GLOBAL-SITE Administrator Guide

version 2.2

MAN-00xx-xx

Service and Support Information

Product Version

This manual applies to product version 2.2 of the GLOBAL-SITE® Controller.

Obtaining Technical Support

Web	tech.f5.com
Phone	(206) 272-6888
Fax	(206) 272-6802
Email (support issues)	support@f5.com
Email (suggestions)	feedback@f5.com

Contacting F5 Networks

Web	www.f5.com
Toll-free phone	(888) 88.BIG.IP
Corporate phone	(206) 272-5555
Fax	(206) 272-5556
Email	sales@f5.com
Mailing Address	401 Elliott Avenue West 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.

Trademarks

F5 and the F5 logo, F5 Networks, BIG-IP, 3-DNS, SEE-IT, GLOBAL-SITE, and EDGE-FX are registered trademarks of F5 Networks, Inc. FireGuard, iControl, and Internet Control Architecture are trademarks of F5 Networks, Inc. In Japan, the F5 logo is trademark number 4386949, BIG-IP is trademark number 4435184, 3-DNS is trademark number 4435185, and SEE-IT is trademark number 4394516. All other product and company names are registered trademarks or trademarks of their respective holders.

Export Regulation Notice

This product may include cryptographic software. Under the Export Administration Act, the United States government may consider it a criminal offense to export this product 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.

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.

Canadian Regulatory Compliance

This class A digital apparatus complies with Canadian I CES-003.

Standards Compliance

The product conforms to ANSI/UL Std 1950 and Certified to CAN/CSA Std. C22.2 No. 950.

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.

This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit (http://www.openssl.org).

This product includes cryptographic software written by Eric Youg (eay@cryptsoft.com).

This product includes software written by Tim Hudson (tjh@cryptsoft.com).

Table of Contents

1
Introduction to the GLOBAL-SITE Controller

2
Working with the First-Time Boot Utility

3

Publishing Documents

4

Increasing Control over Publishing

5

Monitoring Publishing Processes

6
ITCMconsole Command Line Interface

1

Introduction to the GLOBAL-SITE Controller

- Welcome to the GLOBAL-SITE Controller
- GLOBAL-SITE Controller features
- What's new in version 2.2
- Finding additional help

Welcome to the GLOBAL-SITE Controller

The GLOBAL-SITE Controller is a network appliance that organizes and manages distribution and synchronization of file-based content and applications to local and geographically distributed Internet sites. The GLOBAL-SITE Controller is flexible, and has been designed to maximize your file management options. The GLOBAL-SITE Controller can:

- Retrieve file content using a variety of access methods, or it can accept content that is delivered directly to the GLOBAL-SITE Controller itself.
- Replicate content from multiple sources to one or more systems using a variety of delivery mechanisms.
- Deliver content directly from a source to a target, or deliver content indirectly, through another GLOBAL-SITE Controller that has access to the sources.
- Handle errors at multiple levels, with the goal of delivering complete, consistent sets of updates to all targets systems.

The GLOBAL-SITE Controller uses a publishing metaphor, organizing your content as publications, and tracking the delivery of editions. 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.

Application tools

The GLOBAL-SITE Controller supports both a browser-based application and an administrative command line utility. The browser application provides access to all GLOBAL-SITE Controller features and functions. Figure 1.1 shows the Configuration utility home screen for the GLOBAL-SITE Controller. This is the first screen you see when you start the GLOBAL-SITE Controller.

Figure 1.1 *Configuration utility home screen*

The browser application

When you click the **Configure the GLOBAL-SITE Controller using the Configuration utility** link it, opens the Publication List screen (Figure 1.2). This screen gives you an overview of your publications, the status of each, and a quick way to delete any publication. In addition, it provides a link to the Publication detail screens where you can view further details about each publication. Until you create a publication, the list screen is blank.

