

The logo for F5 Networks Global Site Content Controller, featuring the text 'F5 Networks' in a small font above the 'GLOBAL·SITE' text, which is enclosed in a white swoosh. Below the swoosh is the text 'Content Controller' in a small font.

F5 Networks

GLOBAL·SITE™

Content Controller

Administrator Guide

GLOBAL-SITE™ Content Controller Administrator Guide

version 1.1

Service and Support Information

Product Version

This manual applies to version 1.1 of the GLOBAL-SITE™ Content 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 Suite 500 Seattle, Washington 98119

Legal Notices

Copyright

F5 Networks, Inc. (F5) believes the information it furnishes to be accurate and reliable. However, F5 assumes no responsibility for the use of this information, nor any infringement of patents or other rights of third parties which may result from its use. No license is granted by implication or otherwise under any patent, copyright or other intellectual property right of F5 except as specifically described herein. F5 reserves the right to change specifications at any time without notice.

Copyright 1999-2000, F5 Networks, Inc. All rights reserved.

Trademarks

F5, BIG-IP, and 3-DNS are registered trademarks of F5 Networks, Inc. SEE-IT and GLOBAL-SITE are trademarks of F5 Networks, Inc. Other product and company names are registered trademarks or trademarks of their respective holders.

Export Regulation Notice

The GLOBAL-SITE™ Content Controller may include cryptographic software. Under the Export Administration Act, the United States government may consider it a criminal offense to export this GLOBAL-SITE™ Content Controller from the United States.

Export Warning

This is a Class A product. In a domestic environment this product may cause radio interference in which case the user may be required to take adequate measures.

Canadian Regulatory Compliance

This class A digital apparatus complies with Canadian ICES-003.

FCC Compliance

This equipment generates, uses, and may emit radio frequency energy. The equipment has been type tested and found to comply with the limits for a Class A digital device pursuant to Part 15 of FCC rules, which are designed to provide reasonable protection against such radio frequency interference.

Operation of this equipment in a residential area may cause interference in which case the user at his own expense will be required to take whatever measures may be required to correct the interference.

Any modifications to this device - unless expressly approved by the manufacturer - can void the user's authority to operate this equipment under part 15 of the FCC rules.

Acknowledgments

This product includes software licensed from Thomas Pfau under the GNU Library General Public License (© 1991 Free Software Foundation, Inc.), www.gnu.org/copyleft/lgpl.html.

This product includes the Linux operating system distributed and licensed from Red Hat Software, Inc., pursuant to the GNU General Public License (© 1989, 1991 Free Software Foundation, Inc.) www.gnu.org/copyleft/lgpl.html.

This product includes software licensed from DuraSoft GmbH.

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

This product includes software licensed from Richard H. Porter under the GNU Library General Public License (© 1998, Red Hat Software), www.gnu.org/copyleft/lgpl.html.

This product includes the standard version of Perl software distributed and licensed from Red Hat Software, Inc., pursuant to the GNU General Public License (© 1989, 1991 Free Software Foundation, Inc.) www.gnu.org/copyleft/lgpl.html. You may find the most current standard version of Perl software at <http://www.perl.com>.

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. 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. 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™ Content 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. In countries other than the US, Licensee agrees to comply with the local regulations regarding exporting or using cryptographic software.
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

Introduction to the GLOBAL-SITE Controller

Welcome to the GLOBAL-SITE Controller	1-1
Application tools	1-1
Specifications	1-4
GLOBAL-SITE Controller features	1-5
Application and content synchronization	1-5
Version control	1-5
Site recovery and one-button rollback	1-6
Content and administration security	1-6
Visibility into publishing process	1-6
Adaptable open systems and web publishing support	1-7
Secure, succinct network traffic	1-7
Environmental requirements for GLOBAL-SITE Controller hardware	1-8
What's new in version 1.1	1-10
BIG-IP Controller integration	1-10
Non-archived publications	1-11
Scheduled publishing	1-11
Scheduled activation of new content	1-11
Improved status and error reporting	1-12
Multi-disk support	1-12
First-Time Boot utility	1-12
Finding help and other technical support resources	1-13
Release notes	1-13
Online help	1-13
Technical support	1-14

Chapter 2

Publishing Documents

Why GLOBAL-SITE Controller?	2-1
Understanding GLOBAL-SITE objects	2-1
Planning to configure the GLOBAL-SITE Controller	2-3
Defining publications	2-4
Defining sections	2-10
Specifying the subscribers for section content	2-17
Using the publishing cycle	2-21

Getting new content	2-22
Publishing new content	2-24
Displaying status	2-26

Chapter 3

Increased Control Over Publishing

Using enhanced publishing features	3-1
Scheduling the publishing process	3-2
Scheduled publications	3-2
Scheduled activation of new content	3-6
Comparing controlled and independent activation	3-8
Options for controlled activation	3-9
Working with the BIG-IP Controllers and virtual servers	3-9
Guidelines for BIG-IP Controller GLOBAL-SITE Controller cooperation	3-10
Creating a BIG-IP Controller on the GLOBAL-SITE Controller	3-11
Marking BIG-IP Controller nodes as down while updating content with the GLOBAL-SITE Controller	3-14
Creating non-archived publications	3-16
Looking at benefits of non-archived publications	3-17
Planning for non-archived publications	3-17
Considering the details of non-archived publishing	3-17
Using procedures for non-archived publications	3-19
Defining exceptions to sections	3-22
Using basic publishing functions	3-24
Making ongoing configuration changes	3-25
Unavailable subscribers	3-25
Understanding the GLOBAL-SITE Controller phases	3-32
Interpreting phases in the status line	3-33
Using the Phases of Publication table	3-35

Appendix A

Administrative Functions

Using GLOBAL-SITE administrative commands	A-1
Backup strategy for GLOBAL-SITE	A-1
Backing up the GLOBAL-SITE Controller	A-1

Restoring the GLOBAL-SITE Controller	A-1
Command line utility terminology	A-5

Appendix B

Administrative Commands

Using GLOBAL-SITE commands	B-1
gsActivatePubGroupRepos	B-2
gsCancelCommitPhase	B-2
gsCancelCopyPhase	B-3
gsCancelDeleteRepos	B-4
gsCancelDeleteVersions	B-4
gsCancelEdition	B-5
gsCancelPostPubPhase	B-5
gsCancelTestConnection	B-5
gsCancelUpdateFromStage	B-6
gsCancelWanSync	B-7
gsChangeFtpPath	B-7
gsChangePubGroupSchedule	B-8
gsCompletePublish	B-14
gsConnectFtpPathRepeater	B-15
gsConnectPubGroupPath	B-16
gsConnectPubGroupRepos	B-16
gsConnectReposSource	B-17
gsConnectSourceToTarget	B-17
gsCreateBigip	B-18
gsCreateEdition	B-19
gsCreateFtpAccess	B-20
gsCreateFtpPath	B-21
gsCreateNode	B-22
gsCreatePubGroup	B-24
gsCreateRepeater	B-28
gsCreateRepos	B-29
gsDeleteBigip	B-30
gsDeleteFtpAccess	B-31
gsDeleteFtpPath	B-32
gsDeleteNode	B-33
gsDeletePubGroup	B-33
gsDeleteRepeater	B-33

gsDeleteRepos	B-34
gsDeleteVersions	B-35
gsDestroyPubSubscriber	B-36
gsDisconnectFtpPathRepeater	B-37
gsDisconnectPubGroupPath	B-37
gsDisconnectPubGroupRepos	B-38
gsDisconnectReposSource	B-38
gsDisconnectSourceFromTarget	B-39
gsInactivatePubGroupRepos	B-39
gsPausePubGroupSchedule	B-40
gsPrepareToPublish	B-41
gsPublish	B-42
gsResetNodeState	B-43
gsResetPubGroupState	B-44
gsResetPublish	B-44
gsResetReposState	B-45
gsResumePubGroupSchedule	B-45
gsSetBigipInfo	B-46
gsSetFtpAccessInfo	B-47
gsSetFtpAccessPathState	B-48
gsSetNodeInfo	B-49
gsSetPubGroupInfo	B-51
gsSetPublished	B-53
gsSetRepeaterInfo	B-55
gsSetReposInfo	B-56
gsShowAllBigips	B-57
gsShowAllFtpAccess	B-58
gsShowAllNodes	B-58
gsShowAllPubGroups	B-59
gsShowAllRepeaters	B-60
gsShowAllRepos	B-61
gsShowBigip	B-62
gsShowFtpAccess	B-63
gsShowNode	B-63
gsShowPubGroup	B-64
gsShowPubStatus	B-65
gsShowRepeater	B-65
gsShowRepos	B-66
gsStartCommitPhase	B-67
gsStartCopyPhase	B-67

gsStartWanSync B-68
gsUpdateReposFromStage B-68

Appendix C
Reconfiguring the Controller

Changing the host name and/or IP address of a GLOBAL-SITE
Controller C-1
 Editing the network and hosts files C-2
 Resetting the GLOBAL-SITE identifier C-3
 Updating the distributor lists on other GLOBAL-SITE
 Controllers C-4

Glossary

Index

I

Introduction to the GLOBAL-SITE Controller

- Welcome to the GLOBAL-SITE Controller
- GLOBAL-SITE Controller features
- Environmental requirements for GLOBAL-SITE Controller hardware
- What's new in version 1.1
- 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 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 Configuration Utility home page for GLOBAL-SITE Controller.



Figure 1.2 Configuration Utility home page

The browser application

The Publication List screen (shown in Figure 1.1) gives you an 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 (non-archived publications do not show an Editions tab). 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 **navigation pane**, displayed in the left frame, provides links to both help, and application screens. Click the headings Publishing, Distribution, BIG-IPs, or System to see summary help for that group of functions. Click any of the other options on the navigation pane to access an application screen for use.

The screens available from the Distribution, BIG-IPs, and System areas in the navigation pane assist you in managing the administrative functions of the GLOBAL-SITE Controller. Use these screens for viewing or adding GLOBAL-SITE Controllers

that function as distributors, for creating and managing your BIG-IP Controllers and virtual servers, for adding or removing user accounts, and for setting the system time and pausing or restarting the GLOBAL-SITE Publication Scheduler.

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 UNIX command line utility

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

Specifications

The GLOBAL-SITE Controller supports all source and target servers running UNIX, NT, and MAC operating systems. It adapts to your existing server and publishing systems.

Operating systems

The GLOBAL-SITE Controller runs on Redhat_{TM} Linux version 6.1.

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 content 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 not limited to publishing only the most recent content. In archived publications, 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 it easy to view the publication organization, and monitor the publishing process. 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 Publish Progress Display screen, 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 manual publishing or delayed activation options on the Publication 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 consider 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 publisher sends to the distributor only the differences (file changes) between the current edition on the distributor and the new edition.
- ❖ The information is encrypted by the GLOBAL-SITE publisher.

- ❖ 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.

Environmental requirements for GLOBAL-SITE Controller hardware

General guidelines

A GLOBAL-SITE Controller is an industrial network appliance, designed to be mounted in a standard 19-inch rack. To ensure safe installation and operation of the unit:

- ❖ Install the rack according to the manufacturer's instructions, and check the rack for stability before placing equipment in it.
- ❖ Build and position the rack so that once you install the GLOBAL-SITE Controller, the power supply and the vents on both the front and back of the unit remain unobstructed. The GLOBAL-SITE Controller must have adequate ventilation around the unit at all times.
- ❖ Do not allow the air temperature in the room to exceed 40° C.

- ❖ Do not plug the unit into a branch circuit shared by more electronic equipment than the circuit is designed to manage safely at one time.
- ❖ Verify that the voltage selector is set appropriately before connecting the power cable to the unit.



The unit must be connected to Earth ground, and it should have a reliable ground path maintained at all times.



The GLOBAL-SITE Controller contains a lithium battery. There is danger of an explosion if you replace the lithium battery incorrectly. We recommend that you replace the battery only with the same type of battery originally installed in the unit, or with an equivalent type recommended by the battery manufacturer. Be sure to discard all used batteries according to the manufacturer's instructions.



This equipment is not intended for operator servicability. To prevent injury and to preserve the manufacturer's warranty, allow only qualified service personnel to service the equipment.

Guidelines for DC powered equipment

A DC powered installation must meet the following requirements:

- ❖ Install the unit using a 20 Amp external branch circuit protection device.
- ❖ For permanently connected equipment, incorporate a readily accessible disconnect in the fixed wiring.
- ❖ Use only copper conductors.



Install DC powered equipment only in restricted access areas, such as dedicated equipment rooms, equipment closets, or similar locations.

What's new in version 1.1

This section introduces the main features that are new to version 1.1 of the GLOBAL-SITE Content Controller, and tells where you can find detailed information about each feature.

BIG-IP Controller integration

GLOBAL-SITE Controller version 1.1 introduces the first step in integration with BIG-IP Controllers and support for scheduling publishing processes. BIG-IP Controller integration provides several benefits to controlling and maintaining site performance during content update, and introduces a number of new controls over the content activation process. Content activation is one of the final steps in publishing: it occurs after the copying phase, and is when the GLOBAL-SITE Controller changes the temporary files on the subscriber to permanent, active files that customer viewers can access. For details, see *Scheduled activation of new content*, in Chapter 3.

Non-archived publications

Prior to this version, the GLOBAL-SITE Content Controller has been used as a content version control system as well as for propagation of content from section servers to subscriber servers. Version 1.1 introduces non-archived publishing, which gives GLOBAL-SITE Controller customers the option to trade content versioning for improved performance in propagating content.

When creating a new publication, you can now choose whether the publication will be archived. Archived publications include only archived sections. For non-archived sections, the GLOBAL-SITE Controller does not create or keep versions of the file contents, but it does create and keep lists of file dates and sizes so that it can detect changed files. With these lists, the GLOBAL-SITE Controller can store enough information to determine when files have been changed, added, or deleted. See *Creating non-archived publications*, in Chapter 3 for details.

Scheduled publishing

GLOBAL-SITE Content Controller now supports scheduled delivery and activation of new content.

You can set up publishing to occur at regular, pre-determined times and at specific timed intervals. You can also start the publishing process manually, and have it delay until a later, pre-determined time. For details, see *Scheduling the publishing process* in Chapter 3.

Scheduled activation of new content

Along with the ability to schedule the publishing process, the GLOBAL-SITE Controller now allows you to determine how and when the new content will be activated at the subscriber for the customer to see. For details, see *Scheduling the publishing process* in Chapter 3.

Improved status and error reporting

GLOBAL-SITE Controller version 1.1 supports improved status and error reporting, where you can select the type of error reporting you want, and choose to send email alerts, notifying someone of publishing errors. In addition, you now have the ability to compare publication editions (see the Publication Editions screen online help) for archived publications, and to compare section version differences for both archived and non-archived publications (for details, see the online help for the Section Detail screen, Section History tab).

Multi-disk support

Version 1.1 increases the reliability of the GLOBAL-SITE Controller by adding a second disk. This helps protect against controller down time and data loss in case of disk failure, and improves performance by providing two disks to share the work. During installation, the First-Time Boot utility prompts you to choose whether to use the second disk for backup (mirroring the disks for a lesser amount space with error recovery) or increased storage capacity (appending the disks for a greater total of space).

There are different configurations of the GLOBAL-SITE Content Controller; check with your vendor for disk size and total capacity.

First-Time Boot utility

The GLOBAL-SITE Content Controller installation is accomplished through the First-Time Boot utility, a wizard that guides you through a brief series of tasks to set up the GLOBAL-SITE Controller.

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.

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 Publishing, Distribution, BIG-IPs, and System, 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

Table 1.1 Contacting F5 Technical Support

2

Publishing Documents

- Why GLOBAL-SITE Controller?
- Planning to configure the GLOBAL-SITE Controller
- Using the publishing cycle

Why GLOBAL-SITE Controller?

The GLOBAL-SITE Controller is used to manage the process of accumulating a logically consistent set of content, and copying it to a set of content servers in a controlled way.

The basic process involves:

- ❖ Capturing files from one or more sources
- ❖ Storing them on the GLOBAL-SITE Controller for archived publications
- ❖ Distributing those files to a set of content 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.

Understanding GLOBAL-SITE objects

In order to follow the basic GLOBAL-SITE Controller processes, it is important to understand the fundamental GLOBAL-SITE elements. You should also be aware of 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. In archived publications, 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 used in archived publications.

- ❖ An **archived** publication stores a copy of each file for each edition on the GLOBAL-SITE Controller. In **non-archived** publications, the GLOBAL-SITE Controller stores no content, but keeps a version list of dated files.
- ❖ A **distributor** is a remote GLOBAL-SITE Controller that provides access to additional remote content servers. A **publisher** can deliver publications directly to content servers, known as **subscribers**, or deliver to a distributor which then delivers the publication to the subscribers.

It may help you to consider the GLOBAL-SITE elements as relating to a publication like a newspaper.

A **publication** is a set of content like a newspaper or online magazine. Publications are made up of sections.

A **section** is a set of files that stores related content. For example, a section could be an online news article that consists of two files: an HTML file that contains the article text, and a GIF file that contains a related picture. Together these files make up the section that the end user views as a single article.

In archived publications, sections have **versions** (identified with numbers) associated with each file included in the section. For example, you publish a news story section, and then an hour later you review more information about the story. You can revise the HTML file, and publish it as a new version of the section. (In non-archived publications, there are no version numbers, but a date and time show the last time the section was retrieved, and you simply republish the section.)

