable is the user name of access entry.

The **<path>** variable is the source stage path from which this repository will retrieve its new versions.

The second form is the interactive form and is either without any arguments or simply with the **-i** option.

# gsDeleteFtpAccess

**gsDeleteFtpAccess  <nodename> <username>**

## Description

This command removes the FTP access to the primary, repeater, staging, or target server.

The **<name>** variable is the node name on which FTP access is to be deleted.

The **<user>** variable is the user name for FTP logon.

All paths that use this access (this node and user) must have already been deleted (see **gsDeleteFtpPath**).

# gsDeleteFtpPath

**gsDeleteFtpPath node=<nodename> user=<username> path=<directory path>**

**gsDeleteFtpPath [-i]**

## Description

This command removes the FTP path.

The **<node>** variable is the node name to which FTP access entry applies.

The **<user>** variable is the current user name. The **<path>** variable is the full path; for example, **/usr/websrc/seahawks**.

The second form is the interactive form and is either without any arguments or simply with the **-i** option.

The path cannot be deleted if it is recorded as having published content on it.

You must either:

- Actually publish an edition with the null revision (currently version **-1**) for that section, requiring that you set other paths served by that section(s) to **unavailable** so their data will not be removed, also
- Use the command **gsSetPublished** to simply record this path as empty, using an edition with version **-1**, without actually publishing to it.

## gsDeleteNode

**gsDeleteNode <nodename>**

### Description

This command removes a node used by a repeater, staging, or target server.

The node cannot be deleted if there is a repeater by that name or any FTP access entries that reference the node to be deleted.

## gsDeletePubGroup

**gsDeletePubGroup <pubgroup name>**

### Description

This command deletes the publication. All connected objects (paths, sections) must be disconnected first.

## gsDeleteRepeater

**gsDeleteRepeater <node name>**

### Description

This command deletes a repeater. All paths must be disconnected first (see **gsDisconnectFtpPathRepeater**).

# gsDeleteRepos

**gsDeleteRepos <reposname>**

## Description

This command removes repository information. The repository must no longer be connected to any publication (see **gsDisconnectPubGroupRepos**).

The **<reposname>** is a name given to the repository.

**WARNING**

*We do not recommend use of this command, until cleanup of the repository copies on distributors is complete.*

# gsDeleteVersions

**usage: gsDeleteVersions repos=<repository name> [begin=<version number>] [end=<version number>] [keep=<min. number to keep>] [age=<min. age to delete>] [verbose=<y/n>]**

**usage: gsDeleteVersions [-i]**

## Description

This command deletes versions from a section.

The **[begin]** variable is the optional first version number to consider deleting.

The **[end]** variable is the optional last version number to consider deleting.

The **[keep]** variable is the minimum number of versions to keep.

The **[age]** variable is the minimum age in days to delete.

The **[verbose]** variable provides the option to get detailed information.

The second form is the interactive form and is either without any arguments or simply with the **-i** option.

# gsDestroyPubSubscriber

**usage: gsDestroyPubSubscriber name=<pub group>, node=<nodename> user=<username> [accesstype=<type>]**

**usage: gsDestroyPubSubscriber [-i]**

## Description

For the specified publication, this command tries to remove as much of the specified subscriber as possible.

First, an edition is temporarily created that specifies removal of all content for each section. For each path serviced by this subscriber, that edition is recorded as that path's published edition. This effectively means that nothing is published on that path. The path is then disconnected from the section(s) that is its source of content.

Next, the path is disconnected from the publication and from the distributor, after which the path is deleted. Once all paths have been processed, the **FtpAccess** entry **<user>** is deleted, including its node.

Finally, the newly created edition is deleted. There may be failures during the process, such as being unable to delete the user because other publications also use it; consequently, errors are ignored and as much as possible is done.

**WARNING**

*This procedure is dangerous. Be sure you understand the consequences before you use it.*

The second form is the interactive form and is either without any arguments or simply with the **-i** option.

## gsDisconnectFtpPathRepeater

**usage: gsDisconnectFtpPathRepeater name=<nodename> user=<username> path=<directory path>**

**usage: gsDisconnectFtpPathRepeater [-i]**

### Description

This command severs the connection between a path and its distributor, leaving the path disconnected from any distributor.