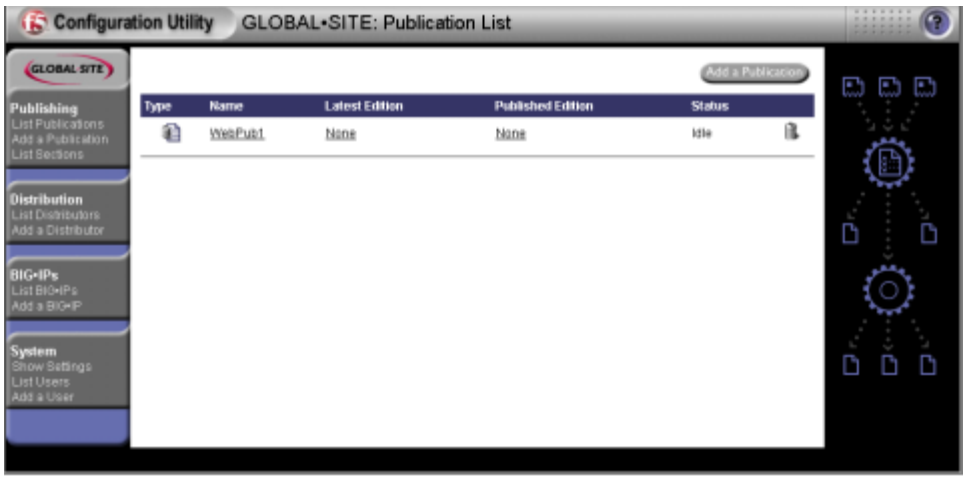

Figure 1.2 *The GLOBAL-SITE Controller Publication List screen*

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 detail screens that display the basic building blocks of your publications, provide the 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 Pub. Options (non-archived publications do not show an Editions tab). Each tab screen provides an overview of the element and access to the status and details for that element. You can typically create, modify, or delete each element, and add or remove that element from a publication from these screens. You also use the publication screens to start and monitor the publication process.

The ***navigation pane***, along the left side of the screen, 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 open the corresponding screen.

The screens available from the **Distribution**, **BIG-IPs**, and **System** navigation pane items 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 units and virtual servers, for adding or removing user accounts, and for setting or viewing system settings.

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 online help for the screen you are viewing.

Specifications

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

Operating systems

The GLOBAL-SITE Controller runs on Redhat$_{TM}$ Linux version 7.0.

The GLOBAL-SITE Controller supports a variety of source and target server operating systems:

- Windows NT® and Windows® 2000 operating systems
- Microsoft® Windows® 95, 98, and Millennium Edition
- UNIX platforms
- Mac® Operating System

Web browsers

You can connect to the GLOBAL-SITE Controller using either of the following popular web browsers:

- Microsoft® Internet Explorer, version 6.0
- Netscape® Navigator, version 4.7X

GLOBAL-SITE Controller features

The GLOBAL-SITE Controller offers a variety of management and control features for content synchronization and distribution.

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 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 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, you can immediately republish any previous edition to production web servers 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. Transferring files using WebDAV with digest authentication or SSL provides additional security as content files are passed between GLOBAL-SITE Controllers and sections or subscribers.

Visibility into the 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 on 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 display the 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 stops at selected times so that you can confirm the status before resuming 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 Mac® 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
- Self-registering DLLs
- 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 the subscribers.

What's new in version 2.2

This section introduces the main features that are new to version 2.2 of the GLOBAL-SITE Content Controller, and tells where you can find detailed information about each feature.

iControl support

The GLOBAL-SITE Controller now supports the content delivery components of the iControl™ SDK. Using the iControl SDK you can get information on distributors or local load balancers, control the delivery process, get publication details and subscriber status, and more. Please see the iControl documentation for details.

Command line support

The GLOBAL-SITE Controller now supports a new command line interface called ITCMconsole, which allows you to configure the controller and gives you greater control over your system. Please see chapter 6, *ITCMconsole Command Line Interface* in this guide for more information, including a complete alphabetic list of the available commands.

SNMP event notification

The GLOBAL-SITE Controller now provides the UC Davis MIB, as well as the GLOBAL-SITE MIB, which displays GLOBAL-SITE specific information. The notification system can be extended with user-defined actions such as sending and email message or running a script. In addition, users can define and raise their own alerts using the new alert notification system. Please see *Working with SNMP and alerts*, on page 7-1 of this guide for more information.

Republish all content

You can now republish all content to your subscribers, or to specific subscribers (for instance to all cache subscribers), regardless of what is on the subscriber machine. This is done on the Deliver screen from the **Publication Option** box. Please see the online help on the Deliver or Deliver Edition screens for more information.