Archived publications also have **editions**, similar to newspaper editions. In the context of the GLOBAL-SITE Controller, an edition of a publication is a set of sections with specific version numbers.

A **subscriber** is the person or business to which the publication is delivered. To publish or distribute to subscribers, you must know their location, and how to get there.

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 to plan what you are going to do, and how you want to do it. This helps save time since you can initially set things up in the way that will be most efficient for you in the long run.

You need to define:

- ❖ Where you are going to look for new content
- ❖ Where you are going to put it
- ❖ Whether you want to create archived or non-archived publications

You can use any file as content as long as you can specify a path to the server and directory that stores the file, and a method of access for the file. The GLOBAL-SITE Controller collects the content from this location, and (for archived publications), stores it as a version of a section. For each archived section, the GLOBAL-SITE Controller stores two things: it stores the content of each file included in the section, and it stores the access information for the section.

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 content servers, referred to as subscribers. However, before you can publish the content, you need to know which servers subscribe to the content, 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***.

For your planning purposes, you need to understand the basic difference between archived and non-archived publications.

- ❖ **Archived publications** allow you to revert to a previous edition or section version by storing copies of the content on the GLOBAL-SITE Controller.

- ❖ **Non-archived publications** provide a faster content replication and delivery, but do not provide versioning abilities.

Remember that archived publications and non-archived publications are mutually exclusive. Archived publications have only archived sections. You can create non-archived sections only within a non-archived publication. You cannot change an archived publication to a non-archived publication, or vice versa, except by first recreating the publication with a new name, and then deleting the old one. For more details about non-archived publications, refer to *Creating non-archived publications* in Chapter 3.

Defining publications

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 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.

Each publication defines two major relationships:

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

The publication organizes the sections you have defined. It is this publication that is delivered, to the 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).

A publication also has a significant attribute that you determine when you create the publication: it can be archived or non-archived. This affects how the publication is delivered, and whether or not versions are stored. Note that, the screens you work with while managing your publication have subtle differences based on the type of publication you are working with.

This chapter of the Administrator Guide deals primarily with archived publications, but you can find more details about non-archived publications in *Creating non-archived publications* in Chapter 3, and in the online help.

Creating the publication

You need to assign a unique name to each publication. This name is a label, and identifies the publication in the various work screens through the browser interface.

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 browser 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. To define how the publication delivery process works, you can set or change the publishing, scheduling, and error handling options now, or later.

To add a publication

Open the New Publication screen.

1. For archived publications, such as the example, use the default setting, where the **Archive publication editions and section versions** check box is checked. For a non-archived publication, clear the check box.

- In an archived publication, the GLOBAL-SITE Controller maintains copies, or editions, of your publications. You can roll back your subscriber content to a previous version, and retrieve old files if necessary.
 - A non-archived publication takes content directly from the source and moves it directly to the subscriber. It does not keep copies of the content on the GLOBAL-SITE Controller; it does not create editions.
2. In the **Name** box, type the name of the publication you are adding.
 3. In the **Description** box, type a description. This is optional, for your own use.

This is the minimum that you have to do to create a new publication. You could click the **Add** button now, and have a new publication. However, if you want to add scheduling and error handling options to the publication, continue with the remainder of the steps here.

Note that you do not have to decide on the publication options now. You can set them at a later time with the Publication Options screen. For more information about publishing delays, scheduling, and error handling options, refer to the online help for the New Publication or Publication Options screen.

4. In the Initiate Publishing Process area, indicate whether you want publishing to be initiated manually, or on a schedule. Accept the default **Manually only** or click the **Scheduled** button.
 - If you want to manually initiate scheduling, either accept **No default delay**, or click the **Wait till ____ by default** button, and set a time.
 - If you want your document publishing to be scheduled, select the days and time you want for publishing.
5. In the Activating New Content area, either:

- Accept the default for **Controlled Activation**, and select options for delayed publication and the proportion of subscribers activated at one time.
 - Click the button for **Independent Activation**.
6. If you are using BIG-IP Controller's virtual servers, select the options you want under BIG-IP Virtual Server Control.
 7. In the Error Handling area, select how you want the GLOBAL-SITE Controller to deal with non self-correcting errors, and specify an email address for error alerts.
 8. Click the **Add** button to complete the addition. The GLOBAL-SITE Publication Sections screen displays.

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




Figure 2.2 The Publication List screen.

9. Click the publication name (**WebPub1** in this example) in this list to view the Publication Editions screen for details about 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. In our example, and for all archived publications, 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. For non-archived publications, you first see the Publication Sections screen, and will never see the Editions tab.

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 tabs. You can click any tab 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. Note that if you are working with a non-archived publication, you will not see the Editions tab.
- ❖ 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 scheduling, error handling, and virtual server controls.

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 content files and a path to the content files.

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

- ❖ You have a server (subscriber); **what content needs to be managed?**
- ❖ **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 in a given path. A

section can stop at any specified subdirectories, known as **exceptions**. Exceptions are directories that are in the path, but are not included in the section.

❖ **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.

❖ **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.

❖ **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 the content for a section 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 needed to access the server.
- ❖ Define the location of the content with a path.

The section you create will be archived, or non-archived, depending on the publication it is part of. You do not have to specify this, but you should remember that a section carries the attribute of the publication it is created for, and cannot be shared between archived and non-archived publications.

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, Figure 2.4, displays.
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.

Add section to publication: [WebPub1](#)

The screenshot shows a web form titled "Settings" for adding a new section to a publication named "WebPub1". The form includes the following fields and controls:

- Section Name:** A text input field containing "sec1".
- Server:** A text input field containing "myserver.f5.com".
- Description:** A text input field containing "test server for internal website".
- Enabled:** A checkbox that is checked, with the label "Enable this section for the publication WebPub1".
- Transfer Method:** A dropdown menu currently set to "FTP".
- User ID:** A text input field containing "website".
- Password:** A text input field with masked characters (asterisks).
- Path:** A text input field containing "/home/webcontent/testsite".
- Buttons:** "Revert" and "Create" buttons are located at the bottom right of the form.

Figure 2.5 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 quotation marks (double quotes), apostrophes (single quotes), slashes, or angle brackets. You cannot change the name of a section later. For our example, type **WebContent1**.

2. In the **Server** 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 (optional).
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 not include it in the edition, 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 that the GLOBAL-SITE Controller can access, via FTP, the server that stores the content. Test Connection verifies that the user ID, password, and path information are valid, by connecting to the specified system.

To test your section connection

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.

Publication: [WebPub1](#)
Test Connection for: [sec1](#)

Retrieval Access Test Results	
Contacting myserver.f5.com via FTP:	Success
Authenticating as gsite:	Success
Paths:	
/home/webcontent/testsite	Success

Figure 2.6 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, click the **Get New Version** button.

The GLOBAL-SITE Controller goes through the process of gathering content from the specified location and storing an initial version of it.

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

Modifying your section configuration later

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

To change your section parameters

1. From any Publication detail screen, click the Sections tab.
The Publication Sections screen displays (see Figure 2.4).

The Publication Sections screen lists all sections for a publication.
2. Click the name of the section you want to work on.

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

Publication: [WebPub1](#)

Section: **sec1**

Section Status: Idle

Test Connection

Get New Version

Source

Section History

Current Settings

Server: myserver.f5.com

Description: test server for internal website

Enabled: ☒ Enable this section for the publication "WebPub1".

Publications: WebPub1

Transfer Method: FTP

User ID: website

Password: ****

Path: /home/webcontent/testsite

Exceptions:

New Exception:

Revert

Save

Figure 2.7 The Section Details screen, Source tab

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

Loading initial content into sections

At this point in our example, or if you are working with an archived publication, you can update the content for a section from its source, or you can create an edition. In non-archived publications, the publishing process updates the sections.

To update the content of a section

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

Publication: **WebPub1**




Figure 2.8 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 Publish Progress Display screen displays, where you can monitor the progress of that process.

This is discussed in more detail in the section *Displaying status*, later in this chapter.

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

Updating: [WebPub1](#)
 Publishing Started: Tue Mar 28 15:32:01 2000
 GLOBAL-SITE Free Space: 239MB/988MB (24% available)
Status: UpdatingSections

Delivery Control:



[Continue](#) [Cancel](#)

Section	Status	Version	Files	MBytes	MB put in GLOBAL-SITE
sec1	Updating	0	0/0	0.00/0.00	0.00
sec2	Updating	0	0/0	0.00/0.00	0.00

[Refresh](#) [Show Logs](#)

Figure 2.9 The Publish Progress Display screen

Specifying the subscribers for section content

A **subscriber** is a 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.

In the example so far, you have defined your web site as a group of sections, and created a single publication to manage them as a group. Now you have to specify which servers will receive the publication.

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.

Specifying a subscriber

1. From any of the Publication detail screens, click the Subscribers tab.
The Publication Subscribers screen displays.
2. Click the **Add a Subscriber** button.
The Add Subscriber screen displays.

Add subscriber to publication: [WebPub1](#)

New Subscriber Settings	
Subscriber Name:	WebServer1
Server:	www.f5.com
GLOBAL-SITE Distributor:	None
Transfer Method:	FTP
User ID:	website
Password:	XXXXXXXXXX
Include Section	Destination Path
☒ sec1	/wwwroot/sec1
☒ sec2	/wwwroot/sec2
Revert Add	

Figure 2.10 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 you can only select currently defined distributors. For more information on setting up a GLOBAL-SITE distributor, see the following section in this chapter, *Defining a distributor*.

4. Leave the **Transfer Method** set at FTP.
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 want to use a secondary GLOBAL-SITE Controller to distribute to remote subscribers, you need to set up a distributor. Use the Add a Distributor screen to set up and add a new GLOBAL-SITE distributor. (In the navigation pane, click **Add a Distributor**.)

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 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 publisher and distributor pair. 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.

Once you have your distributors set up, you can easily specify a distributor for any subscriber that you add. Use the Add Subscriber screen. (To get to the Add Subscriber screen: from the Publication List screen select the publication you are working with, on the next Publication screen click the Subscribers tab, on the Publication Subscribers screen, click the **Add a Subscriber** button.)

To add a distributor with a subscriber

1. Add the subscriber as you did earlier in *To add a subscriber*.
2. In the **GLOBAL-SITE Distributor** box, select the distributor that you want to use for this subscriber.
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

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

Publication: [WebPub1](#)

Test Connection for: [WebServer1](#)

Delivery Access Test Results	
	Contacting www.f5.com via FTP: Success
	Authenticating as website: Success
Paths:	
	/wwwroot/sec1 Success
	/wwwroot/sec2 Success

Figure 2.11 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 area 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 subscribers (servers).

Within the publication cycle, there are two independent activities:

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

These two activities are typically executed sequentially, but they do not have to 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.

For archived publications, there are four ways to update a section's content at any time; all open the Publish 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. On the Create a New Editions screen, select the **Get a new version** option from the version list box, and click the **Create Edition** button.
- ❖ The fourth is to click the **Deliver Edition** button from the Publication Editions screen. On the Deliver screen, and select **Create new Edition** from the **Edition** box, and click the **Deliver Edition** button. (This also publishes the newly retrieved files.)

In all of these cases for an archived publication, 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 most recent version 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, it does not create a new version.

The Section History tab on the Section Detail screen lists all versions currently stored in the section. (This tab does not appear on the Section Detail screen for non-archived publications.)

Publication: [WebPub1](#)
 Section: **sec1**
 Section Status: Idle

Test Connection Get New Version

Source Section History

Show the differences between versions and

Name	Version	Creation Date	Delete
sec1	2	Tue Mar 28 15:32:05 2000	
sec1	1	Wed Mar 8 19:23:30 2000	

Figure 2.12 The Section Detail screen, Section History tab

In the example, we have already stored two versions. You can view the contents of each version by clicking the number under the **Version** heading.

Keep in mind that multiple publications can share a single section. For instance, an image could be used by numerous publications.

Creating the very first version of a section may take some time, as all of the content 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.

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 for archived publications 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.

For details of the non-archived publication process, and of the enhanced scheduling options, refer to the online help, or to *Scheduling the publishing process*, and *Creating non-archived publications* in Chapter 3.

Creating 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 Edition** button.

The Create a New Edition screen displays.

Create new edition of publication [WebPub1](#)

Section	Version
sec1	2: Tue Mar 28 15:32:05 2000 ▾
sec2	2: Tue Mar 28 15:32:05 2000 ▾

[Create Edition...](#)

Figure 2.13 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 a New Edition 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. Or click the version number and date for a section to view the Section File Listing screen for that version of that section.

Displaying status

Some of the publishing phases can be lengthy, and are therefore designed as asynchronous processes. So that you can follow the process, the overall publication status is always displayed wherever the publication is listed. On the Publication detail screens (Editions, Sections, Subscribers, Options), for example, you can see that it is currently **Idle**. You can also see that in the Publication List screen. Click the publication status in either of those screens to display the delivery details screen.

Publication: **WebPub1**

Status: **Idle** Deliver Edition... Show Details

Editions
Sections
Subscribers
Pub. Options

Create Edition...
Compare Editions

Edition	Last Published	Remove
▼ 1: Tue Mar 28 15:53:18 2000	not yet	
Section	Version	
sec1	2: Tue Mar 28 15:32:05 2000	
sec2	2: Tue Mar 28 15:32:05 2000	

Figure 2.14 The Publication Editions screen showing **Idle** status

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

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

1. Click the **Deliver Edition** button, for archived publications, or the **Deliver** button for non-archived publications.
The Deliver screen displays.

Deliver publication: **WebPub1**

Edition: 1- Tue Mar 28 15:53:18 2000 Start Delivery: Now

Show Delivery List
Activate New Content: As soon as possible

Include	Section	Path	Current Edition
WebServer1 (www.f5.com)			
☒	sec1	/wwwroot/sec1	None
☒	sec2	/wwwroot/sec2	None

Deliver Edition...

Figure 2.15 The Deliver screen

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

- 2. Click **Deliver Edition**. This starts the delivery process. For non-archived publications, the Deliver screen shows **Deliver** buttons, which have a similar function. The Publish Progress Display screen displays.




Figure 2.16 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 2.16 shows details from an early phase in the process, and the details are about the sections. Figure 2.17 shows details from a later phase of the process, and those details are about the subscriber path.

Publishing: [WebPub1](#)

Publishing Started: Tue Mar 28 15:57:26 2000

GLOBAL-SITE Free Space: 239MB/988MB (24% available)

Status: PostPublishPhaseDone

Delivery Control:



[Continue](#)

Host	Section	Status	Version	Files	MBytes
Distributor: python.dev.net					
www.f5.com	sec1	gsDataCommitDone	-1 -> 2	0/0	0.00/0.00
www.f5.com	sec2	gsDataCommitDone	-1 -> 2	0/0	0.00/0.00

[Refresh](#)

[Show Logs](#)

Figure 2.17 Publication Progress Display; Status: **CommitPhaseDone**

Using 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 does not show as **Idle**, various control buttons are displayed on the detail status screens. The cancel button function is crucial.

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**.

If you click the **Cancel** delivery control button 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** button .

You can also try again for the subscriber sections that failed by clicking the **Retry** button . You can stop the publishing with the **Reset** button , which stops the process and returns the publication status to **Idle**.

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 3.1, *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.

3

Increasing Control over Publishing

- Using enhanced publishing features
- Scheduling the publishing process
- Working with BIG-IP Controllers and virtual servers
- Creating non-archived publications
- Defining exceptions to sections
- Using basic publishing functions
- Making ongoing configuration changes
- Understanding the GLOBAL-SITE Controller phases

Using enhanced publishing features

Once you have read through instructions on how to set up a basic archived publication, you will be ready to create your own publications. You may want to add some additional control to the publishing process, or you may want to create a non-archived publication. The GLOBAL-SITE Content Controller provides several features that help you in administering and managing your various publications. Some features you can use to augment basic publishing include:

❖ **Scheduling the publishing process**

Scheduling the delivery and activation of content on the subscribers allows you more freedom, offers additional control, and provides a consistent delivery of content to your subscribers.

- Scheduled initiation of publishing provides options for daily, weekly, or continuous delivery of your publications.
- Scheduled activation of new content offers opportunities for immediate or delayed activation after copying new content to the subscribers, and for activating content on all subscribers at once, in groups, or independently.

❖ **Working with BIG-IP Controllers and virtual servers**

The integration of the GLOBAL-SITE Controller with the BIG-IP Controller provides several benefits to controlling and maintaining site performance during content update. This feature gives you the ability to restrict web access to subscribers in conjunction with a BIG-IP Controller. This reduces the risk of exposing a possibly inconsistent mix of old and new content to web users during the update process.

❖ **Creating non-archived publications**

Non-archived publication gives you the option to trade content versioning for improved performance in propagating content. Non-archived sections do not leave versions of the file contents on the GLOBAL-SITE Controller, but they do create version lists of dated files so that the GLOBAL-SITE Controller can detect file changes. Using non-archived publications tends to provide a somewhat faster delivery of content to subscribers.

❖ **Defining exceptions to sections**

Creating exceptions to sections provides you with additional flexibility in designing and implementing your publications. Exceptions allow you to exclude a subdirectory from a section when publishing the section.