The second form is the interactive form and is either without any arguments or simply with the **-i** option.

## gsDisconnectPubGroupPath

**gsDisconnectPubGroupPath name=<pub group>, node=<nodename> user=<username> [accesstype=<type>] path=<directory path>**

**gsDisconnectPubGroupPath [-i]**

### Description

This command disconnects the association between a content path and a publication leaving the path unserviced by any publication. Any connection between the path and a source section to be published on that path must be severed first (see **gsDisconnectSourceFromTarget**).

The second form is the interactive form and is either without any arguments or simply with the **-i** option.

# gsDisconnectPubGroupRepos

**gsDisconnectPubGroupRepos <pubgroup name> <repository name>**

**gsDisconnectPubGroupRepos [-i]**

## Description

This command disconnects the association between a section and a publication. The section must first be disconnected from any content targets (see **gsDisconnectSourceFromTarget**).

The second form is the interactive form and is either without any arguments or simply with the **-i** option.

# gsDisconnectReposSource

**usage: gsDisconnectReposSource name=<reposname> node=<nodename>**
**[accesstype=<access type string>] access=<username>**
**path=<access path>**

**usage: gsDisconnectReposSource [-i]**

## Description

This command severs the relationship between a section and its staging (source) path, leaving the section without a source/stage.

The second form is the interactive form and is either without any arguments or simply with the **-i** option.

## gsDisconnectSourceFromTarget

```
gsDisconnectSourceFromTarget name=<pub group>
   repository=<repository> node=<nodename> user=<username>
   [accesstype=<type>] path=<directory path>

gsDisconnectSourceFromTarget [-i]
```

### Description

This command disconnects the source section from the target content path, leaving both the section and path still associated with the publication.

The second form is the interactive form and is either without any arguments or simply with the **-i** option.

## gsInactivatePubGroupRepos

```
gsInactivatePubGroupRepos <pubgroup name> <repository name>
gsInactivatePubGroupRepos [-i]
```

### Description

This command inactivates a section for the specified publication. While the section is inactivated, it cannot be used in any new editions created. It also will not be updated by **gsUpdateReposFromStage**, unless it is specifically included by the command line.

The second form is the interactive form and is either without any arguments or simply with the **-i** option.

## gsPrepareToPublish

**gsPrepareToPublish <pub group name> <edition number>**

### Description

This command prepares a publication and the distributors to be used for a publishing cycle using the specified edition.

The publication is left in the state **PubPrepDone**.

Normally, the **gsPublish** command should be used rather than this command. Use this command only if you only want to prepare to publish, and then pause.

## gsPublish

**usage: gsPublish [-v] [-d] <pub group name> <edition number>**

### Description

This command directs the publication to start publishing, or to resume publishing from some previously stopped publishing cycle. If **<edition number>** is zero, the publication will first update all active sections, and then create a new edition using the latest version from each section.

The optional **-v** switch can be included to give additional verbose printout.

The optional **-d** switch causes the process to fork (used primarily for testing at F5 Networks).

## gsResetPubGroupState

**gsResetPubGroupState <pub group name> <state>**

### Description

This command is used to set the pubgroup to **Idle** immediately after preparing a publication with **gsPrepareToPublish**, thus canceling the publishing process before proceeding further. Use one of the phase-specific **gsCancel** commands for all other states.

## gsResetReposState

**gsResetReposState <repos name> <state>**

### Description

This command is used to set the repositories to **Idle** after updating sections with **gsUpdateReposFromStage** only if some abnormal occurrence left the section in a state other than **Idle**.

# gsSetFtpAccessInfo

**gsSetFtpAccessInfo node=<nodename> user=<username> [newname=<new username>] [description=<string>] [password=<new password>]**

**gsSetFtpAccessInfo [-i]**

## Description

This command allows modification of the FTP information.

The **<node>** variable is a node name to which FTP access entry applies.

The **<user>** variable is the current user name.

The **<newname>** is a new user name.

The **<description>** a new description.

The **<password>** is new clear text password.