Controlling servers with the GLOBAL-SITE agent (Linux)

The GLOBAL-SITE Controller now supports the Linux agent (Red Hat 6.1 and later), in addition to the Microsoft Windows 32 agent and the Solaris 2.6 agent. You can download this agent from the first page you see when you open the GLOBAL-STIE Controller, under the Additional Software Downloads heading. For details about the GLOBAL-SITE agent, see *Controlling servers with the GLOBAL-SITE agent*, on page 4-20 and *Understanding file transfer methods*, on page 4-18 of this guide.

Windows® Media Technology™ push

Windows Media Technology (WMT) push is now available for non-archived publications. Please see *Supported URL schemes*, on page 4-11, for more information.

Dual NIC support

The GLOBAL-SITE Controller now supports two network interface cards that can be configured from the First-Time Boot utility (see *Configuring the interfaces*, on page 2-3), or from the ITCMconsole (see *Configuring interfaces*, on page 6-19).

Licensing control

The GLOBAL-SITE Controller starts only if authorized to do so by a licensing software key. You can obtain a licensing key from a Technical Support representative.

Gigabit Ethernet card option

The GLOBAL-SITE Controller now supports a Gigabit Ethernet network interface card, which provides higher throughput. This is available from F5 Networks as a separately-purchased option on the controller.

Update restrictions

The GLOBAL-SITE Controller now restricts non-archived updates to 16 sections at a time, and restricts archived updates to 5. These limits are set to overlap the network and CPU load when the controller updates multiple sections together.

What's new from version 2.1

The following features were new in the GLOBAL-SITE Controller version 2.1, but were not included in an update of the Administrator Guide at that time.

Custom scripts

The GLOBAL-SITE Controller can now run scripts on a subscriber server while delivering a publication. For details on how to activate this feature, view the online help on the Add Subscriber screen. (To get to this screen, open the Configuration utility, click a publication, click the Subscribers tab, click **Add a Subscriber**, and click the Agent Actions tab.)

Solaris 2.6 agent

The GLOBAL-SITE Controller now supports the Solaris 2.6 agent, in addition to the Microsoft Windows 32 agent. You can download this agent from the first page you see when you start the GLOBAL-SITE Controller, under Additional Software Downloads.

Finding additional help

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 by clicking **Release Notes** on the GLOBAL-SITE Controller home screen. 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***
  The Configuration utility home screen has a link to the PDF version of this administrator guide. If you upgrade your GLOBAL-SITE Controller software at a later time, the upgrade may include an updated PDF version of the guide.

  A PDF version of this administrator guide is also 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.

- **Getting started help**
  Each item 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.

2

Working with the First-Time Boot Utility

- Using the First-Time Boot utility
- Gathering configuration information
- Starting the First-Time Boot utility
- Configuring settings for the GLOBAL-SITE web server
- Configuring the LDAP server
- Confirming your configuration settings
- Committing your configuration settings to the system

Using the First-Time Boot utility

The First-Time Boot utility is a wizard that walks you through a brief series of tasks to complete the required configuration for the GLOBAL-SITE Controller. The utility has you do things like define a root password and configure IP addresses for the network interfaces.

The First-Time Boot utility's three phases are: configure, confirm, and commit. Each phase walks you through a series of screens, so that you can configure the following settings:

- End-user software license
- Root password
- GLOBAL-SITE password
- Settings for optional technical support access
- Enable support account
- Host name
- Default route
- interfaces
- Nameserver
- NTP server
- Time zone
- System time
- SSH
- Telnet
- FTP log file access
- GLOBAL-SITE Identifier
- Settings for the GLOBAL-SITE web server
- Self signed SSL certificate information
- LDAP server

First, you configure all of the required information. Next, you have the opportunity to correct, if necessary, and confirm each individual setting that you have configured. Last, your confirmed settings are committed and saved to the system.