Scheduling the publishing process

One of the things that can save you time, and add reliability to your updates, is the ability to schedule the publishing of your publications. GLOBAL-SITE Content Controller supports both scheduled delivery and activation of new content:

- ❖ You can set up publishing to occur at regular, pre-determined times and at specific timed intervals.
- ❖ You can determine how and when the new content will be activated at the subscriber for the customer to see.

To understand the full extent of the GLOBAL-SITE Controller scheduling features, you should understand a bit about how publications are delivered to the subscribers. In most cases, publishing happens in two phases: copying content from the source to the subscribers, and then activating the new content on the subscriber. When the GLOBAL-SITE Controller delivers the content to the subscribers, it is initially in a temporary location, and not yet available to the viewers. Activation moves this content to its permanent location on the subscriber and makes the content available, ready to be accessed.

Scheduled publications

Setting up schedules for your publications can help automate the publishing of new content on a regular basis, make your update process less dependent on human intervention, and add a degree of regularity to the publishing process. The process of creating scheduled publications is simple, and can be done when you initially create the publication, or later in the lifetime of the publication.

Reviewing scheduling features

The GLOBAL-SITE Controller offers several features and choices for scheduling the publishing process. You can provide as much automation for your publishing as you want, and you can always pause the schedule for a publication, or revise the settings at a later time. Some of the scheduling features include:

- ❖ The continuous update process sends the changes to sections and publications to your most out of date subscribers first, and continues the cycle until all out of date subscribers are updated.
- ❖ You can set the schedule to run, or not run, for specific days of the week.
- ❖ You can make changes to a publication only when the publication is in an idle state. However, you can pause the schedule at any time. If the publication is not yet publishing, you can immediately make configuration changes. If the publication is publishing, it continues its current process until done, and then pauses its schedule. When the schedule pauses, you can make changes.
- ❖ On a scheduled publication, you must pause its schedule before you can do manual deliveries.
- ❖ Scheduled updates automatically publish to out-of-date subscribers (which may be out of date because the last publishing process did not complete).
- ❖ You can set delays for starting the publishing process and activating new content for all manual deliveries. Default times are configured and editable in the Publication Options screen.
- ❖ If a publishing process runs longer than the scheduled interval, the schedule ignores all missed update times except the most recent, and automatically begins the missed update as soon as it finishes the current update. For instance, on a schedule with 5-minute interval updates, if an update takes 18 minutes, the scheduler will miss three update intervals (one at 5-, one at 10-, and one at 15-minutes) during the 18-minute process. But after completing the 18-minute process, the scheduler will start again without delay in response to the immediately preceding scheduled interval (the one missed at the 15-minute mark).

- ❖ Once you have set the scheduling features for a publication, use the System Settings screen to turn on the GLOBAL-SITE publication scheduler so that scheduled publishing can take place. This screen also allows you to pause and restart scheduling for the entire GLOBAL-SITE system.

Using scheduling options

Most of the settings for scheduled publications and activation of content are made in the New Publication or Publication Options screen. For procedural details, see the online help for those screens.

Configuration Utility GLOBAL-SITE: New Publication

GLOBAL SITE

Publishing
List Publications
Add a Publication
List Sections

Distribution
List Distributors
Add a Distributor

BIG-IPs
List BIG-IPs
Add a BIG-IP

System
Show Settings
List Users
Add a User

New Publication Information:

☒ Archive publication editions and section versions

Name:

Description:

Initiating Publishing Process

☒ Manually only

☐ No default delay

☐ Wait till by default

☐ Scheduled

Su Mo Tu We Th Fr Sa

☐ ☐ ☐ ☐ ☐ ☐ ☐

☐ Every minutes

☐ Daily at

Activating New Content

☒ Controlled Activation

Before activating manually published content:

☒ No default delay

☐ Wait for user input

☐ Wait till by default

of subscribers activated at the same time

☐ Activate through one distributor at a time

☐ Independent Activation

BIG-IP Virtual Server Control:

☐ Disable virtual server nodes while activating content

☐ Time out persistent connections after seconds

Error Handling

When an error occurs during the publishing process:

☐ Continue publishing after logging the error

☐ Terminate publishing and log the error

☒ Pause publishing, log the error and wait for user input

Send e-mail notification to:

Add

Figure 3.1 The New Publication screen

The scheduling options available for initiating the publishing process include:

- ❖ Manual initiation of the publishing process to begin immediately, or at a later pre-determined time
- ❖ Publishing that runs automatically one or several days a week
- ❖ Publishing that occurs at regularly timed intervals
- ❖ Publishing at one or more pre-set times during the day

Of course, you can override manual publication scheduling settings when you deliver the publication or edition. Use the Deliver screen setting to change the time, or choose to deliver **Now**.

Deliver publication: [WebPub1](#)

Edition:

Start Delivery:

Show Delivery List

Activate New Content:

Include	Section	Path	Current Edition
WebServer1 (www.f5.com)			
☒	sec1	/wwwroot/sec1	None
☒	sec2	/wwwroot/sec2	None

Deliver

Now
1:AM
2:AM
3:AM
4:AM
5:AM
6:AM
7:AM
8:AM
9:AM
10:AM

Figure 3.2 The Deliver screen showing the Start Delivery options

On the all the Publications detail screens, you have the option to pause and resume scheduled publishing for scheduled publications.

After scheduling your publications, you may want to look at the System Settings screen, where you can check, or set the GLOBAL-SITE Controller system time, and stop or start the GLOBAL-SITE scheduler. (Check the online help on that screen for details.) You may also want to look at the scheduled activation options that are available on the Publication Options and New Publication screens.

Scheduled activation of new content

Content activation is one of the final steps in publishing: it occurs after the copying phase, and is when the GLOBAL-SITE Controller changes the temporary files on the subscriber into permanent, active files that customer viewers can access.

The GLOBAL-SITE Content Controller provides a number of controls over the content activation process. These controls fit into two categories:

❖ **Independent activation mode**

This new mode is useful for a publication with more than one section, when it is important that new content from each section is activated on a subscriber as soon as possible, regardless of whether other sections are updated. (In controlled activation, all sections are copied before the activation phase begins.)

❖ **Controls for the controlled activation mode**

With these controls you can activate new content on groups of subscribers, rather than on all subscribers at once. Activating some subscribers instead of all at one time minimizes the disruption of the complete web site during content update.

Whether your publication is working on a schedule, or manually published, you can determine how the new content will be activated on the subscriber. For manually initiated delivery of content, you can delay the activation of new content until a predetermined time, or wait and give the go-ahead manually.

Publication: **WebPub1**Status: **Idle**[Deliver Edition...](#)[Show Details](#)

Editions	Sections	Subscribers	Pub. Options
Rename			
Publication Information:			
Name: WebPub1 Description: This is where you can type in a complete description of the publication			
Initiating Publishing Process		Activating New Content	
☒ Manually only ☒ No default delay ☐ Wait till 1:AM by default		☒ Controlled Activation Before activating manually published content: ☐ No default delay ☐ Wait for user input ☒ Wait till 1:AM by default 100% at the same time ☐ Activate distributor at a time ☐ Independent	
☐ Scheduled Su Mo Tu We Th Fr Sa ☒ ☒ ☒ ☒ ☒ ☒ ☒ Every 20 minutes ☐ Daily at: 1:AM		☐ BIG-IP Virtual ☐ Disable virt ☐ Time of connections after 120 seconds	
Error Handling When an error occurs during the publishing process: ☐ Continue publishing after logging the error ☐ Terminate publishing and log the error ☒ Pause publishing, log the error and wait for user input Send e-mail notification to:			

Figure 3.3 Use the Publication Options screen for activating new content

You use the New Publication or Publication Options screen for settings that determine how content is activated as part of the publishing process. The two major choices are:

- ❖ Independent activation of sections on subscribers

- ❖ Controlled activation of all subscribers, or a group of subscribers all at once

Controlled activation includes additional options:

- Activating manually published content with a delay, with user input, or at a particular hour
- Activating subscribers all at once, or half of them, or a quarter of them at a time
- Activating subscribers through one distributor at a time

The Deliver screen also provides the opportunity for independent activation of sections on subscribers.

The Publishing Progress Display screen provides the **Continue** button so you can by-pass a publishing initiation or activation delay.

Comparing controlled and independent activation

In order to understand the independent activation features, it is important to understand the controlled activation process.

- ❖ **In controlled activation**, the publishing process moves all the publication together through the process strictly one phase at a time. A new phase is not started until the previous phase is complete. Therefore, a section's content cannot be activated on a subscriber until all other sections' contents had first been copied to the subscriber.
- ❖ **Independent activation** allows each section of each subscriber to independently move into the activation phase as soon as that section's copy phase is completed.

Imagine that you have files that are being made available for download from different sources. You may want the files to be made visible to the customers as soon as they are ready, rather than waiting until all of them are ready to become available. In this case, it is acceptable to activate files for a given section as soon as they are copied to a subscriber. You would want to use independent activation.

Both controlled and independent modes of activation are available for either manual or scheduled publications.

Options for controlled activation

With version 1.1 of the GLOBAL-SITE Controller, you can control how subscribers are activated. You will find these options on the New Publication and Publication Options screens under **Activating New Content**. (For details on how to use these options, see the online help on either screen.) The default for activating is that the GLOBAL-SITE Controller tries to activate all subscribers simultaneously.

You can change the default behavior, and specify that only some of the subscribers should be activated together. In addition to all subscribers, you can choose one half, or one quarter. This option is most useful in conjunction with control of the BIG-IP Controller.

When you are accessing some of the subscribers via a distributor, you can use the **Activate through one distributor at a time** check box to specify that the GLOBAL-SITE Controller should activate subscribers through one distributor at a time. If you use this option, the subscriber activation percentage specifies what fraction per distributor to activate together.

Working with the BIG-IP Controllers and virtual servers

One of the more powerful additions to the GLOBAL-SITE Controller for version 1.1 is the integration of functionality with the BIG-IP Controller. Typically, BIG-IP Controllers are used to load balance an array of servers. You can increase your publishing scope, and control over your publications if you set up a virtual server subscriber on the BIG-IP Controller, and use the GLOBAL-SITE Controller to control it and the nodes behind it as a single entity.

When the GLOBAL-SITE Controller accesses a virtual server subscriber on a BIG-IP Controller, you have the option to maintain the subscriber site performance by removing the virtual server nodes from user access while updating the contents of the site. The restriction of web access to subscribers in conjunction with a BIG-IP Controller reduces the risk of exposing a possibly inconsistent mix of old and new content to web users during the update process.

The GLOBAL-SITE Controller Administrator Guide does not provide details on BIG-IP Controller functionality. You may want to refer to your BIG-IP Controller user guides for details of setting up the BIG-IP Controller and virtual server nodes.

Use the GLOBAL-SITE Controller browser interface to:

- ❖ set up a BIG-IP Controller as a virtual server
- ❖ remove unused BIG-IP Controllers
- ❖ change settings for current controllers
- ❖ mark BIG-IP Controller nodes as down during publishing

Guidelines for BIG-IP Controller GLOBAL-SITE Controller cooperation

Before you create a virtual server subscriber on the GLOBAL-SITE Controller, you will want to ensure that you have followed these guidelines:

- ❖ The GLOBAL-SITE Controller must be configured on the same BIG-IP interface as the virtual server nodes. This is typically the internal interface (the network interface through which a BIG-IP Controller distributes connections). In newer releases of the BIG-IP Controller, that interface must have **admin** access.
- ❖ The BIG-IP Controller must have the **bigorb-server** installed for the GLOBAL-SITE Controller to talk to. (Your vendor may install it for you.) For security reasons, the **bigorb-server** responds only to requests from a GLOBAL-SITE Controller on an **internal** or **admin** interface.

- ❖ The GLOBAL-SITE Controller must have FTP access to the virtual server nodes using the same node IP address that the BIG-IP Controller has configured for the virtual server that the GLOBAL-SITE Controller will in turn use to create a virtual server subscriber.
- ❖ You need to define a virtual server on the BIG-IP Controller for port 2792 (for GLOBAL-SITE Controller to GLOBAL-SITE Controller communication, using the F5-GLOBAL-SITE protocol) if it will be used as a distributor for another GLOBAL-SITE Controller.
- ❖ If the GLOBAL-SITE Controller (that you are using with the BIG-IP Controller) is to be used as a publisher, you must configure on the BIG-IP Controller a virtual server to it for port 443 (HTTPS) for access to the web browser. You will probably also need a secure NAT (SNAT) for the GLOBAL-SITE Controller to get FTP access to systems via the BIG-IP Controller, such as section sources.
- ❖ Once you have completed the preceding tasks, you need to inform the GLOBAL-SITE Controller about the BIG-IP Controller via the Add a BIG-IP screen, which will also indicate whether communication works or not.

Creating a BIG-IP Controller on the GLOBAL-SITE Controller

To add a new BIG-IP Controller to the GLOBAL-SITE system, use the Add a BIG-IP screen (from the navigation pane, click Add a BIG-IP). You can add a BIG-IP Controller only if the BIG-IP Controller is up and available locally or through a distributor.

Current BIG-IP Settings:

BIG-IP Name:

BIG-IP Description:

GLOBAL-SITE Distributor:

BIG-IP Virtual Servers:

192.168.15.143:20:21:23:80:443:2792
192.168.15.144:20:21:23:80:443

Update

Figure 3.4 The Add a BIG-IP screen

To add a BIG-IP Controller

1. Type in the name of the BIG-IP Controller that you want to add. This is the internal IP address, or alias of the BIG-IP Controller.
2. Type in a description for the BIG-IP. This identifier may make it easier for you to recognize each controller.
3. Select a BIG-IP distributor from the list of available controllers.
4. Click the **Add** button to add this BIG-IP distributor. The BIG-IP List screen displays.

Once you have added the BIG-IP Controller, you can see it listed on the BIG-IP List screen, where you can delete any controller that has no references. You can also review or change the BIG-IP Controller settings on the BIG-IP Detail screen. If the BIG-IP Detail screen does not show any virtual servers, then the BIG-IP Controller is set up incorrectly.

Current BIG-IP Settings:

BIG-IP Name:

BIG-IP Description:

GLOBAL-SITE Distributor:

BIG-IP Virtual Servers:

192.168.15.143:20:21:23:80:443:2792
192.168.15.144:20:21:23:80:443

Figure 3.5 The BIG-IP Detail screen

Setting up a virtual server subscriber

If you have one or more BIG-IP Controllers set up on the GLOBAL-SITE Controller, you have the option of setting up a virtual server subscriber. If the GLOBAL-SITE Controller is not aware of any BIG-IP Controllers, your option is to set up a normal server subscriber. To set up any server, start with the Publication Subscriber screen.

Publication: **WebPub1**

Status: Idle

Add subscriber type:

Name	Current Edition	Server	Status
 WebServer1	1	Virtual Server	Idle 

Figure 3.6 Selecting a virtual server subscriber on the Publication Subscriber screen

To add a new subscriber

1. If it is visible, from the **Add subscriber type** box, select the type of subscriber you want to add, Server or Virtual Server.
2. Click the Add a Subscriber button.
Depending on the type of server you are adding, either the Add Subscriber screen or Add Virtual Server Subscriber screen displays.
3. Fill in the form with subscriber access details, sections to include, and section paths.
4. Click **Add**.

Marking BIG-IP Controller nodes as down while updating content with the GLOBAL-SITE Controller

When you publish content to a virtual server subscriber, the GLOBAL-SITE Controller can have the BIG-IP Controller disable web access to the individual nodes of the virtual server while the GLOBAL-SITE Controller is activating the content. This prevents web clients from seeing a mix of old and new content. Once the content has been successfully activated, the BIG-IP Controller is instructed to re-enable web access. Use the BIG-IP Virtual Server Control options on the Publication Options screen for these settings.

Publication: WebPub1

Status: Idle Deliver Edition... Show Details

[Editions](#)
[Sections](#)
[Subscribers](#)
[Pub. Options](#)

[Rename](#)

Publication Information:

Name: WebPub1

Description:

Initiating Publishing Process

☒ Manually only

☒ No default delay

☐ Wait till by default

☐ Scheduled

 Su Mo Tu We Th Fr Sa

☒ ☒ ☒ ☒ ☒ ☒

☒ Every minutes

☐ Daily at:

Activating New Content

☒ Controlled Activation

 Before activating manually published content:

☒ No default delay

☐ Wait for user input

☐ Wait till by default

 of subscribers at the same time

☐ Activate through one distributor at a time

☐ Independent Activation

BIG-IP Virtual Server Control:

☒ Disable virtual server nodes while activating content

☒ Time out persistent connections after seconds

Error Handling

When an error occurs during the publishing process:

☐ Continue publishing after logging the error

☐ Terminate publishing and log the error

☒ Pause publishing, log the error and wait for user input

Send e-mail notification to:

Figure 3.7 BIG-IP Virtual Server Controls on the Publication Options screen

If you set up a virtual server subscriber, you can specify that 50% of the subscribers are to be activated together, and specify that access to them be disabled during the activation phase. Using this method, no single virtual server node will be enabled with a mix of old and new content at the same time. Because they are accessed via a

virtual server address, any single subscriber may show old pages while the first files are being updated, but will then have a new page once the new content is available.

A final option, under **BIG-IP Virtual Server Control**, reduces the likelihood of mismatched content. The **Time out persistent connections** check box under **BIG-IP Virtual Server Control** can temporarily enable simple persistence before disabling any of the nodes. This helps to lock each web user into a single node, and reduce the chances that a BIG-IP Controller will load balance the user from a system with old content to one with new content.