The second form is the interactive form and is either without any arguments or simply with the **-i** switch.

# gsSetFtpAccessPathState

**gsSetFtpAccessState node=<nodename> user=<username> path=<full path> state=<state>**

**gsSetFtpAccessState [-i]**

## Description

This command is used to set the state of a path to either **Idle** or **Unavailable**.

Setting the path to **Idle**, allows publishing to that path.

Setting the path to **Unavailable** prohibits any publishing to that path.

The second form is the interactive form and is either without any arguments or simply with the **-i** switch.

# gsSetNodeInfo

```
gsSetNodeInfo name=<nodename> [newname=<new nodename>]
   [description=<string>] [cpu=<string>] [opsys=<string>]

gsSetNodeInfo [-i]
```

## Description

This command allows the previously-entered information to be changed.

The **<newname>** variable is the new node name of the node. It can be the alias. If the node is to be a distributor node, be aware of the ramifications of renaming the node (see **gsCreateNode**).

The **<description>** variable is any additional description other than cpu type and operating system, for example, Western Region Production Machine.

The **<cpu>** variable is a string describing the cpu, for example, Intel PII or Alpha 8400.

The **<opsys>** variable is a string describing the operating system, for example, Linux or Windows NT.

An example:

```
[saxon@saxon mgmtdi]$ gsSetNodeInfo -i
Enter Current Node Name: saxon55

Enter New Node Name [.=saxon55]: .
Enter CPU description [.Pentium II]:.
Enter Operating System description [.=Windows NT]: WNT
Enter additional description [.=Personal Workstation]:.
set info completed.
[saxon@saxon mgmtcli]$
```

Notice the use of the period (**.**) to denote no change; the current value is displayed in brackets, showing what the entry of a period (.) will leave as the value.

The second form is the interactive form and is either without any arguments or simply with the **-i** switch.

# gsSetPubGroupInfo

```
gsSetPubGroupInfo name=<pubgrp name> [newname=<new pubgrp name]
   [description=<description>] [protocol=<publishing protocol>]
   [pubstop=<pubstop type> ...] [puberror=<puberror option>]

gsSetPubGroupInfo [-i]
```

## Description

This command is used to modify publication information.

The **<newname>** variable is used to change the publication's name.

The **<description>** variable is used to change the description.

The **<pubstop>** variable is used to change the pubstops, and can be repeated as often as necessary (see **gsCreatePubGroup**).

The **<puberror>** variable is used to change the action taken when a phase detects that an error occurred (see **gsCreatePubGroup**).

The second form is the interactive form and is either without any arguments or simply with the **-i** switch.

# gsSetPublished

**usage: gsSetPublished pubgroup=<name> edition=<edition number>**
**[[node=<nodename> user=<username>] path=<directory path>]**
**[setedition=no]**

**usage: gsSetPublished [-i]**

## Description

This command sets information as to the edition published without actually publishing the edition.

If **<node>**, **<user>**, and **<path>** are specified, only that path is set to say that **<edition>** (and the appropriate version/section) is published on it.

If only **<node>** and **<user>** are specified, all paths for that node/user pair are set for this publication.

If **<node>** and **<user>** are not specified, all paths for this publication are set.

If **<set edition>** is yes, **<edition>** is set as the published edition for this publication.

**WARNING**

*Use this command with CAUTION.*

The second form is the interactive form and is either without any arguments or simply with the **-i** switch.

# gsSetRepeaterInfo

**usage: gsSetRepeaterInfo name=<nodename> [description=<string>]**
**[identifier=<installation ID>] [secret=<encryption string>]**

Where the installation ID is either decimal or hexadecimal (0xffffffff).

**usage: gsSetRepeaterInfo [-i]**

## Description

This command changes the distributor information.

The second form is the interactive form and is either without any arguments or simply with the **-i** switch.

# gsSetReposInfo

**gsSetReposInfo name=<reposname> [newname=<new repos name]**
**[description=<description>] [[exception=<full path> ... ] |**
**[[exception+=<full path> ... ] [exception-=<full path> ... ]]]**

**gsSetReposInfo [-i]**

## Description

This command modifies repository information.