Gathering configuration information

Before you run the First-Time Boot utility on a specific GLOBAL-SITE Controller, you should have the following information ready to enter:

- Passwords for the root system, for the GLOBAL-SITE web server, for the GLOBAL-SITE Controller, and for technical support access (optional)
- Host name for the GLOBAL-SITE Controller

- A default route (typically a router's IP address)
- Settings for the network interfaces, including IP addresses, media type, and custom netmask
- The IP address or IP address range for remote administrative connections

Starting the First-Time Boot utility

The First-Time Boot utility starts automatically when you turn on the GLOBAL-SITE Controller (the power switch is located on the front of the controller). The first screen the controller displays is the License Agreement screen. You must scroll through the screen, read the license, and accept the agreement before you can move to the next screen. If you agree to the license statement, the next screen you see is the Welcome screen. From this screen, simply press any key on the keyboard to start the First-Time Boot utility, and then follow the instructions on the subsequent screens to complete the process.

Note

You can change any configurations at a later time using the ITCMconsole command line utility. See Chapter 6, ***ITCMconsole Command Line Interface****, for more information.*

Defining a root password

You gain administrative access to the GLOBAL-SITE Controller by using a root password. The root password must contain a minimum of 6 characters, but no more than 15 characters. Passwords are case-sensitive, and we recommend that your password contain a combination of uppercase and lowercase characters, as well as special characters. Once you enter a password, the First-Time Boot utility prompts you to confirm your root password by typing it again. If the two passwords match, your password is immediately saved. If the two passwords do not match, you receive an error message asking you to re-enter your password.

WARNING

The root password is the only setting that is saved immediately, rather than confirmed and committed at the end of the First-Time Boot utility process. You can change the root password from the ITCMconsole, after the First-Time Boot utility completes and you reboot the GLOBAL-SITE Controller. (To change the root password after initial configuration, you use the ITCMconsole command ***user****. See* ***user****, on page 6-33 for more information.) You can change other system settings when the First-Time Boot utility prompts you to confirm your configuration settings or at a later time using the ITCMconsole.*

Defining a GLOBAL-SITE password

Enter a password that is between 6 and 15 characters long. You only use the GLOBAL-SITE password when you log into the GLOBAL-SITE Controller as **gsite** to access the GLOBAL-SITE command line utilities. Those utilities are separate from the ITCMconsole. You can configure everything you need on the web-based user interface or from the ITCMconsole and should rarely, if ever, need to access this other command line utility.

Allowing Tech Support

If you want to allow Technical Support access to your GLOBAL-SITE Controller, choose to enable a Tech Support account. The next screen asks you to set and confirm the tech support password, which can be between 6 and 15 characters long. If you do not want a tech support account, you go to the next screen.

Defining a host name

The host name identifies the GLOBAL-SITE Controller itself. Host names must be in the format of a fully-qualified domain name. Host names may contain letters, numbers, and the symbol for dash (**-**), however, they may not contain spaces. For example, if the controller's label is **controller1**, then you define the host name as **controller1.yourdomain.com**.

Configuring a default route

If a GLOBAL-SITE Controller does not have a predefined static route for network traffic, the controller automatically sends traffic to the IP address that you define as the default route. Typically, a default route is set to a router's IP address.

Configuring the interfaces

The First-Time Boot utility lists available interfaces that are available to configure. You need to type the IP address and netmask, and choose the media type for the interface. Use the arrow keys to move to the interface you want to configure. You must at least configure the **eth0** interface. Use the keyboard key **C** to continue with the First-Time Boot utility once you are through configuring interfaces.

- **IP address**
  This is the IP address of the controller itself.
- **Netmask**
  Type a netmask appropriate to the subnet indicated by the IP address.

- **Media type**
  The media type options depend on the network interface card included in your hardware configuration.

Configuring nameserver

Enter an external nameserver to forward DNS requests to.

Configuring NTP clocks

You can synchronize the time on your GLOBAL-SITE Controller to a public time server by using Network Time Protocol (NTP). NTP is built on top of IP and assures accurate, local timekeeping with reference to clocks located on the Internet. This protocol is capable of synchronizing distributed clocks, within milliseconds, over long periods of time. If you choose to enable NTP, make sure UDP port 123 is open in both directions when the GLOBAL-SITE Controller is behind a firewall.

Configuring a time zone

Configuring a time zone ensures that the clock for the GLOBAL-SITE Controller is set correctly, and that dates and times recorded in log files correspond to the time zone of the system administrator. Scroll through the time zone list to find the time zone closest to your location. Note that one option may appear with multiple names.