If you choose the **Independent Activation** mode, the BIG-IP Controller is instructed to disable and re-enable web access for each subscriber section independently. We do not recommend that you use the persistence option in conjunction with independent activation mode.

Creating non-archived publications

Non-archived publication is basic, fundamental content replication. It takes content directly from the user's source, and moves it directly to the location where it can be accessed by the customer-viewers. It does not keep copies of the content on the GLOBAL-SITE Controller.

Non-archived publication uses the same process as archived publications to determine which files have changed since the last publication. Once it determines which files have changed, it copies them new files to the subscribers, and activates them. Activating files is the process of changing them from temporary status, not viewable by users, to permanent on the subscriber, and available to the viewer.

Non-archived publication is the right choice when you are looking for speed of publishing over guaranteed backup and the ability to revert to a previous version of a publication. Because the GLOBAL-SITE Controller does not keep copies of the files, the

publishing process is not constrained by the size of the GLOBAL-SITE Controller disk, but only by the subscriber capacity.

Looking at benefits of non-archived publications

Non-archived publishing retains many of the benefits of archived publishing:

- ❖ The GLOBAL-SITE Controller still replicates only content that has changed since the last publishing cycle.
- ❖ You can cancel the process at any time.
- ❖ Files are still moved through a secure channel to all distributors.
- ❖ You can schedule publishing.
- ❖ You can configure the publishing process to pause or continue when encountering errors.

Planning for non-archived publications

For your planning purposes, remember that archived publications and non-archived publications are mutually exclusive. Archived publications have only archived sections. You can create non-archived sections only within a non-archived publication.

To change an archived publication to a non-archived publication, you need to recreate the publication with a new name, and delete the old one.

Considering the details of non-archived publishing

Before creating non-archived publications, there are a few things you might want to think about. These details can affect the efficiency and performance of your non-archived publications.

- ❖ How do you organize your files and directories?
- ❖ What is the usual size of your sections?
- ❖ What is the typical amount of activity in your files?

- ❖ What is the capacity of your subscriber server?

Organizing files and directories

Before starting to copy content to the subscriber, the GLOBAL-SITE Controller reviews all content at the source, looking for things that have changed between the new content and what the subscriber has. It does this by comparing a list of files, sizes, and dates (that is tracked by the GLOBAL-SITE Controller) to the list of files, sizes, and dates on the source. The time it takes for this comparison depends on how the content is organized. The fewer directories you have in the section, the less time it takes to complete the comparison. For instance, publishing fewer than 100 directories seems rapid, and publishing more than 100 directories takes noticeably longer. Because the important element here is not the number of files but the number of directories, you may want to structure your publications with this in mind.

Reviewing section size

During the content copy, for each file that has changed, the entire file is copied to any distributors, and then to any subscribers. (Compare this to the archived publishing process, where only the changes are sent to the distributor, and the entire file is sent to subscribers.) Each file within a section is copied sequentially from the source to the subscriber. During this process, it is the larger sections that determine how fast the copy phase can go. Therefore, it is more efficient to have roughly equal section sizes. This will improve the performance of the copy process over having several small and one really large section.

This information is provided to help you set expectations based on your own file structure and organization, and not as a suggestion that you change your publication structure.

Checking file activity

During the update process, the directory list is created and then the files are retrieved. During the time that files are being copied to the subscriber, each file is compared to how it was at the initial comparison. If the file has changed (or is changing between these

times), the GLOBAL-SITE Controller ignores it, and does not copy it to the subscriber during the current publishing cycle. The subscriber keeps the old file, which will be updated on the next publishing cycle. The reason that the GLOBAL-SITE Controller ignores it at this point, is that the GLOBAL-SITE Controller assumes the file may be in flux, and possibly incomplete.

When you set up your publication, you need to watch for very active files, and be careful to schedule publishing for times when they are most likely to be static. If you publish when these files are changing, you may find that the subscribers do not get an updated version of that file. For additional details, refer to Scheduled publications.

Assessing subscriber capacity

During the copy phase, the changed files are copied into a temporary directory of the subscriber. These files are not available to the viewers until they are activated. During this time, the subscriber server contains all the old content, and all the changed content. Therefore, the subscriber server capacity must be large enough for the old content and the changes to co-exist temporarily.

Using procedures for non-archived publications

Creating, managing, and publishing non-archived publications is very much like dealing with archived publications. There are differences on various application screens. Most noticeable is the fact that when you are dealing with non-archived publications, you do not see the Editions tab as you do when working with archived publications. And, since non-archived publications have no editions, and no versions, you will notice other, little differences on some of the other screens. These differences are noted in the online help pages for each screen.

New Publication Information:

☐ Archive publication editions and section versions

Name:

Description:

Initiating Publishing Process

☒ Manually only

☒ No default delay

☐ Wait till by default

☐ Scheduled

Su	Mo	Tu	We	Th	Fr	Sa
☒	☒	☒	☒	☒	☒	☒

☒ Every minutes

☐ Daily at:

Activating New Content

☒ Controlled Activation

Before activating manually published content:

☒ No default delay

☐ Wait for user input

☐ Wait till by default

of subscribers activated at the same time

☐ Activate through one distributor at a time

☐ Independent Activation

BIG-IP Virtual Server Control:

☐ Disable virtual server nodes while activating content

☐ Time out persistent connections after seconds

Error Handling

When an error occurs during the publishing process:

☐ Continue publishing after logging the error

☐ Terminate publishing and log the error

☒ Pause publishing, log the error and wait for user input

Send e-mail notification to:

Figure 3.8 New Publication screen for non-archived publication

To create a non-archived publication

1. Start at the New Publication screen. (From the navigation pane, click Add a Publication, or from the Publication List screen, click the **Add a Publication** button.)
2. Clear the **Archive publication editions and section versions** check box.

3. Complete the remainder of the screen and click the **Add** button. Refer to the online help for this screen if you need additional details.

Remember that you do not have to specify all options when you create the publication: you can return later and set scheduling and error-handling options.

4. The Publication Sections screen displays.
Note that the Editions tab does not display, as there will be no editions of a non-archived publication.

Publication: **non archived publication**

Status: Idle

Deliver Show Details

Sections Subscribers Pub. Options

Edit Section List Create New Section

Name	Last Retrieved
No sections are currently associated with this publication	

Figure 3.9 Publication Sections screen for non-archived publication

To create a non-archived section

On the Publication Sections screen for your non-archived publication, click the **Create New Section** button.

The Create a New Section screen displays.

Any section you create here is automatically non-archived.

◆ Note

*Note that in the Edit Section List screen, all sections listed in the **Available Sections** list are non-archived sections.*

Procedures for dealing with non-archived sections are detailed in the online help for the specific screens. Most procedures are the same for archived and non-archived, the online help pages point out any differences.

Defining exceptions to sections

After you have set up your basic publications, and created various sections, 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.

The most challenging part of setting up exceptions is planning the publication in advance, so that the sections, and section exceptions work to your benefit. Remember that, when first created, each section includes all files and directories at the given path. If you want to exclude a subdirectory from the section, you must specify it as an exception. Note that you cannot exclude files, only directories.

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**).

Publication: [WebPub1](#)

Section: **sec1**

Section Status: Idle

Test Connection

Get New Version

Source

Section History

Current Settings

Server:

myserver.f5.com

Description:

sec1

Enabled:

☒

Enable this section for the publication "WebPub1".

Publications:

WebPub1

Transfer Method:

FTP

User ID:

website

Password:

XXXXXXXXXX

Path:

/home/webcontent/testsite

Exceptions:

New Exception:

/home/webcontent/testsite/images

Revert

Save

Figure 3.10 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 **Save** 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.

Using basic publishing functions

In addition to enhanced publishing features, the GLOBAL-SITE Controller includes basic functionality that can help you in controlling and managing your publications. One of these is the ability to make ongoing changes to your publications. The other outlines the foundations of the GLOBAL-SITE publishing process. The last two sections in this chapter discuss these two aspects of publishing functionality.

❖ **Making ongoing configuration changes**

Once you have your publications set up and working well, you may find that changes in content, directories, subscribers, and servers may necessitate changes to the organization of your

publications. The GLOBAL-SITE Controller provides a variety of processes for maintaining reliable publications and efficient publishing.

❖ **Understanding the GLOBAL-SITE Controller phases**

Insight into the operations that support the GLOBAL-SITE Controller browser interface can increase your understanding of how the publishing process works, and facilitate troubleshooting of potential errors. Familiarity with the phases of the publishing process also help you grasp the significance of the status messages provided throughout the publishing process.

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.

Publication: [WebPub1](#)
Subscriber: **WebServer1**

Status: **Idle**

Test ConnectionRename

Current Edition: **None**

Current Settings

Server:

GLOBAL-SITE Distributor:

Transfer Method:

User ID:

Password:

Include	Section	Destination Path	Version	Status
☒	sec1	/wwwroot/sec1	None	Idle
☐	sec2	/wwwroot/sec2	None	Idle

CloneRevertSave

Figure 3.11 The Subscriber Detail screen

To disable a section for a particular subscriber

For each subscriber that is offline, you need to exclude all sections. Make sure the publication is in an **Idle** state. On the Subscriber Detail screen:

1. Under Include, clear the check box next to each section.
2. When you have cleared the check boxes for each section, click the **Save** 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 offline, when it returns online, you need to include all sections in order to publish to it. First, make sure that the publication status is **Idle**, since you cannot change publication information unless the publication is in an **Idle** state. Then, on the Subscriber Detail screen:

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

You can also use the Deliver screen to disable the path for a section.

- ❖ For archived publications, for each subscriber, clear the **Include** check box for each section to keep that subscriber from receiving the edition.
- ❖ For non-archived publications, under Action, select **Ignore** for each section you do not want published to that subscriber.

Refer to online help for the Deliver screen for more details.

For each subscriber that is offline, you need to exclude the subscriber from the publication, for each section, before attempting to deliver the edition.

To disable a subscriber for a particular edition delivery

For archived publications, on the Deliver screen:

1. Under each subscriber, clear the check box next to 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 are ignored by the attempt to publish this edition. Each must be returned to **Include** status before you can publish to it.

To disable a subscriber for a particular publication

For non-archived publications, on the Deliver screen:

1. Under each subscriber, in the Action box, select **Ignore** for each section listed.
2. When you click the **Deliver** button, the publication does not go to that subscriber.

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

For each subscriber that was offline, when it returns online you need to include each subscriber for each section before you can publish to it.

To enable a subscriber for an edition of a publication

For archived publications, on the Deliver screen:

1. Click a check into the **Include** check box 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 are included in the next published edition.

To enable a subscriber for a non-archived publication

On the Deliver screen:

1. In the Action column, under each subscriber, select **Deliver** for each section that you want to deliver.
2. When you have changed the box for each section that you want to include, click the **Deliver** button.

All subscriber sections that are selected with a check mark are included in the next published publication.

Checking subscribers

If you run one or more publishing cycles while a subscriber is offline, the offline subscriber almost certainly has an old version of the content. When there are out-of-date subscribers, the GLOBAL-SITE Controller automatically brings them into synchronization.

While the GLOBAL-SITE Controller is synchronizing the subscribers for a publication, you may occasionally see an entry in your Publication List that looks unfamiliar to you. These ***child publications*** are temporary, created solely to bring out-of-date subscribers back in sync.

You cannot work with, or edit, a child publication: you can only watch it complete its task and evaporate. If you click the publication link, it opens the Publishing Progress Display screen.

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).

The Publication Subscribers screen displays.

- ❖ For archived publications, it shows the currently published edition for each subscriber if all paths are up-to-date.
- ❖ For non-archived publications, it shows the last published date.

Or, for archived publications, you can use the Subscriber Detail screen, which shows the version published to each path, while the initial Deliver screen shows the edition.

Publication: [WebPub1](#)
Subscriber: **WebServer1**

Status: **Idle** Test Connection Rename

Current Edition: **1**

Current Settings

Server:

GLOBAL-SITE Distributor:

Transfer Method:

User ID:

Password:

Include	Section	Destination Path	Version	Status
☒	sec1	/wwwroot/sec1	2	Idle
☒	sec2	/wwwroot/sec2	2	Idle

Clone Revert Save

Figure 3.12 Subscriber Detail screen

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, if for instance, the path has been removed, 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 under Name, locate the subscriber you want to remove.
2. Click the delete button  to its right.



Figure 3.13 *The Publication Subscribers screen*

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

Freeing up disk space by deleting editions and unused versions

When working with archived publications, you may occasionally need to purge old versions from sections to recover disk space. You may want to base your decision on the amount of change between versions.

To purge versions

Purging versions 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 want to roll back to at some point, or to publish to other subscribers.
2. Once you have deleted all unnecessary editions, switch to the Section Detail screen and click the **Section History** tab, then click the **Remove Unused Versions** button.




Figure 3.14 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.

Understanding the GLOBAL-SITE Controller phases

When creating publications and running the publishing process from the browser interface, you will see publication status indicators. These indicators of where the publication is, or how the publishing process is going, relate to various states and phases that are more clearly seen if you use the command line utility. You need not use the command line utility, but you may find it helpful to understand the basic phases that it goes through during the entire publishing cycle. This section discusses the phases of the command line utility in relation to what you may see in the browser interface while monitoring your publishing process.

Interpreting phases in the status line

Various GLOBAL-SITE screens provide status updates on the state of publications during the processing activities (the Publication List screen is a good example). In different places, 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 combination of the section and subscriber path states).
- ❖ The detailed section or subscriber path state, depending on which object is being processed

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

During a publication delivery with no errors, the normal sequence of states follows this progression:

- ❖ **Idle**

The publication is in a stable state, either ready for updating or publishing.

- ❖ **UpdatingSections**

Sections are updated if you have specified this 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.

❖ **PostPublishPhase**

After the publishing process has completed, or failed, the GLOBAL-SITE Controller attempts to clear up any files or directories created to support publishing and, if necessary, to rollback the subscriber(s) to the previous state.

Other common phases that you are likely to see on the status line of your publication include:

❖ **Scheduled**

This indicates that the publication has been scheduled, but it is waiting to publish, and not currently publishing.

❖ **Schedule Paused**

This indicates that the publication is a scheduled publication, but the schedule has been paused. It can also indicate that the publication has been newly created, or has been newly scheduled, and the schedule has not yet been activated.

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** button.

You can also try again for the subscriber sections that failed by clicking the **Retry** button. You can stop the publishing with the **Reset** button, which stops the process and returns the publication status to **Idle**.

Using the Phases of Publication table

Table 3.1, *The GLOBAL-SITE Controller Phases of Publication*, following, describes what is happening to the publishing 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.
EditionFailed	N/A	The process has failed while creating a new edition from the latest section versions.
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 .
Scheduled	N/A	The publication is scheduled to do something (publish or activate content) at a certain time.
UpdatingSectionsIncomplete	mix of gsReposStageUpdateDone , gsReposStageUpdateFailed and Cancelled	When multiple sections are being updated, not all completed.
UpdatingSectionsFailed	gsReposStageUpdateFailed	All section updates failed.

Table 3.1 *The GLOBAL-SITE Controller Phases of Publication*

Overall Publication State	Section or Path State	Description
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. Reset 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.
WANSync	gsWanSync	Each distributor is updating its section(s) with the received changes.
WANSyncDone	gsWanSyncDone	Ready to proceed to CreateChangeTrees .
WANSyncSuspended	gsWanSyncDone	Proceeds to CreateChangeTrees after you click Continue . Retry retries the operation and Cancel cancels publishing and proceeds to PostPublishPhase cleanup. Only stops here if the corresponding Publication Option is set.
WANSyncCanceled	some canceled	Clicked Cancel during phase. Cannot proceed with delivery. Retry retries the operation. Reset takes the publication into the PostPublishPhase cleanup.

Table 3.1 The GLOBAL-SITE Controller Phases of Publication

Overall Publication State	Section or Path State	Description
WANSyncFailed	gsWanCDirFailed	An error is displayed for the specific cause, typically lack of disk space on the publisher to store the temporary change file. Retry retries the operation. Reset takes the publication into the PostPublishPhase cleanup.
WANSyncFailed	gsWanSyncFailed	An error occurred while the distributor was updating its sections, typically lack of disk space on the distributor. Retry retries the operation. Reset takes the publication into the PostPublishPhase cleanup.
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 .
ChangeTreesSuspended	gsChangeTreesSuspended	Proceeds to CopyPhase after you click Continue . Click Retry to retry the change tree; click Cancel to cancel the publishing. Only stops here if the corresponding Publication Option is set.
ChangeTreesCanceled	Canceled	The process was canceled by the user. The publication cannot proceed further. Retry retries the operation and Reset takes the publication into the PostPublishPhase cleanup.

Table 3.1 The GLOBAL-SITE Controller Phases of Publication

Overall Publication State	Section or Path State	Description
ChangeTreesFailed	gsCreateCDirFailed	An error, typically lack of disk space, will be displayed. The delivery process cannot proceed. Retry retries the operation and Reset takes the publication into the PostPublishPhase cleanup.
CopyPhase	gsDataCopy	Each GLOBAL-SITE Controller copies new and changed files to temporary subdirectories on each subscriber
CopyPhaseDone	gsDataCopyDone	Ready to move to CommitPhase
CopyPhaseSuspended	gsDataCopyDone	Will proceed to CommitPhase after clicking Continue . Retry retries the operation and Reset takes the publication into the PostPublishPhase cleanup. User selected the Publication Option to stop at this point.
CopyPhaseCanceled	Canceled plus various	User clicked the Cancel control, stopping subscriber path copies in progress, as well as those not started yet. If the Continue control is clicked, only subscribers that successfully completed the copy phase for all sections can proceed to CommitPhase . Clicking Reset sets the publication state to PostPublishPhase cleanup, and Retry retries the operation.
CopyPhaseFailed	gsDataCopyFailed	An error is displayed for the specific cause, typically communication failure or lack of disk space on the subscriber. Click Retry to retry the operation, Reset takes the publication into PostPublishPhase cleanup.