The <**reposname**> variable is a name given to the repository.

Use of <**exception**> precludes use of <**exception+** > or <**exception-** >, and replaces the section's exception paths with the specified paths (can be used multiple times in one command).

Use of <**exception+** > and/or <**exception-** > specifies adding and removing exception paths, respectively (can be used multiple times in one command).

The second form is the interactive form and is either without any arguments or simply with the **-i** switch.

## gsShowAllFtpAccess

**gsShowAllFtpAccess [-v]**

### Description

This command prints on screen a list of all <**node**> and <**user**> pairs.

The **-v** option prints out detailed information for each, including a clear-text password.

## gsShowAllNodes

**gsShowAllNodes [-v]**

### Description

This command prints, on screen, a list of all nodes.

The **- v** option prints out detailed information for each node.

## gsShowAllPubGroups

**gsShowAllPubGroups [-v]**

### Description

This command prints, on screen, a list of all publications.

The **-v** option prints out detailed information for each publication.

# gsShowAllRepeaters

**gsShowAllPubGroups [-v]**

## Description

This command prints, on screen, a list of all distributors.

The **-v** option prints out detailed information for each distributor.

Sample response with **-v** option:

```
Node:           globalsite4.your.net
Type:           Repeater
Description:    Local node distributor
State:          Idle
Mode:           Direct

Installation ID: 0x4
Encryption Secret:
no paths currently connected to this repeater...
```

Another partial sample response:

```
Node:           globalsite2.your.net
Type:           Repeater
Description:
State:          Idle
Mode:           Direct

Installation ID: 0x2c
Encryption Secret: GlobalSiteTest
        Repeats to following trees:
        Node: 30.30.30.30
        Type: FTP
```

```
FTP Access Entry: gsite
Path: /gSITE/prod/targ2/images

Node: 30.30.30.30
Type: FTP
Ftp Access Entry: gsite
Path: /gSITE/prod/targ2
```

## gsShowAllRepos

**gsShowAllRepos [-v]**

### Description

This command prints on screen a list of all sections.

The **-v** option prints out detailed information for each section.

## gsShowFtpAccess

**gsShowFtpAccess <nodename> <username>**

### Description

This command displays FTP access information for one specific user, including clear-text password.

The <**nodename**> variable is a node name to which FTP access entry applies.

The <**username**> variable is the user name.

## gsShowNode

**gsShowNode <nodename>**

### Description

This command displays information pertaining to a specific primary, repeater, staging, or target server.

## gsShowPubGroup

**gsShowPubGroup <pubgroup name>**

### Description

This command displays information pertaining to a specific publication.

## gsShowPubStatus

**gsShowPubStatus <pubgroup name>**

### Description

This command displays information pertaining to the status of a specific publication step.

## gsShowRepeater

**gsShowRepeater <node name>**

### Description

This command shows the status of a repeater (in the browser interface, the *distributor*).

## gsShowRepos

**gsShowRepos <repository name>**

### Description

This command displays repository information.

The <**reposname**> is a name given to the repository.

## gsStartCommitPhase

**gsStartCommitPhase <pub group name>**

### Description

This command starts the commit phase of a publishing cycle. The previous phase, copy phase, must be completed, and is indicated by the publication state.

## gsStartCopyPhase

**gsStartCopyPhase [-v] [-w | -r]] <pub group name>**

### Description

This command starts the copy phase of a publishing cycle. The previous phase, WANSync phase, must be completed, and is indicated by the publication state.

The **-w** option stops processing after the change trees are prepared, but before they are actually copied to the publication targets. It can be used with the **-v** option, which displays additional status information. After the change trees are prepared, the user is queried as to whether to continue with the actual copy or not.

If the copy phase is not performed immediately (after saying "no" to the **-w** prompt), then the copies can be initiated without recreating the changed files by using the **-r** option to resume the copy phase. It can be used with the **-v** option, which displays additional status information.

The **-v** option alone proceeds with both change trees, and the copying of changed files to the publication targets, providing additional status information during the process.

Using none of the options proceeds with change trees and copying of changed files to the publication targets without providing the additional status information.