Configuring system time

You must update the system time if it is more than 100 minutes off in order for the NTP daemon to work correctly.

Configuring remote administration

When you configure remote administration, the first screen you see is the Configure SSH screen, which prompts you to type an address for **ssh** command line access.

The First-Time Boot utility prompts you to enter a single IP address, or a range of IP addresses, from which the GLOBAL-SITE Controller can accept administrative connections (either remote shell connections, or connections to the GLOBAL-SITE web server). The following example demonstrates how to specify a range of IP addresses. This example allows remote administration from all hosts on the 192.168.2. subnetwork:

```
192.168.2.
```

Configuring services

Choose to configure any or all of the following:

- **Telnet**
  Turn Telnet on or off, and set allowed hosts for Telnet. The default is off.
- **FTP**
  Turn FTP on or off. The default is off. Turning FTP on enables FTP access to the GLOBAL-SITE Controller, therefore we recommend you choose to keep FTP turned off.

Entering the GLOBAL-SITE Identifier

You should have received a GLOBAL-SITE Identifier from your support representative. Enter that number here. If you do not have an Identifier, please contact Technical Support.

Configuring settings for the GLOBAL-SITE web server

The GLOBAL-SITE web server requires that you define a domain name for the server, a user ID, and a password. The GLOBAL-SITE web server hosts the web-based Configuration utility. The information that you configure in these screens allows you to access the Configuration utility from a web browser on your workstation.

The First-Time Boot utility guides you through a series of screens to set up web server access:

- The first web server screen prompts you for a user name (the default is **gsite**) and a password. The password does not show on screen as you type it. The utility prompts you to enter the password again for confirmation purposes.
- The next web server screen prompts you to set web access. Enter the IP addresses that are permitted access to your web-based user interface, or type **all** to allow all access.
- The certification screen prompts you to enter the management port number (the default is **443**), country, state, city, company, and division information used for the authentication certificate.

Configuring GLOBAL-SITE web server passwords

You can add additional web users or change the password for this web user account with the ITCMconsole **webuser** command.

Configuring the LDAP server

You can now set up a common authentication and authorization data store for the GLOBAL-SITE Controller you use. Choose **Standalone or Master server** if the controller is to be the LDAP server, or **Remote LDAP server** if you already have an LDAP server you want to use. If you choose a remote LDAP server, you must provide the host name and password for that server.

Confirming your configuration settings

At this point, you have entered all the configuration information, and now you confirm each setting. Each confirmation screen displays a setting and prompts you to either accept or re-enter it. If you choose to edit the setting, the utility displays the original configuration screen in which you defined the setting the first time. When you finish editing the item, you return directly to the Confirmation screen for that item, and continue the confirmation process. Note that once you accept a setting in the Confirmation screen, you do not have another opportunity to review it before the commit phase.

You confirm or edit the settings in the same order that you configured them: Once you have confirmed the last setting, the First-Time Boot utility moves directly into the commit phase, where you are not able to make any changes.

Committing your configuration settings to the system

Once you confirm all of the configuration settings, the First-Time Boot utility saves the configuration settings. If you want to change any information at a later time, you can change certain settings in the web-based user interface, or you can change the settings using the ITCMconsole command line utility.

3

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 files and/or file lists on the GLOBAL-SITE Controller for 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 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 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 can 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, plan what to do, and plan how to do it. This helps save time because you can set things up initially in the way that is the most efficient for you in the long run.

You need to define:

- Where to look for new content
- Where to put new content
- Whether to store content on the GLOBAL-SITE Controller

You can use any file as content as long as you can specify a path to the server and directory that store 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: the content of each file included in the section and 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 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*, on page 4-15.

Defining publications

The key to configuring a GLOBAL-SITE Controller is:

- Determine what the logical publications are
- Determine which sections are in each publication

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 a 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:

- The sections that are published together
- 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 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 can be archived or non-archived. You determine this attribute when you create the publication. 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*, on page 4-15, and in the online help.

Creating the publication

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

For our example, assume you are creating a publication called **WebPub1**.

◆ 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 guide. The screen names are shown in the figure caption for each screen.