Table 3.1 The GLOBAL-SITE Controller Phases of Publication

Overall Publication State	Section or Path State	Description
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 .
CommitPhaseSuspended	gsCommitPhaseSuspended	Proceeds to PostPublishPhase after you click Continue . Click Retry to retry the commit; click Cancel to cancel the publishing.
CommitPhaseCanceled	canceled plus others	The Cancel control was clicked, which stopped all commits in progress, as well as any not yet started. Avoid this if possible, as it leaves any incomplete subscribers in an indeterminate state. Click Retry to retry the commit, and Reset to move into PostPublishPhase cleanup.
CommitPhaseFailed	gsDataCommitFailed	Error message indicates the problem, usually communication failure. Click Retry to retry the commit. Click Cancel to cancel publishing, and attempt to go back to previous content. Click Continue to complete publishing and leave completely successful subscribers with the newly published content. (If activated independently, all successful paths retain newly-published material.)

Table 3.1 The GLOBAL-SITE Controller Phases of Publication

Overall Publication State	Section or Path State	Description
PostPublishPhase	various phase examples, all with options [...] of Done / Failed gsDelCDIrs , .. gsDataCopy , .. gsCopyCleanPath , .. gsDataCommit , .. gsCommitCleanPath , .. gsRollbackDone/Failed gsRollbackDoneClean.. gsRollbackFailedClean..	This phase occurs after publishing is completed, or after any failure. It cleans up any distributor or subscribers. When a phase fails, this attempts to roll back the subscriber to where it was before the publishing started.
PostPublishPhaseCanceled	as in PostPublishPhase	The Cancel button was clicked while the PostPublishPhase was in progress. Click Reset to quit the PostPublishPhase or Retry to try it again.
PostPublishPhaseDone	as in PostPublishPhase	The PostPublish phase has completed.
PostPublishPhaseFailed	as in PostPublishPhase	The process has failed during the PostPublish phase. Click Reset to quit the PostPublishPhase or Retry to try again.
PostPublishPhaseReady	as in PostPublishPhase	The controller is ready for the PostPublish phase, and will shortly begin that phase.
DeleteVersions	N/A	The section is deleting versions.
TestConn	N/A	The system is currently testing the connection.
RemovingFromDistributors	N/A	A section is being deleted. First the distributor is checked, and information is removed.
Unavailable	N/A	The publication is not available for publishing, and may be in the process of deletion. Can also be seen when a section is not available for updating. Generally indicates a deletion has been attempted; the deletion may have failed, and should be tried again.

Table 3.1 The GLOBAL-SITE Controller Phases of Publication

Understanding post publish phases

All publishing goes through a post-publish phase. In this phase, the GLOBAL-SITE Controller attempts to clear up any files or directories created to support publishing and, if necessary, to rollback the subscriber(s) to the previous state. The GLOBAL-SITE Controller attempts rollback only if there was a failure, or if publishing was cancelled.

Although the cleanup process may not be done for subscriber paths if the publishing failed, post publishing is always done for sections, regardless of how far the publishing progressed. At the end of publishing, section states move into the post publishing phase, and progress until the post publishing phase is done. If the section's final state is **Failed**, the next publishing process attempts to clean up before continuing.

The post publishing phase moves through many stages, and these can be reflected on the detail screen. It will help you to know what these states are, and what you can expect if you see them as a publication status. Some of these are included in Table 3.1, *The GLOBAL-SITE Controller Phases of Publication*. Details on the rollback state are provided in Table 3.2.

Rollback state	Explanation
gsRollbackDone	Rollback completed successfully, but cleaning the path was not attempted, probably because the post publish phase was cancelled.
gsRollbackFailed	Rollback failed and path cleanup was not attempted, probably because post publish phase was canceled.
gsRollbackDoneCleanDone	Both rollback and path cleanup were successful.

Table 3.2 Description of the rollback state

Rollback state	Explanation
gsRollbackFailedCleanDone	Rollback failed but path cleanup was successful.
gsRollbackDoneCleanFailed	Rollback was successful, but path cleanup failed.
gsRollbackFailedCleanFailed	Both rollback and path cleanup failed.

Table 3.2 Description of the rollback state

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 restoring the GLOBAL-SITE Controller in 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.
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.
2. Update any changes made since the last backup of the database that are pertinent or needed immediately.
3. 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 preceding scenarios are complex and challenging. We recommend that you request assistance from a field engineer.

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.

polled

Refers to one of the scheduled publishing options, used but not named, in the browser interface. ***Manual*** publishing is user-initiated, ***daily*** publishing happens at specified times during specified days of the week, ***polled*** publishing is continuous publishing, where the publishing process occurs at timed intervals, for instance, every 20 minutes, throughout the specified days of the week.

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

The server where the GLOBAL-SITE Controller delivers content to be viewed by 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 **gsInactivatePubGroupRepos**). 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.

gsCancelDeleteRepos

gsCancelDeleteRepos [-v] <section name>

Description

This command is used to cancel deletion of sections. It is normally used only when the deletion is taking an abnormal length of time. It will assume, in most cases, that some success has already been attained and mark the section as UNUSABLE. Deletion can then proceed again, at a later time.

The optional **-v** switch generates verbose output.

gsCancelDeleteVersions

gsCancelDeleteVersions <section name>

Description

This command cancels version deletion. It is to be used sparingly and only when version deletion appears to be taking too much time and/or there is no discernible progress.

gsCancelEdition <pub group name>

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

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

gsCancelPostPubPhase [-v] <publication name>

The optional **-v** switch generates verbose output.

```
gsCancelTestConnection <node name> <user name>
                        <destination name>
```

This command is used to cancel a test connection that appears to be making no progress.

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**.

gsChangePubGroupSchedule

Changing an existing Manual schedule

The following command syntax should be used when changing an existing **Manual** schedule:

```
gsChangePubGroupSchedule name=<publication name>  
[schedule=Manual] [startpub=<hh:mm>|-1]  
[activate=<hh:mm>|-1]
```

Changing to a Manual schedule

The following command syntax should be used when changing the schedule from **Daily** or **Polled** to **Manual**:

```
gsChangePubGroupSchedule name=<publication name>  
schedule=Manual [startpub=<hh:mm>|-1]  
[activate=<hh:mm>|-1]
```

Changing an existing Daily schedule

The following command syntax should be used when changing an existing **Daily** schedule:

```
gsChangePubGroupSchedule name=<publication name>  
[schedule=Daily [startpub==<hh:mm>...]  
[daysofweek=<day1, day2, ...>]  
[activate==<hh:mm>|-1]
```

Changing to a Daily schedule

The following command syntax should be used when changing the schedule from **Manual** or **Polled** to **Daily**:

```
gsChangePubGroupSchedule name=<publication name>  
schedule=Daily startpub==<hh:mm>  
[startpub=<hh:mm>...]  
[daysofweek=<day1, day2, ...>]  
[activate=<hh:mm>|-1]
```


Changing an existing Polled schedule

The following command syntax should be used when changing an existing **Polled** schedule:

```
gsChangePubGroupSchedule name=<publication name>  
    [schedule=Polled] [pollinterval=<hh:mm>]  
    [daysofweek=<day1, day2, ...>]  
    [activate==<hh:mm> | -1]
```

Changing to a Polled schedule

The following command syntax should be used when changing the schedule from **Manual** or **Daily** to **Polled**:

```
gsChangePubGroupSchedule name=<publication name>  
    schedule=Polled pollinterval=<hh:mm>  
    [daysofweek=<day1, day2, ...>]  
    [activate==<hh:mm> | -1]
```

Using the interactive form

The following command syntax is used to run the interactive form of the command:

```
gsChangePubGroupSchedule [-i]
```

Command usage

This command allows the user to change the scheduling parameters of a publication. The command can change parameters within the existing schedule type or change the publication to a completely different schedule type.

Manual type publications

A publication of schedule type **Manual** has only two parameters associated with it:

- ❖ A time of day to actually start publishing
- ❖ A time of day to activate (commit) the published content

The user manually starts the process by instructing the GLOBAL-SITE Controller primary (publisher) to deliver an edition. If properly instructed to do so, the appropriate sections will be updated and a new edition created (an already-existing edition can be specified, thereby bypassing these two steps).

Once the edition is determined (and created if necessary), the manually-scheduled publication then checks to see if a particular time of day has been specified for the actual process of publishing to begin. If so, the publication goes into Scheduled state until the next occurrence of that time of day (if that time of day occurred after the user manually initiated the process, but before the edition was created, publishing will begin immediately).

Once the content has been delivered to all target subscribers, the manually-scheduled publication then determines whether or not an activation time of day has been specified. If not, activation begins immediately. If so, activation waits until the next occurrence of that time of day (if that time of day passed after the user-initiated process began, activation occurs immediately).

Both of these times can be temporarily overridden when the user initiates the process: Each can be set to a different time or can be effectively removed for this one publish operation.

Daily type publications

A publication of schedule type **Daily** can have as many as 24 times of day specified, all of them effective on specified days of the week. On each day of the week that has been selected, the publication will automatically, at each selected time of day, update its active sections, create an edition based on the latest versions from each of its active sections, and begin publishing. Any activation time of day will be ignored as the publication will attempt to complete the publishing process as quickly as possible.

Polled type publications

A publication of schedule type **Polled** has a specified interval that is effective on specified days of the week. On those days, the publication will attempt to update and publish according to the frequency specified by the polling interval. For example, the

interval may be every ten minutes. Again, any activation time is ignored as the publication will attempt to publish as quickly as possible.

The **gsChangePubGroupSchedule** command allows an activation time to be specified for **Daily** and **Polled** publications. However, this time is not used as long as the publication remains on its **Daily** or **Polled** schedule. However, if the publication's scheduling is **paused**, it is temporarily on a manually-activated schedule and uses this specified activation time unless it is overridden when the publish is initiated. For more information, see the **gsPausePubGroupSchedule** and **gsResumePubGroupSchedule** commands.

◆ **Note**

*A publication's scheduling **must** be paused before using this command, unless it is a manually-scheduled publication.*

Each change must include the publication name (**name=<publication name>**) and, optionally, the desired schedule type (**schedule=<Manual|Daily|Polled>**). If the schedule is not entered, it is assumed that the schedule type is not changing. Additional parameters can be necessary depending upon whether the schedule type remains the same or there is a change to a completely new schedule type.

If a manual schedule's parameters are being changed (the publication remains **Manual**), there are two optional parameters:

- ❖ **startpub=** followed by a two-digit hour, a colon and a two-digit minute. This specifies a new time of day for manually-activated publishing to begin. This parameter is sometimes referred to as PubPrep time. For example, the parameter value **startpub=16:30** causes manually-activated publishing to occur at 4:30 P.M. (any update and edition creation is done immediately). Alternately, a **-1** can be entered, which removes any scheduled time and publishing starts immediately after manual activation (and any updating of sections and edition creation). An example is the parameter value **startpub=-1**

❖ **activate=**, followed by either a valid two-digit hour, colon, and two-digit minute or a **-1**, with the same results as the first parameter, except that this affects activation time.

If either of these parameters is not specified, that part of the schedule is unchanged.

If changing a **Polled** or **Daily** schedule to a **Manual** schedule, the two parameters alone are again the only additional parameters (other than **name=** and **schedule=**), again optional. Of course, **schedule=Manual** will be mandatory in order to know to change to **Manual**.

If you are making changes to an already **Daily** schedule, the additional parameters are 0-24 of the **startpub=** parameter values (**-1** is not valid here - only hour and minute), an optional **daysofweek=** parameter, and an optional **activate=** parameter. If no **startpub=** parameters are present, no change is made to the times of day for **Daily** scheduling; however, if one time is entered, all desired times must be entered. For example:

```
gsChangePubGroupSchedule name=DailyPub
    schedule=Daily startpub=08:00 startpub=11:30
    startpub=16:30
```

The publication DailyPub continues its **Daily** schedule and begins updating sections and publishing at 8:00A.M., 11:30 A.M., and 4:30 P.M..

The **daysofweek=** parameter specifies the days of the week for which the "startpub=" times are valid. For example, the above command might also have the following parameter entry:

```
daysofweek=Mon,Wed,Fri
```

This parameter value specifies that publishing is to occur only on Monday, Wednesday, and Friday of each week. The list of days is always a three-character day (Sun, Mon, Tue, Wed, Thu, Fri, Sat) and the days are always separated by a comma, with no spaces.

If this parameter is not entered, the currently-specified days of the week are not changed.

The **activate=** parameter is as described earlier and, again, has no effect unless scheduling is paused.

Changing from **Manual** or **Polled** to **Daily** is the same with the following exceptions:

- ❖ **schedule=Daily**. Mandatory, in order to know we are changing to Daily.
- ❖ **startpub=<hh:mm>**. One instance is mandatory (and up to 23 can optionally be entered) as Daily will require at least one start time.
- ❖ **daysofweek=** Optional, If not present and changing from Polled, the same days that the old schedule had will be used. If changing from Manual and there is no "daysofweek", all seven days are used.

If an already Polled schedule is being changed, the **daysofweek=** and **activate=** parameters are optional and **pollinterval=** is optional. The first two are as described above for their optional entry. If not specified, there is no change and if **activate=** is specified, it is potentially used only when the normal schedule is paused.

The **pollinterval=** parameter is much like a **startpub=** parameter except that a number of minutes is being entered rather than a specific time of day. The number of minutes is used as the interval between the start of each update/create edition/publish sequence, and is expressed as hh:mm, where **hh** is the number of hours and **mm** is the number of minutes. For example, **pollinterval=01:30** sets a polling interval of an hour and a half or every 90 minutes. If not specified, the poll interval is unchanged.

If changing a **Manual** or **Daily** schedule to a **Polled** schedule, the only differences from above are:

- ❖ **schedule=Polled**. Mandatory, in order that it be known to change to Polled
- ❖ **pollinterval=**. Mandatory, as there is a change to a completely new type of schedule

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.

◆ Note

*An alternative is to use the **gsPublish** command.*

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.

gsCreateBigip

```
gsCreateBigip name=<bigipname> [description=<string>] user=<string>  
password=<string>
```

```
gsCreateBigip [-i]
```

Description

This command creates an object representing a BIG-IP Controller. The **<name>** variable is a string that is usually the distributor name followed by a colon, followed by the address of the BIG-IP Controller. The **<description>** variable is any description string you want. The **<user>** and **<password>** variables are still required by the command, but not used at this time.

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>] [bigip=<string> virtserv=<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 <bigip> variable indicates the name of the BIG-IP Controller that is servicing this node (For more information, see the **gsCreateBigip** command). When using this variable, the <virtserv> variable is also required. The name is usually the name

of the distributor that publishes to this BIG-IP Controller's virtual server, followed by a colon and the IP address of the BIG-IP Controller.

The **<virtserv>** variable is a string that indicates the name of the virtual server being used. This name is normally an IP address with a colon followed by a port number. When using this variable, the **<bigip>** variable is also required.

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

gsCreatePubGroup

Creating a Manual scheduled publication

The following command syntax should be used when creating a **Manual** scheduled publication:

```
gsCreatePubGroup name=<publication name> [pubstop=<pubstop  
type>...] [puberror=<puberror option>] [description=<string>]  
[pubtype=<Archived or -Archived>] [mail=<string>]  
[independentactivation=y | [independentactivation=n]  
[singledistactivation=<y/n>][subscriberactivation=<1-100>]  
[activate=<hh:mm>]] ] [[disablevirtservers=y  
persisttimeout=<1-120>] | [disablevirtservers=n]]  
[[schedule=Manual] [startpub=<hh:mm>]]
```

Creating a Daily scheduled publication

The following command syntax should be used when creating a **Daily** scheduled publication:

```
gsCreatePubGroup name=<publication name> [pubstop=<pubstop  
type>...] [puberror=<puberror option>] [description=<string>]  
[pubtype=<Archived or -Archived>] [mail=<string>]  
[independentactivation=y | [independentactivation=n]  
[singledistactivation=<y/n>][subscriberactivation=<1-100>]  
[activate=<hh:mm>]] ] [[disablevirtservers=y  
persisttimeout=<1-120>] | [disablevirtservers=n]]  
schedule=Daily startpub=<hh:mm> [startpub=<hh:mm> ...]  
[daysofweek=<day1,day2,...>]
```

Creating a Polled scheduled publication

The following command syntax should be used when creating a **Polled** scheduled publication:

```
gsCreatePubGroup name=<publication name> [pubstop=<pubstop  
type>...] [puberror=<puberror option>] [description=<string>]  
[pubtype=<Archived or -Archived>] [mail=<string>]  
[independentactivation=y | [[independentactivation=n]  
[singledistactivation=<y/n>][subscriberactivation=<1-100>]
```



```
[activate=<hh:mm>]] ] [[disablevirtservers=y  
persisttimeout=<1-120>] | [disablevirtservers=n]]  
schedule=Polled pollinterval=<hh:mm>  
[daysofweek=<day1,day2,...>]
```

Using interactive mode

The following command syntax should be used when using interactive mode:

```
gsCreatePubGroup [-i]
```

Description

This command creates a new publication. The newly-created publication starts out as paused from scheduling.