## gsStartWanSync

**gsStartWanSync [-v] <pub group name>**

### Description

This command starts the **WANSync** phase of publishing. Publication preparation must be successfully completed (see **gsPreparetoPublish**).

The optional **-v** switch can be included to give additional verbose printout.

## gsUpdateReposFromStage

**gsUpdateReposFromStage <pub group name> [reposname ...]**

### Description

By specifying the optional list of repositories, this command initiates an update operation only for those repositories specified, as opposed to all active repositories in the pubgroup, which is the default.

Specifying the optional list can override a repository's inactive status temporarily by including it in the specified list, or override its active status temporarily by not including it in the list.

**Example:**

**gsUpdateReposFromStage Pub1 section1 section2**

Only Section 1 and Section 2 are being updated.

**gsUpdateReposFromStage Pub1**

All sections on the active list are being updated.

d to as ***section.***

# Glossary

**commit**

To move files from a temporary to a permanent directory on a *subscriber* (subscribing server).

**current edition**

The edition of a publication published to a subscriber and available to its users. All subscribers should have the same current edition. There may be multiple editions of a global/SITE publication on the publisher, but only one is currently published to a subscriber.

**deliver edition**

The final step in the global/SITE Controller publication process, where the content, a specific edition of a publication, is transferred to the target servers, or *subscribers*.

**distributor**

A global/SITE Controller that is remotely located and provides the publisher access to additional subscribers.
Compare to *publisher.*

**edition**

A complete collection of specific versions of selected sections. An edition is an instance of a publication that indicates which versions of selected sections are to be included, and the edition is ready to be delivered (*published*), or has already been delivered to subscribing servers.
A publication may have multiple editions, but only one is currently published to the publication's subscribers. Past editions may be kept for backup, history, comparison, and other purposes.

**exception**

A directory, located within a path of a section, that is to be excluded from the section when creating a section version. To exclude a directory from a section, you must specify the absolute path of the directory to be excluded. You can exclude only directories, not files.

**FTP**

File Transfer Protocol. A method that global/SITE Controller uses to collect and publish file-based content to internet sites.

**global/SITE identifier**

The unique, pre-assigned numeric identifier that distinguishes one global/SITE Controller from all others. You must have this identifier in order to add a global/SITE Controller as a distributor. If you change the global/SITE Identifier, it makes your section data unreachable, and prevents communication between your global/SITE Controllers.

**key phrase**

A phrase that is shared by the pair of publish and distribute global/SITE Controllers. Both the publisher and the distributor must use the same key phrase in order to communicate. You must have the key phrase in order to add a global/SITE Controller as a distributor.

**publication**

Each global/SITE Controller publication is a collection of information about subscribers, content, and publication options. It maps and records where to get source information (*content*), and where to store it, and how and where to deliver (*publish*) it. A publication may maintain multiple editions simultaneously; it defines which *versions* of which *sections* are used in each *edition*, and which go to each *subscriber.*

**publication options**

Currently, publication options allow you to have the global/SITE Controller pause during the publishing process so that you can check status.

**publication targets**

See *subscribers.*

**publish**

The process of identifying specific versions of sections to be included in an edition and then delivering this content to subscribers. You publish editions of a publication in order to make the subject matter available to the users of your internet site.
See *deliver edition.*

**publisher**

The global/SITE Controller that is sending content to the subscribers. A publisher may send content directly to subscribers, or to another global/SITE Controller functioning as a distributor that passes the publication to the subscribers.
See also, *distributor.*

**repeater**

See *distributor.*

**repository**

See *section*.

**section**

The source content retrieved via one path (server and directory) and access method (authorization and protocol). Different updates are saved as unique section versions, which may be used in different publications and publication editions.

**subscriber**

The location reached via a specific path (server and directory) and access method (authorization and protocol) where content is delivered, or *published*. The subscriber is the place that content is published *to* for access by customers. In the global/SITE Controller, the subscriber record contains a destination path for each section of a publication.

**version**

A version is one particular instance of a section that differs from other earlier or later instances of that section due to changes or modifications. Sections can have multiple versions which are numbered and dated to identify them as unique. You can specify the version of any section that you want included in an edition of a publication.