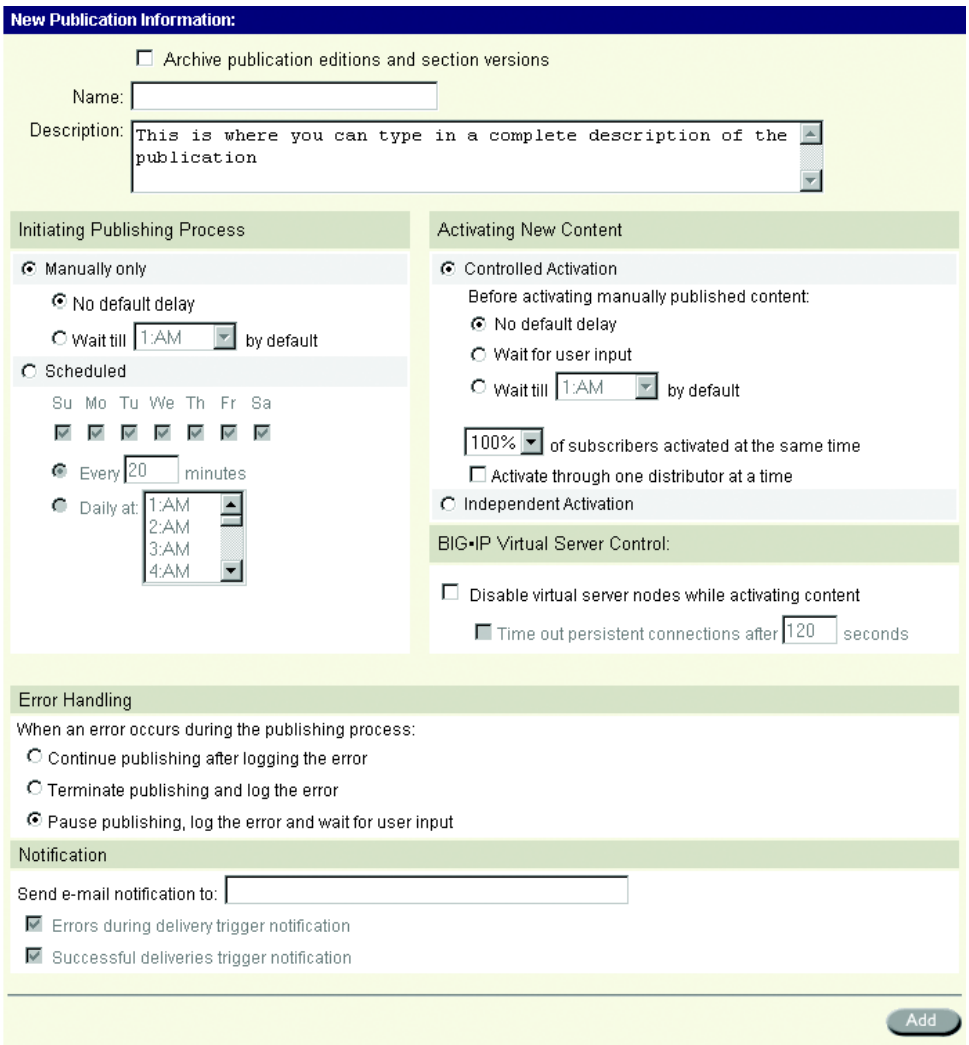

Figure 3.1 *The New Publication screen*

There are boxes 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

1. In the navigation pane, click **List Publications**.
   The Publication List screen opens.
2. Click the **Add a Publication** button in the top right corner of the screen.
   The New Publication screen opens.

3. Make the necessary changes as required. For further information regarding this screen, click the Help button found in the upper right corner of the screen.

Once you have created the new publication, you can click its name to open the Publication Editions screen (for archived publications) or the Publication Sections screen (for non-archived publications).

Note

You cannot change the archived/non-archived option once you create a publication. For all other options, you can set or change them at a later time from the Publication Options screen. For more information about publication options, click the Help button on the Publication Options screen.