The following table lists and describes the arguments available with the **gsCreatePubGroup** command.

name=	Used to specify the name of the pub group.
pubstop=	Can be used as many times as necessary to specify the desired pub stops (None , AfterEditionCreation , AfterPubPrep , AfterWANSync , AfterChangeTrees , AfterCopy , AfterCommit). These may not be considered useful with publications that are scheduled and the AfterCopy pub stop will be ignored if an activation time has been scheduled. However, if that time is in the future, the publication will effectively stop until the scheduled time arrives. Note that only AfterCommit is supported via the browser-based interface.
puberror=	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.
pubtype=	Should be either Archived or -Archived . The default, if this argument is not entered, is Archived .

Table B.1 gsCreatePubGroup arguments

mail=	Should be a string of e-mail addresses to which electronic mail will be sent when certain problems occur with the publication. For example: mail="support@yourcompany.com supervisor@yourcompany.com"
indepentacti vation=	Requires a y for Yes or an n for No. If selected (with a y), the following parameters either cannot be used or will have no effect: singledistactivation= , subscriberactivation= , or activate= . Selecting independent activation causes each path that is published to be activated as quickly as possible, regardless of whether or not other paths on the same subscriber are ready for activation. The result is that each path for a subscriber is activated with maximum speed, but consistency with other paths being published on that subscriber may be lost until the other paths have been successfully activated, too. (Activation simply means that new material copied to the subscriber path is made visible to those accessing the data). If not specified, the default value for this argument is "n", not selected.
singledistac tivation=	Requires a y for Yes or n for No. As already mentioned, this argument and independentactivation= argument are mutually exclusive, that is, both cannot be y . If y , activation is performed on only one distributor at a time. When a distributor completes activation, activation is started on the next applicable distributor until all applicable distributors have completed activation of all paths on all subscribers that they distribute to. If not specified, the default value for this argument is "n", not selected.
subscriberac tivation=	Requires a number in the range of 1-100, which represents the percentage of subscribers for a particular distributor to be activated simultaneously. If independent activation has been selected, this argument is not considered valid. The default, if not entered, is 100. This argument is used when activating subscribers behind a BIG-IP Controller.
activate=	Requires a time of day, such as 16:30 , at which time <i>manually-activated</i> publications are to be activated. This argument is not useful if independent activation was selected.
disablevirt servers=	Requires a y for Yes or an n for No. Selecting this feature specifies that virtual server nodes (behind a BIG-IP Controller) be disabled while their content is being activated. The default if there is no entry is n
persisttimeo ut=	Used only if disablevirtservers= is y . It is a number in the range of 1-120 that specifies the number of seconds after which to timeout persistent connections (again, used with a BIG-IP Controller). The default value is 120.

Table B.1 gsCreatePubGroup arguments

schedule=	Specifies which type of scheduling is normal for this publication. If not entered, the default is Manual. The values to be used are either Manual , Daily , or Polled . For more information, see the gsChangePubGroupSchedule command for an explanation of these types of schedules.
startpub=	Used only with Manual or Daily schedules. It is a time of day in the form HH:MM, for example, 15:30. For a Manual publication, this represents the time of day at which a manually-initiated publish is to actually begin the PubPrep phase (actual publishing beginning after any updating of sections and edition creation). If not entered for a Manual publication, PubPrep begins as soon as possible. For Daily schedules, this argument must be entered once and can be entered up to 23 more times. Each entry represents a different time of day at which to begin updating sections, create an edition, and then publish it.
pollinterval=	Used only by Polled schedules and is mandatory for them. It is the interval in minutes to elapse between each initiation of the section updating, edition creation, publishing cycle. For example, an entry of 01:00 results in polling every 60 minutes.
daysofweek=	Valid only with Daily and Polled schedules. It specifies the days of the week on which the schedules are to be active, for example, "daysofweek=Mon,Wed,Fri" specifies that scheduling be done only on Monday, Wednesday, and Friday. If this entry is not specified, scheduling takes place on all seven days of the week. The current values are: Sun, Mon, Tue, Wed, Thu, Fri, and Sat. The values must be separated by a comma with no intervening spaces.

Table B.1 *gsCreatePubGroup arguments*

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.

Table B.2 gsCreateRepeater results

gsCreateRepos

```
gsCreateRepos name=<reposname> [type=FTPStage|FTPna]
[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. FTPStage is available, as well as FTPna for non-archived source.

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.

gsDeleteBigip

gsDeleteBigip <name>

Description

This command deletes any BIG-IP Controller that is not currently referred to by a node.

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. When a node is deleted, any connection to a BIG-IP Controller is automatically deleted. However, if a BIG-IP Controller is removed, all nodes must have been previously disconnected from that controller.

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. An attempt is made to clean up all information about this publication on all distributors known to the primary (publisher).

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.

This command also removes information from the distributors.

This command attempts to delete information for this section on all distributors through which this repository is currently published. The repository itself will not be removed from the primary until this happens. If for some reason there is a problem with a remote distributor such that it is not certain that the information has been removed, the section is marked as UNUSABLE and deletion can be tried again at the user's convenience.

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 unserved 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=<repos.name>

Description

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

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.

gsPausePubGroupSchedule

gsPausePubGroupSchedule <pub name>

Description

This command suspends the normal scheduling.

Pausing a Daily or Polled schedule puts the schedule into Manual mode and an Activation Time previously specified is used unless overridden. For more information, see `gsCreatePubGroupSchedule` and `gsChangePubGroupSchedule`.

Pausing a Manual schedule cancels any PubPrep time or Activation time previously scheduled

If a PubPrep time was scheduled, the publication's state is set to `EditionReady`, so that when publishing is initiated, it will skip updating and edition creation and immediately begin publishing. It also cancels activation time, if that was scheduled. If only activation was scheduled, it sets state to `CopyDone` so that publishing will resume with the activation or commit phase of publishing. Subsequent publishes will ignore the normal PubPrep time and Activation time, unless an overriding value is presented. For more information, see the `gsResumePubGroupSchedule` command.

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] [-p -1|-2|hh:mm] [-a -1|-2|hh:mm] <pub group name> <edition number>`

`gsPublish -r [-v] [-d] <publication name>`

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.

The optional **-p** switch is used to override manual PubPrep time, while the optional **-a** switch is used to override manual Activation time. For both switches, the **-1** argument causes the already-specified time to be ignored. The **-2** argument, which is the default if either of these switches is not specified, causes the already-specified time to be used. The **hh:mm** argument causes this time of day to be used instead of the already-specified time. Note that these changes to the already-specified time are one-time only changes and revert back to the normal time after this operation is complete.

The optional **-r** switch with the **<publication name>** argument retries a failed/cancelled publication phase.

gsResetNodeState

gsResetNodeState <node name> <state>

Description

This command is used to set a node state to **Idle** if a problem occurred during a Test Connection operation and has left the node in the **TestConn** state.

This command can also be used to set the node to any supported state.

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.

gsResetPublish

gsResetPublish [-v] [-d] <publication name>

Description

This command ensures that post publication processing is done after any cancellations/failures and then resets to IDLE.

The optional **-v** switch generates verbose output.

The optional **-d** switch forks the process. This switch is used for debugging only.

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**.

gsResumePubGroupSchedule

gsResumePubGroupSchedule <pub name>

Description

This command resumes the normal scheduling. The normal scheduling parameters are reactivated to their previous values, unless they have been changed in the interim. For more information, see the **gsPausePubGroupSchedule** command.

gsSetBigipInfo

```
gsSetBigipInfo name=<bigipname> [newname=<new name>]  
    [description<description>] [user=<string>] [password=<string>]  
gsSetBigipInfo [-I]
```

Description

The primary use of this command is to change the name (<**new name**=>) or description (<**description**=>) of a BIG-IP Controller.

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

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>]  
[virtserv=<string, .=<as is>, space=<no virtual server or bigip>]  
bigip=<string, .=<as is>]
```

gsSetNodeInfo [-i]

Description

This command allows the previously-entered information to be changed. Also, because information on a node is eliminated when the node is first updated, this command can be used to set new information on that updated node.

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.

The **<virtserv>** variable is a string that usually indicates the IP address of the virtual server being used. If this argument is not entered, then the **<bigip>** variable should not be used.

The **<bigip>** variable indicates the name of the BIG-IP Controller that is servicing this node.

An example:

```
$ gsSetNodeInfo
Enter Current Node Name: mynode.dev.net

Enter New Node Name [.=mynode.dev.net]: .
Enter CPU description [.=Pentium III]:.
Enter Operating System description [.=Windows NT]: WNT
Enter additional description [.=Personal Workstation]: Group Workstation
Enter new Virtual Server [.=100.100.100.1:80,space=<no virtual server>]:.
Enter new bigip [.=globalsite2.dev.net:30.30.30.1]:.
set info completed.
```

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>] [pubstop=<pubstop type>...]
puberror=<puberror option>]
[independentactivation=<y/n>][singledistactivation=<y/n>]
[disablevirtservers=<y/n>] [subscriberactivation=<1-100>]
[persisttimeout=<no. of seconds>] [mail=<string/space=none>]
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 **<independentactivation>** variable is used to turn independent activation on or off.

The **<singledistactivation>** variable is used to turn the distributor activation method on or off.

The **<disablevirtservers>** variable is used to turn on or off the disable virtual servers feature.

The **<subscriberactivation>** variable is used to change the percentage of subscribers behind a BIG-IP Controller that are turned off at one time while being activated with new content.

The **<persisttimeout>** variable is used to change the number of seconds for persistence timeouts.

The **<mail>** variable is used to change the recipients of mail.

For a more detailed description of some of these variables, see the **gsCreatePubGroup** command.

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

gsSetPublished

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

usage: gsSetPublished [-i]

Description

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

The **pubgroup** argument is the name of the publication.

The **edition** argument is the edition number containing the version or versions to be set as having been published.

The **[node]** argument is the node name of the subscriber.

The **[user]** argument is the user name used to access the subscriber's paths.

The **[path]** argument is the directory path for the subscriber.

The **[setedition]** argument is set to either <yes> or <no>. This argument is used to set or not set the edition as the published edition for this publication. If <yes> is specified, then **edition** is set as the published edition for this publication. If <no> is specified, then the published edition that was already recorded remains as the published edition.

If **[node]**, **[user]**, and **[path]** are specified, then only that path is set to say that **edition** (and the appropriate version/section) is published on it. Any other paths for the same subscriber are not affected.

If only **[node]** and **[user]** are specified, then all paths for that node/user pair (subscriber) are set for this publication. This option is useful when a new subscriber has been added that the user is certain contains the section versions contained by **edition**. This option is also useful when an existing subscriber is switched to a new GLOBAL-SITE distributor.

If **[node]** and **[user]** are not specified, then all paths for this publication are set. This option is useful when bringing up a new publication where it is certain that all subscribers already contain the material specified by edition. In this case, setting the **[setedition]** option to **<yes>** is appropriate.

 WARNING

Use this command with CAUTION. This is because no check is made to verify the actual content.

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.

gsShowAllBigips

gsShowAllbigips [-v | -h | -d]

Description

This command prints, on screen, a list of all BIG-ips that have been created (see the **gsCreateBigip** command).

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

The **-h** switch displays the header only (name, local object ID, object type:version, description, and state). This switch is useful when a lengthy printout is not needed.

The **-d** switch formats the instantiation as if it were to be written to disk but instead dumps the instantiation to the screen. This switch is useful for debugging.

gsShowAllFtpAccess

gsShowAllFtpAccess [-v | -h | -d]

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.

The **-h** switch displays the header only (name, local object ID, object type:version, description, and state). This switch is useful when a lengthy printout is not needed.

The **-d** switch formats the instantiation as if it were to be written to disk but instead dumps the instantiation to the screen. This switch is useful for debugging.

gsShowAllNodes

gsShowAllNodes [-v | -h | -d]

Description

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

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

The **-h** switch displays the header only (name, local object ID, object type:version, description, and state). This switch is useful when a lengthy printout is not needed.

The **-d** switch formats the instantiation as if it were to be written to disk but instead dumps the instantiation to the screen. This switch is useful for debugging.

gsShowAllPubGroups

gsShowAllPubGroups [-v | -h | -d]

Description

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

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

The **-h** switch displays the header only (name, local object ID, object type:version, description, and state). This switch is useful when a lengthy printout is not needed.

The **-d** switch formats the instantiation as if it were to be written to disk but instead dumps the instantiation to the screen. This switch is useful for debugging.

gsShowAllRepeaters

gsShowAllPubGroups [-v | -h | -d]

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

The **-h** switch displays the header only (name, local object ID, object type:version, description, and state). This switch is useful when a lengthy printout is not needed.

The **-d** switch formats the instantiation as if it were to be written to disk but instead dumps the instantiation to the screen. This switch is useful for debugging.

gsShowAllRepos

gsShowAllRepos [-v | -h | -d]

Description

This command prints on screen a list of all sections.

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

The **-h** switch displays the header only (name, local object ID, object type:version, description, and state). This switch is useful when a lengthy printout is not needed.

The **-d** switch formats the instantiation as if it were to be written to disk but instead dumps the instantiation to the screen. This switch is useful for debugging.

gsShowBigip

gsShowBigip [-h | -d]

Description

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

Using no switch prints out detailed information for each publication.

The **-h** switch displays the header only (name, local object ID, object type:version, description, and state). This switch is useful when a lengthy printout is not needed.

The **-d** switch formats the instantiation as if it were to be written to disk but instead dumps the instantiation to the screen. This switch is useful for debugging.

gsShowFtpAccess

gsShowFtpAccess [-h | -d] <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.

The **-h** switch displays the header only (name, local object ID, object type:version, description, and state). This switch is useful when a lengthy printout is not needed.

The **-d** switch formats the instantiation as if it were to be written to disk but instead dumps the instantiation to the screen. This switch is useful for debugging.

gsShowNode

gsShowNode [-h | -d] <nodename>

Description

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

The **-h** switch displays the header only (name, local object ID, object type:version, description, and state). This switch is useful when a lengthy printout is not needed.

The **-d** switch formats the instantiation as if it were to be written to disk but instead dumps the instantiation to the screen. This switch is useful for debugging.

gsShowPubGroup

gsShowPubGroup [-h | -d] <pubgroup name>

Description

This command displays information pertaining to a specific publication.

The **-h** switch displays the header only (name, local object ID, object type:version, description, and state). This switch is useful when a lengthy printout is not needed.

The **-d** switch formats the instantiation as if it were to be written to disk but instead dumps the instantiation to the screen. This switch is useful for debugging.

gsShowPubStatus

gsShowPubStatus [-h | -d] <pubgroup name>

Description

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

The **-h** switch displays the header only (name, local object ID, object type:version, description, and state). This switch is useful when a lengthy printout is not needed.

The **-d** switch formats the instantiation as if it were to be written to disk but instead dumps the instantiation to the screen. This switch is useful for debugging.

gsShowRepeater

gsShowRepeater [-h | -d] <node name>

Description

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

The **-h** switch displays the header only (name, local object ID, object type:version, description, and state). This switch is useful when a lengthy printout is not needed.

The **-d** switch formats the instantiation as if it were to be written to disk but instead dumps the instantiation to the screen. This switch is useful for debugging.

gsShowRepos

gsShowRepos [-h | -d] <repository name>

Description

This command displays repository information.

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

The **-h** switch displays the header only (name, local object ID, object type:version, description, and state). This switch is useful when a lengthy printout is not needed.

The **-d** switch formats the instantiation as if it were to be written to disk but instead dumps the instantiation to the screen. This switch is useful for debugging.

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.

C

Reconfiguring the Controller

Changing the host name and/or IP address of a GLOBAL-SITE Controller

When changing the IP address or the host name of a GLOBAL-SITE Controller, you must leave the GLOBAL-SITE identifier the same. The GLOBAL-SITE identifier is independent of the host name of the controller.

WARNING

If you change the GLOBAL-SITE identifier, the controller cannot communicate with other GLOBAL-SITE Controllers in a publisher/distributor configuration.

In order to change the host name or IP address on a GLOBAL-SITE Controller, you need to have the following information at hand:

- ❖ The new IP address and host name
- ❖ The current GLOBAL-SITE identifier assigned to the box. You can find this in either the List Distributors screen or the System Settings screen.
- ❖ A list of all other GLOBAL-SITE Controllers that use this controller as a distributor.

The process of changing the IP address or host name of a controller includes three separate tasks. You must:

- ❖ Edit both the network and the hosts files, replacing the old host name and domain name with the new names, and then verify that the GLOBAL-SITE Controller is using the new names.
- ❖ Reset the GLOBAL-SITE identifier on the controller.

- ❖ Update the distributor lists on the other GLOBAL-SITE Controllers that reference the original host name. For each of these controllers, delete the old distributor, and create a new distributor. If you have no additional GLOBAL-SITE Controllers, then you do not need this step.

◆ Note

*If you change only the IP address, you have fewer tasks to perform. In the edit task, you need edit only the **/etc/hosts file**, and then in the update task you update only the GLOBAL-SITE Controllers that recognize the changed controller by IP address.*

Editing the network and hosts files

The first task in this process is to change the name of the GLOBAL-SITE Controller. It is only possible to do this using the command line utility on the controller itself.

To change the host and domain names

1. To start, log in to the GLOBAL-SITE Controller at the console as **root**.
2. Edit the **/etc/sysconfig/network** file, replacing the fields labeled **HOSTNAME** and **DOMAINNAME** with the correct values.

A sample **/etc/sysconfig/network** file is shown:

```
NETWORKING=yes
FORWARD_IPV4=false
HOSTNAME=gshost1
DOMAINNAME=f5.com
GATEWAY=192.168.101.254
GATEWAYDEV=eth0
```

Perhaps you want to change **gshost1** to **myhost1**, and change **f5.com** to **myshop.com**.

3. Edit the **/etc/hosts** file, replacing the old host name with the new host name.

A sample **/etc/hosts** file is shown:

```
127.0.0.1 localhost localhost.localdomain
192.168.101.30 gshost1.f5.com gshost1
```

In the above sample, you would change:

gshost1.f5.com gshost1

to read:

myhost1.myshop.com myhost1

4. After changing the host name, reboot the GLOBAL-SITE Controller by typing **/sbin/init 6** at the command line prompt.
5. After the controller reboots, it should be using the correct host name and IP address.
Verify this by logging in as **root**, and running the following commands:

hostname (shows current host name, and should match the new name you entered in step 2)

dnsdomainname (shows current DNS domain)

If everything is correct, you can move on to the next task where you reset the GLOBAL-SITE identifier.

If there is anything not correct, start again at step 1.

Resetting the GLOBAL-SITE identifier

The second task in this process is to reset this controller so that it recognizes its new name. Again, do the task on the controller itself, using the command line utility.

Remember that in the command line utility, the distributor is called a repeater.

To reset the GLOBAL-SITE identifier for this controller

1. Log on as the user **gsite**:
su - gsite

2. Then set the distributor (repeater) information for this GLOBAL-SITE Controller, and delete the old distributor (repeater) information from the database. For both the **<newhostname>** and the **<oldhostname>**, use the fully qualified domain name.

```
/usr/local/bin/gsSetRepeaterInfo name=<newhostname>  
    identifier=<identifier>
```

```
/usr/local/bin/gsDeleteRepeater <oldhostname>
```

This GLOBAL-SITE Controller now has the correct host name and IP address, and its GLOBAL-SITE identifier is correct.

3. To verify the new name and IP address, use the **gsShowRepeater <newhostname>** command. The GLOBAL-SITE Controller converts the data you enter in decimal to a hexadecimal number. You should see a line that shows the installation ID in hexadecimal, like this:

```
Installation ID: 0X12345
```

This number should be something other than a zero; this is the GLOBAL-SITE identifier.

Updating the distributor lists on other GLOBAL-SITE Controllers

The next task in this process is to add the new host name to the distributor lists on the other GLOBAL-SITE Controllers that referenced the old name. They must be updated with the new server name on their distributor lists. If you have no additional GLOBAL-SITE Controllers, you do not need to do this.

You can use the GLOBAL-SITE Controller browser interface to delete the old distributor name and create a new one. For each GLOBAL-SITE Controller that needs to reference the new name, you need to log on to that GLOBAL-SITE Controller, and update the information.

To update the other controllers, you log on to each individual controller and complete the following three tasks:

- ❖ Create the new distributor.
- ❖ Reassign to the new distributor all subscribers that reference the old distributor.
- ❖ Delete the old distributor.

When you reassign the subscribers, note that you reconnect to the new distributor any subscribers that had previously been set up with the old distributor. If there are no subscribers, you can skip this task.

If you have multiple GLOBAL-SITE Controllers, you then log off the first, and log on to the next GLOBAL-SITE Controller and repeat the process.

Creating a new distributor

You create a new distributor using the new host name, the original GLOBAL-SITE identifier, and the key phrase shared between the publisher and distributor pair.

To create the new distributor

1. Log on to the GLOBAL-SITE Controller that needs to recognize the changed name.

You may want to verify the GLOBAL-SITE identifier and key phrase for this controller before continuing. To do so, in the navigation pane click List Distributors.

The Distributor List screen displays.

Click the name of the old GLOBAL-SITE Controller in order to view the Distributor Detail screen and verify the GLOBAL-SITE identifier and key phrase. Then continue to add the new distributor.

2. In the navigation pane, click Add a Distributor.

The Add a Distributor screen displays.

3. In the **Host Name** box, type the new unique identifier host name or IP address for the original GLOBAL-SITE Controller that you changed. (You have this information from the first task where you changed the names by editing files.)
4. In the **Description** box, type in an optional description. This makes it easy for you to identify the distributor. You can use the same description that you had before the name change.
5. In the **GLOBAL-SITE Identifier** box, type the unique, pre-assigned identifier for the GLOBAL-SITE Controller that you are re-adding as a distributor. This identifier has not changed. You can find the GLOBAL-SITE identifier by looking at the host's List Distributor screen if you need a reminder.
6. In the **Key Phrase** box, type in the key phrase that is shared by the publisher and distributor pair. (Both the publisher and the distributor must have the same Key Phrase.)

This should not have changed.
7. Click **Add** to complete the addition.

The GLOBAL-SITE Distribution List screen displays, with the re-added distributor listed at the bottom.

Once you re-add the distributor, you can move on to the next task of reassigning subscribers to the distributor. If the distributor does not use subscribers, you can skip this task and go to the final task of deleting the old distributor.

Reassigning subscribers to the new distributor name

Once you have created the new distributor, and before you delete the old distributor, you need to connect to the new distributors any subscribers that are set to the old distributor.

You can undertake this task at a later time, but doing it now may make it easier to track your changes, and you can do it all before logging out of this GLOBAL-SITE Controller and into another one, if you have multiple distributors.

To reassign a subscriber

When you reassign a subscriber, you redefine the path that points to the subscriber.

1. In the navigation pane, click List Distributors.
The Distributor List screen displays.
2. On the Distributor List screen, click the name of the distributor you have just changed.
The Distributor Detail screen displays.
3. On the Distributor Detail screen, review the subscribers.
For each subscriber on this distributor, you will change the path from the old distributor to the new distributor, using the next two steps).
4. Click the first subscriber name.
The Subscriber Detail screen displays.
5. Change the GLOBAL-SITE Distributor information from the old distributor name to the new name. (Refer to the online help for the Subscriber Detail screen if you need procedural details for changing subscriber settings.) When you have finished the changes, click the Save button.
6. Click the browser's **Back** button or back arrow until you return to the Distributor Detail screen, then select the next subscriber.
7. Repeat steps 4, 5, and 6 for each subscriber that references the changed distributor.

When you have completed the task of reassigning the subscribers to the new host name, delete the old distributor on this GLOBAL-SITE Controller. After that, you can move to the next GLOBAL-SITE Controller that needs to recognize the changed distributor.

Deleting the old distributor

Now that you have added the new distributor, and reassigned all subscribers referencing it, you are ready to delete the old distributor. Once you have finished this task, you can move on to another GLOBAL-SITE Controller if you have multiple controllers that reference the changed distributor.

To delete the old distributor

You should still be logged on to the GLOBAL-SITE Controller that needs to recognize the changed name.

1. In the navigation pane, click List Distributors.
The Distributor List screen displays.

You may want to click the name of this GLOBAL-SITE Controller, in order to view the Distributor Detail screen and verify that no subscribers are still attached to it.

2. On the Distributor List screen, locate the name of the old GLOBAL-SITE Controller for which you just changed the host name or IP address. Click the delete button  at the far right of that name.

This deletes the old distributor from this GLOBAL-SITE Controller.

3. Log off this GLOBAL-SITE Controller.

If there are additional distributors that recognize the changed GLOBAL-SITE Controller, you need to log onto the next GLOBAL-SITE Controller, and repeat the tasks of adding the distributor with the new name, reassigning the subscribers to it, and deleting the reference to the old name.

Glossary

activate new content

In the publishing process, activating new content means moving content just delivered to a subscriber from its temporary location to its permanent location on the subscriber, and making it available to the viewer.

See also, *controlled activation* and *independent activation*.

archived publications

In archived publications, the GLOBAL-SITE Controller stores a copy of the file contents of each edition of a publication using its RCE version control system. This archiving of the contents allows the user to revert to a former publication edition and/or section version.

commit

To move files from a temporary to a permanent directory on a *subscriber* (target server). See *activate new content*.

controlled activation

In controlled activation, the publishing process moves all content together through a series of phases for a collection of subscribers. For instance, before the content for any one section can be activated on one or more subscribers, all content for all sections is copied to all subscribers. Compare to *independent activation*.

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

Part of the final process in the GLOBAL-SITE Controller publishing process, where the content, a specific edition of a publication, is copied to the target servers, or *subscribers*.

distributor

A GLOBAL-SITE Controller that may be located in a remote data center which provides the publisher access to additional subscribers. Compare to *publisher*.

disk mirroring

When setting up the GLOBAL-SITE Controller, you can choose to mirror the two disks. This writes duplicate data to both disks, allowing you to use half of your disk space, but providing error recovery in case either of the two disks should fail.

edition

In archived publications, a complete collection of specific versions of selected sections. An edition is an instance of a publication that indicates two things: 1) which versions of selected sections are to be included, and 2) that the edition is ready to be delivered (*published*), or has already been delivered to subscribing servers.

An archived 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.

Non-archived publications do not have editions stored on the GLOBAL-SITE Controller.

exception

A directory, located within the 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.

First-Time Boot utility

A utility that walks you through the initial system configuration process. The First-Time Boot utility runs automatically when you turn on a controller for the first time.

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 provided by F5 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.

independent activation

In independent activation, the GLOBAL-SITE Controller moves content through the publishing phases section by section, independent of its association with a subscriber or subscribers. One section can be copied and then activated for one subscriber before another section is copied. Compare to *controlled activation*.

key phrase

A phrase that is shared by the pair of publisher and distributor 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. (Unlike the GLOBAL-SITE identifier, the key phrase is user-created, and not provided by your vendor.)

non-archived publications

Non-archived publication allows for basic, fundamental content replication. A non-archived publication does not keep copies of the content (versions) on the GLOBAL-SITE Controller as an archived publication does. Rather, it takes content directly from the user's source, and moves it directly to the location where it will be accessed by customer. However, non-archived publication does create and keep version lists of dated files so that it can detect changed files.

passphrase

A string of words and characters that you type in to authenticate yourself as a user. Passphrases are similar to passwords, but longer. Passphrases are considered to be more secure because of their greater length.

persistence

A series of related connections received from the same client. When persistence is turned on, the BIG-IP Controller does not load balance the connections; instead it returns the clients to the node that they previously connected to.

private key

One of two keys used in asymmetric cryptography. The private key can be used to encrypt data which can then be decrypted using the public key, which is shared with all involved parties.

public key

The second of two keys used in asymmetric cryptography. The public key, which is shared by all involved parties, is used to decrypt data which has been encrypted with the private key.

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 determine how the GLOBAL-SITE Controller handles: initiating the publishing process, scheduling the publishing, and activating the new content. It also provides settings for how the GLOBAL-SITE Controller deals with the BIG-IP Controller virtual servers, and how it handles errors.

publication targets

See *subscribers*.

publish

The process of identifying specific versions of sections to be included in an edition and then delivering and activating this edition to subscribers. 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 server that content is published *to*. In the GLOBAL-SITE Controller, the subscriber information includes a destination path for each section of a publication.

version

A version exists in archived publications only. A version is one particular instance of a section (in an archived publication) that differs from other earlier or later instances of that section due to changes or modifications. In archived publications, 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.

Non-archived publications do not have versions. See *non-archived publications*.

Index

A

- access methods
 - specifying 2-3
- access settings
 - providing 2-11
- activation
 - specifying subscribers 3-15
- activation defaults
 - changing 3-9
- activation settings 3-4
- activation. See also content activation
- Active Server Pages 1-7
- Add a BIG-IP screen 3-11
- Add a Distributor screen
 - viewing 2-19
- Add Subscriber screen
 - viewing 2-20
- Add subscriber screen
 - viewing 2-17
- admin interfaces 3-10
- administration
 - and command line utility 1-4
- administrative tasks
 - navigating to 1-3
- Adobe PDF 1-7
- application delivery
 - ensuring 1-6
- application management 1-6
- application rollback 1-6
- applications
 - protecting for transfer 1-6
 - supported 1-7
 - synchronizing 1-7
- archived publications
 - and archived sections 1-11
 - and editions 2-9
 - and project history 1-5
 - and versions 3-29

- as attribute 2-5
 - changing to non-archived 3-17
 - comparing editions 1-12
 - creating 2-6
 - defined 2-2
 - enabling subscribers for 3-28
 - including sections 3-27
 - vs. non-archived 2-3
- archived sections 1-11
- archiving
 - and sections 2-11
 - choosing 1-11
- asynchronous processes 2-26
- authorization 2-11

B

- BIG-IP Controller
 - setting up and adding 3-11
- BIG-IP Controller settings
 - reviewing or changing 3-12
- BIG-IP Controllers
 - managing 3-10
 - setting up and adding 3-9
 - viewing 3-12
- BIG-IP Detail screen 3-12
- BIG-IP List screen
 - deleting controllers 3-12
- BIG-IP Virtual Server Control 2-8
- BIG-IPs option
 - and help 1-3
- bigorb-server
 - as prerequisite 3-10
- browsers 1-5
 - connecting to home page 1-1

C

- Cancel delivery operation
 - and CopyPhase state 2-29, 3-34
- Cancelled state 2-29
- changes
 - tracking and saving 1-5
- child publications
 - as temporary 3-29
 - editing 3-29
- cleanup process 3-41
- command line utility 1-4
- commit phase
 - described 1-10
- commit. See content activation
- CommitPhase
 - See also publication states
- CommitPhase phase
 - defined 3-34
- configuration changes
 - reasons for 3-25
- configurations
 - changing 3-3, 3-24
 - synchronizing 1-7
- configurations supported 1-12
- Connection Test screen 2-14
 - viewing 2-21
- content
 - capturing for publication 2-21
 - collecting from source 2-3
 - comparing to version numbers 2-22
 - copying from source 3-2
 - delivering 2-9, 2-10, 2-24
 - gathering and storing 2-22
 - moving 2-7
 - organizing into sections 2-21
 - protecting for transfer 1-6
 - publishing 1-5, 2-18
 - storing as section version 2-3
 - synchronizing 1-7
 - tracking 1-1
- content activation
 - and controlled activation mode 3-8
 - and web access 3-14
 - as publishing phase 3-2
 - delaying 3-1, 3-3, 3-6
 - described 1-10, 3-6
 - determining 1-11, 3-2
 - immediate 3-1
 - on subscriber groups 3-6
 - scheduling 3-1, 3-2
- content activation controls
 - listed 3-6
- content activation process
 - determining 3-6
- content activation settings 3-4
- content activation speed 3-6
- content comparisons
 - and directories 3-18
 - purpose of 3-18
- content copying 3-18
- content delivery
 - ensuring 1-6
 - scheduling 3-1, 3-2
 - speeding up 3-1
- content distribution 1-5
- content distribution process 2-24
- content location
 - defining 2-11
- content management 1-6
- content mismatches 3-16
- content replication 3-16
- content rollback 1-6
- content servers
 - defined 2-13
- content sources 2-11
- content synchronization 1-5
- content update methods 2-22

- content updates 2-16
 - and performance 2-23, 3-10
 - for sections 2-11, 2-16
- content versions
 - keeping and creating 1-11
 - problems with 3-29
- contents
 - and versions 2-23
 - delivering 2-17
- continue delivery 2-29, 3-34
- continuous publications 3-1
- control buttons
 - displaying 2-29
- controlled activation mode
 - described 3-6, 3-8
- controller configuration
 - and virtual server nodes 3-10
 - example 2-4
- controller integration
 - benefits of 1-10
 - for scope and control 3-9
- controller phases
 - understanding 3-25
- controller reliability 1-12
- controllers
 - and web publishing tools 1-7
 - creating and updating 1-4
 - integrating 3-1
 - setting up 2-10
 - using secondary 2-19
 - verifying login to server 2-20
 - verifying login to servers 2-19
 - viewing or adding 1-3
- controller-to-controller communication 3-11
- copy phase
 - and changed files 3-19
- CopyPhase state
 - defined 3-34
 - See also publication states
- Create a New Edition screen 2-22, 2-25

- Create a New Section screen
 - viewing 2-12, 3-21
- CreateChangeTrees state
 - defined 3-33
 - See also publication states

D

- daily publications 3-1
- data
 - gathering 2-14
 - pulling from servers 1-5
 - synchronizing 1-5
- data loss
 - protecting against 1-12
- data management requirements 2-10
- date
 - displaying 2-26
- delayed activation option 1-6
- delays
 - setting 3-3
- Deliver screen 2-28
 - and disabling subscribers 3-27
 - and editions 2-22
 - and independent activation 3-8
 - and path disabling 3-27
 - using 3-5
 - viewing 2-27
- Deliver status 3-28
- delivery 2-5
- delivery control buttons 2-30, 3-34
- Delivery Detail screen
 - and subscribers 3-29
- delivery location 3-2
- delivery process
 - starting 2-28
- delivery status
 - viewing 2-27
- destination paths

- for sections 2-19
- detail states
 - table of 3-35
- directories
 - and post-publish phase 3-41
 - deploying 1-5
 - excluding 3-22
- directory cleanup 3-41
- directory lists
 - creating 3-18
- directory names
 - specifying 2-3
- directory paths
 - creating 3-24
 - example 3-24
- DIRLIST 3-18
- disk array
 - expanding 1-12
- disk drive types 1-12
- disk failure
 - protecting against 1-12
- disk mirroring 1-12
- disk size
 - and publishing 3-16
- disk space
 - and down time 1-12
 - recovering 3-31
- distribution 1-5
 - of applications 1-1
 - of file-based content 1-1
- Distribution List screen
 - viewing 2-20
- Distribution option
 - and help 1-3
- distribution profiles
 - defining 1-5
- distribution rules
 - setting up 1-5
- distributors

- and subscriber activation 3-9
 - and version removal 3-32
- defined 2-2
- pairing with publishers 1-7
- setting up and selecting 2-18, 2-19
- specifying 2-20

E

- Earth ground 1-9
- Edit Section List screen 3-21
- edition creation
 - manual 2-25
- edition creation methods 2-24
- edition details
 - viewing 1-2
- edition numbers
 - assigning 2-26
- edition status 1-6
- editions
 - and version removal 3-32
 - comparing 1-12
 - creating 1-3, 2-24
 - defined 2-1
 - maintaining 2-7
 - modifying and deleting 1-3, 3-31
 - publishing 2-24
 - republishing 1-6, 3-29
 - selecting previous 1-6
 - viewing 2-9
- Editions tab
 - and publication types 3-19, 3-21
- efficiency
 - maintaining 3-25
- email alerts
 - sending 1-12
- error alerts 2-8
- error conditions 3-33
- error notification 1-12
- error reporting 1-12

error-handling options
 setting 2-6, 2-10

exception paths
 and multiple sections 3-24
 creating 3-22

exceptions
 creating 3-2, 3-22, 3-23
 defined 2-11
 described 3-22
 example 3-22

F

F5-GLOBAL-SITE protocol 3-11

Failed state 3-41

failed subscribers 3-3

file activation
 defined 3-16

file availability 3-19

file changes
 determining 3-16

file cleanup 3-41

file comparisons 3-18

file contents
 storing 2-2

file copying 3-18

file sets
 updating 2-1

file-based content
 supported 1-7

files
 activating 3-8
 and post-publish phase 3-41
 capturing, storing, distributing 2-1
 delivering 2-1
 deploying 1-5
 listing 3-1
 removing from subscribers 2-25
 updating 2-1
 using as content 2-3

firewalls
 and publishing 1-7

First-Time Boot utility
 described 1-12

for sections
 updating 2-16

FTP access
 establishing 2-13, 2-20
 to systems 3-11
 to virtual server nodes 3-11

FTP accounts 2-12

FTP passwords 2-13

FTP user names 2-13

H

hardware
 usage guidelines 1-8

hardware installation
 planning 1-8

help button 1-4

help, technical support 1-14

home page
 illustration of 1-2

host names
 changing C-1

HTML 1-7

I

identifier
 See publication name

Idle state 2-27, 2-29, 3-3
 defined 3-33
 See also publication states

Idle status 3-27

Include status 3-26, 3-27

independent activation mode
 and web access 3-16

- described 3-6, 3-8
- indicators
 - for status 3-32
- information
 - storing 1-11
- integration
 - with BIG-IP Controllers 1-10
- internal interfaces 3-10
- Internet sites
 - distributed 1-1
- interval updates 3-3
- IP addresses
 - changing C-1

J

- Java Scripts 1-7

L

- lithium battery 1-9
- locator
 - described 1-4
- logical publications
 - determining 2-4

M

- manual deliveries 3-3
- manual publications 3-5, 3-6, 3-9
- manual publishing 1-6, 2-7
- Microsoft Windows applications 1-7
- mirroring 1-12
- mismatched content 3-16

N

- navigation pane
 - described 1-3

- network traffic
 - reducing 1-5
- New Publication screen
 - and activation 3-5
 - and content activation 3-4, 3-7
 - and non-archived publications 3-20
 - opening 2-6
 - using 2-5
- non-archived publications
 - and Publication Sections screen 2-9
 - and versions 3-29
 - as attribute 2-5
 - benefits of 3-16, 3-17
 - comparing versions 1-12
 - creating 2-6, 3-1, 3-20
 - creating, managing, publishing 3-19
 - deciding to use 3-17
 - described 2-2, 3-16
 - enabling and disabling subscribers 3-28
 - ignoring sections 3-27
 - vs. archived 2-3
- non-archived sections
 - creating 3-17, 3-21
- normal server subscribers 3-13

O

- offline subscribers
 - excluding sections 3-26
 - including sections 3-27
 - marking 3-25
- operating system platforms 1-4, 1-7
- organization
 - changing 3-24
- out-of-date subscribers 3-3, 3-29

P

- parameters
 - changing 2-14
- password validity 2-13

- passwords 2-11
- path validity 2-13
- paths
 - disabling 3-25
 - identifying 2-11
 - listing for sections 2-14
- performance
 - controlling and maintaining 1-10, 3-1, 3-10
 - improving 1-12
- post-publish phase
 - and sections 3-41
 - described 3-41
- post-publishing states 3-41
- procedures
 - finding online 1-14
- processes
 - monitoring 2-16
- profiles
 - defining 1-5
- project data
 - synchronizing 1-5
- project history availability 1-5
- publication
 - and firewalls 1-7
 - creating 2-5
- publication configuration
 - changing 3-24
- publication cycles 2-21
- publication delivery 2-5
 - stopping 2-29
- Publication Editions screen
 - creating editions 2-22
 - delivering editions 2-22
 - described 1-2, 1-3, 2-9
 - viewing 2-8, 2-26
 - viewing section details 2-26
- publication efficiency
 - maintaining 3-25
- publication information
 - changing 3-27
 - retaining 3-30
- publication links 3-29
- Publication List screen
 - described 1-2
 - for status 2-26
 - unfamiliar entries 3-29
 - viewing 1-1
 - viewing new publications 2-8
- publication management 2-5
- publication name 2-5
- Publication Options screen 2-7
 - and activation 3-5
 - and content activation 3-7
 - and defaults 3-3
 - confirming status 1-6
 - for web access 3-14
 - setting options 2-10, 3-4
- publication organization
 - viewing 1-6
- publication planning 3-22
- publication reliability
 - maintaining 3-25
- publication screens
 - using 2-9
- Publication Sections screen
 - and content delivery 2-10
 - and non-archived sections 3-21
 - starting 2-16
 - updating content 2-22
 - viewing 2-12, 2-15, 3-21
- publication settings
 - establishing 2-24
- publication state
 - resetting 2-29, 2-30
- publication states 3-33
 - and control buttons 2-29
 - listed 3-33
 - table of 3-35

- publication status
 - and publishing cycles 2-30
 - and screens 1-6
 - changing 2-28
 - updating 2-27
 - viewing 1-1, 2-26
- publication status indicators 3-32
- Publication Subscribers screen
 - and removal 3-30
 - and server setup 3-13
 - and versions 3-29
 - listing subscribers 2-10
 - viewing 2-17
- PublicationPreparation state
 - defined 3-33
 - See also publication states
- publications
 - and multiple sections 3-6
 - and relationships 2-4
 - and subscriber synchronization 3-29
 - archiving 1-11
 - automating 3-2
 - changing 2-4, 3-3
 - controlling and managing 3-24
 - creating 3-1
 - defined 2-1
 - defining 2-4
 - deleting 1-2
 - delivering 2-2, 3-2
 - described 1-1
 - disabling 2-13
 - example 2-5
 - multiple 2-23
 - publishing and updating 1-1
 - receiving 2-17
 - scheduling 3-2, 3-19
 - setting up 3-19
 - viewing 1-1, 2-9
- publications control
 - increasing 3-9
- publications details
 - viewing 1-2
- publications options
 - accessing 1-2
 - creating, modifying, deleting 1-3
- publications overview 1-2
- publications process
 - starting and monitoring 1-3
 - understanding 1-3
- publications processes
 - viewing status 2-26
- publications status 1-2
- Publish Progress Display screen 2-16
 - by-passing delays 3-8
 - changing status 2-28
 - described 1-6
 - opening 3-29
 - updating content 2-22
 - viewing 2-28
- Publish Progress screen
 - described 1-6
- publisher/distributor configuration 1-7
- publishers
 - defined 2-2
 - pairing with distributors 1-7
- publishing
 - of applications 1-1
 - of file-based content 1-1
 - setting up 3-2
- publishing automation 3-3
- publishing behavior 1-7
- publishing cycle transitions 2-30
- publishing cycles
 - defined 2-3
 - running 3-29
- publishing editions
 - comparing 1-12
- publishing error notification 1-12
- publishing intervals 1-11
- Publishing option
 - and help 1-3

- publishing options
 - setting 2-6
- publishing phases
 - post-publishing 3-41
 - understanding 3-32
- publishing processes
 - and missed updates 3-3
 - and non-archived publications 3-16
 - asynchronous 2-26
 - controlling 1-5
 - delaying 1-11, 3-3
 - initiating 3-5
 - managing 2-1
 - monitoring 1-6
 - pausing and restarting 3-5
 - phase by phase 3-8
 - regulating 3-2
 - scheduling 1-10, 3-1, 3-3, 3-5
 - starting manually 1-11
 - stopping 1-6, 2-29, 3-34
- publishing scope
 - increasing 3-9
- publishing speed 3-16
- publishing states
 - distributors vs. subscriber paths 3-35
- publishing status
 - confirming 1-6
 - viewing 2-27
- publishing times 1-11

R

- rack mounting 1-8
- related file sets
 - defining 2-1
- relationships
 - between sections and web servers 2-17
 - establishing 2-1
- reliability
 - adding 3-2
 - maintaining 3-25

- remote controllers 2-2
- reset delivery 2-30, 3-34
- retry delivery 2-30, 3-34
- rollback
 - for subscribers 3-41
- rollback states 3-41
- root paths
 - creating 3-24
- rules
 - setting up 1-5

S

- schedule settings 3-5
 - making 3-4
 - revising 3-3
- Schedule state
 - See also publication states
- scheduled publications 3-6, 3-9
- scheduled publishing
 - pausing and restarting 3-5
 - specifying 2-7
- Scheduled state
 - defined 3-34
- scheduled updates 3-3
- SchedulePaused state
 - defined 3-34
 - See also publication states
- schedules
 - and missed updates 3-3
 - pausing 3-3
 - setting 3-2, 3-3
- scheduling
 - pausing and restarting 3-4
- scheduling options
 - setting 2-6, 2-10
- screen controls
 - listing 1-14
- secondary controllers

- using 2-19
- section contents 3-22
- section creation
 - example 2-12
- section destinations
 - specifying 2-16
- Section Detail screen
 - and exceptions 3-22, 3-23
 - and source detail 2-13
 - changing parameters 2-14
 - listing versions 2-23
 - removing versions 3-31
 - testing connections 2-14
 - updating content 2-22
 - viewing 2-26
- section details
 - viewing 2-28
- Section Details screen 2-13
- section exceptions
 - creating 3-2
 - described 3-22
- Section File Listing screen
 - described 1-6
 - viewing 2-26
- section groups
 - managing 2-17
- section names
 - syntax 2-12
- section parameters
 - changing 2-15
- section paths
 - listing 2-14
- section sizes
 - and efficiency 3-18
- section sources 2-11
- section states 3-33
- section status 1-6
- section unavailability 2-29
- section updates
 - stopping 2-29
- section version differences
 - comparing 1-12
- section versions
 - accessing 1-2
 - creating 2-14
 - defining 2-25
- section/web server relationships 2-17
- sections
 - accessing 1-2
 - adding to publications 2-13
 - and archiving 2-11
 - and destination paths 2-19
 - and directory paths 2-13
 - and post-publish phase 3-41
 - and versions 2-26
 - changing 2-15
 - changing parameters 2-15
 - contents 2-10
 - creating 2-11, 2-12
 - creating, modifying, deleting 1-3
 - defining 2-4
 - described 2-1
 - determining 2-4
 - displaying 2-25
 - excluding 3-26
 - including and ignoring 3-27
 - including for subscribers 2-19
 - including in editions 2-13, 2-19
 - managing 2-9
 - naming 2-11
 - number needed 2-10
 - organizing into publications 2-21
 - re-enabling for subscribers 3-27
 - setting up 3-24
 - sharing 2-23
 - showing source details 2-13
 - specifying passwords 2-19
 - testing 2-13
 - updating content 2-16
 - viewing and modifying 1-1
- security
 - maintaining 1-7

- security mechanisms 1-6
- server
 - specifying 2-11
- server access 2-11, 2-13, 2-19
- server names 2-18
- server paths
 - specifying 2-3
- server subscribers
 - adding
- servers
 - and subscribers 2-18
- simple persistence
 - enabling 3-16
- simultaneous activation 3-9
- site performance
 - controlling and maintaining 1-10
- SNATs (Secure Network Address Translations) 3-11
- source content
 - reviewing 3-18
- source server names
 - listing 2-14
- source server paths
 - listing 2-21
- source servers
 - defined 2-13
- SSL encryption
 - using 1-6
- states
 - distributors vs. subscriber paths 3-35
 - for rollback 3-41
 - listed 3-33
 - table of 3-35
- static relationships
 - establishing 2-1
- status
 - changing 2-28
 - confirming 1-6
 - providing 3-33
 - updating 2-27
 - viewing 1-1
- status indicators 3-32
- status reporting 1-12
- storage 2-5
- storage capacity
 - increasing 1-12
- subscriber servers 2-18
- subdirectories
 - excluding 2-13, 3-22
 - excluding from sections 3-2
- subscriber access
 - disabling during activation 3-15
- subscriber activation
 - as option 3-8
 - specifying 3-15
- subscriber capacity
 - and publishing 3-16
- Subscriber Detail screen
 - and sections 3-27
 - for disabling paths 3-25
 - viewing 2-21
- subscriber path details
 - viewing 2-28
- subscriber path states 2-28, 3-33
- subscriber paths 3-29
 - cleaning up files 3-41
- subscriber rollback 3-41
- subscriber server capacity 3-19
- subscriber unavailability 2-29
- subscribers
 - accessing 1-2, 3-9
 - adding 2-18, 2-20
 - and copy phase 2-29
 - and virtual servers 3-9
 - changing 2-19
 - checking versions 3-29
 - configuring 2-17
 - creating, modifying, deleting 1-3

- defined 2-17
- described 2-1
- enabling and disabling 3-28
- excluding 3-27
- identifying 2-3
- including 3-28
- listing 2-10, 2-18
- marking offline 3-25
- removing 3-30
- See also offline subscribers
- See also server subscribers
- See also virtual server subscribers
- synchronizing 3-29
- updating 3-3
- synchronization
 - of applications 1-1
 - of file-based content 1-1, 1-5
- synchronization types 1-7
- System option
 - and help 1-3
- System Settings screen
 - and scheduling 3-4
 - checking time 3-5
 - stopping and starting scheduler 3-5
- system time
 - setting 1-3

T

- target servers. See subscribers
- Technical Support 1-14
- temporary files
 - changing to permanent 1-10
- Test Connection function 2-13, 2-20
- test results
 - displaying 2-14, 2-21
- time
 - displaying 2-26
- Transfer Method 2-13, 2-19

- troubleshooting
 - performing 1-4

U

- unavailable subscribers
 - marking 3-25
- update process
 - and web access restriction 3-10
- update processes
 - and humans 3-2
 - sending changes 3-3
- update reliability 3-2
- updates
 - and intervals 3-3
 - transmitting 1-5
- UpdatingSections state
 - defined 3-33
 - See also publication states
- user accounts
 - removing 1-3
- user ID validity 2-13
- user IDs 2-11, 2-18
- user interface screens
 - described 1-4
- user interface types 1-1

V

- ventilation 1-8
- version access
 - providing 1-5
- version information
 - retaining 3-30
- version lists
 - creating 3-1
 - keeping 2-2
 - keeping and creating 1-11
 - purpose of 1-11

- version numbers
 - assigning 2-26
- version storage 2-5
- versions
 - accepting default 2-26
 - checking current 3-29
 - creating 2-14, 2-23
 - defined 2-1
 - keeping 1-11
 - listing 2-23
 - listing, creating, removing 2-25
 - purging 3-31
 - storing 2-14, 2-22
 - time to create 2-23
 - viewing and modifying 1-1, 2-23
- virtual server controls
 - setting 2-10
- virtual server nodes
 - removing 3-10
- virtual server subscribers
 - controlling 3-9
 - creating 3-11
 - setting up 3-13
- virtual servers
 - showing 3-12
 - using as distributor 3-11
 - using as publisher 3-11
- voltage 1-9
- and content 1-5
- defining distribution profiles 1-5
- selecting editions 1-6
- web publishing tools
 - and controllers 1-7
- web site disruption
 - minimizing 3-6
- web users
 - locking into single node 3-16
- website paths
 - listing 2-21
- weekly publications 3-1

W

- WANSync state
 - defined 3-33
 - See also publication states
- web access
 - and updates 3-1
 - disabling and re-enabling 3-14, 3-16
 - disabling during activation 3-14
 - restricting 3-1, 3-10
- web browsers 1-5
- web publishers