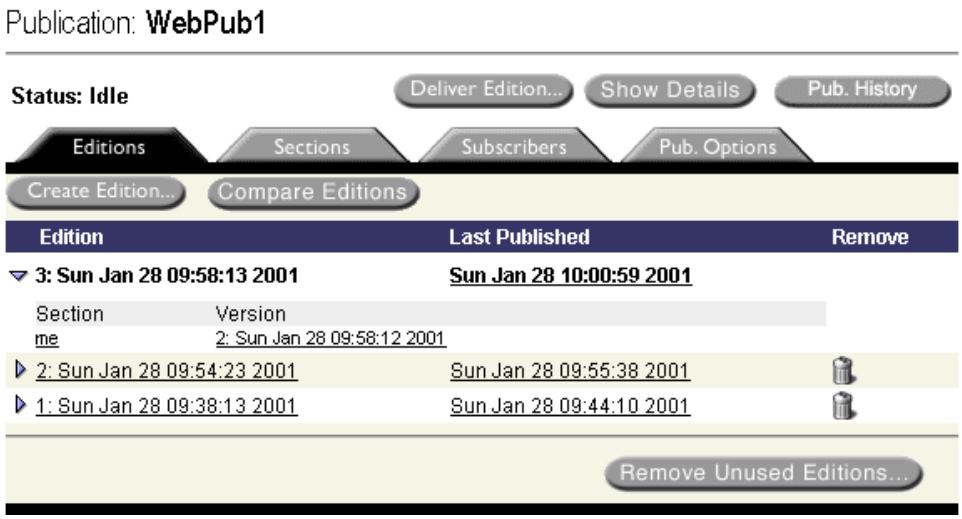

Figure 3.2 *The Publication Editions screen*

Using the Publication detail screens

The Publication detail 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 see. There are 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 never see the Editions tab.

Click any tab to display a detail screen of the same name for that aspect 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 do not see the Publication Editions screen.
- 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 is delivered.
- The Publication Options screen is where you set options for the delivery process including scheduling, error handling, and virtual server controls.

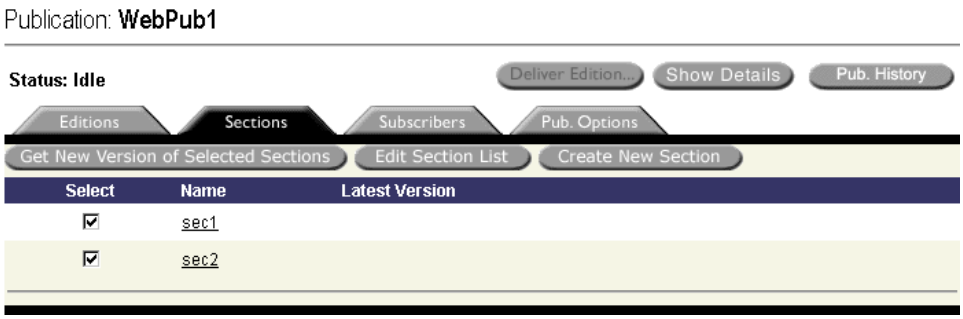

Figure 3.3 *The Publication Sections screen*

Defining sections

Before delivering content, you need to tell the GLOBAL-SITE Controller where to find it. You do that by clicking the Sections tab and using the Publication Sections screen.

A GLOBAL-SITE Controller section includes both content files and a path to the content files.

Before you 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 (path, directory, and file names are all case sensitive). A section can exclude any specified subdirectories, known as ***exceptions***. Exceptions are directories that are in the path, but are not included in the section. You can also include or exclude all files with specific file extensions, such as **.cgi** and **.html**.

- **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. Be sure to distinguish names using characters, not just uppercase versus lowercase letters; not all operating systems are case-sensitive.
- **What is the path to each section?**
  You must identify the specific directory path on the server that is the source of the section's content. Note that some operating systems are case sensitive; be sure to get the paths exactly right.
- **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 created a publication, is creating sections. Updating the content for a section is one of the most frequent tasks you perform on a GLOBAL-SITE Controller. To create a section, you:

- Give the section a name
- Specify the server that contains the section
- Provide access settings: the user ID and password for the server
- Define the location of the content with a path

The section you create can 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 **sec1**, which gets its content from the directory **/home/webcontent/testsite** on the system **myserver.f5.com** using the FTP account **website**.

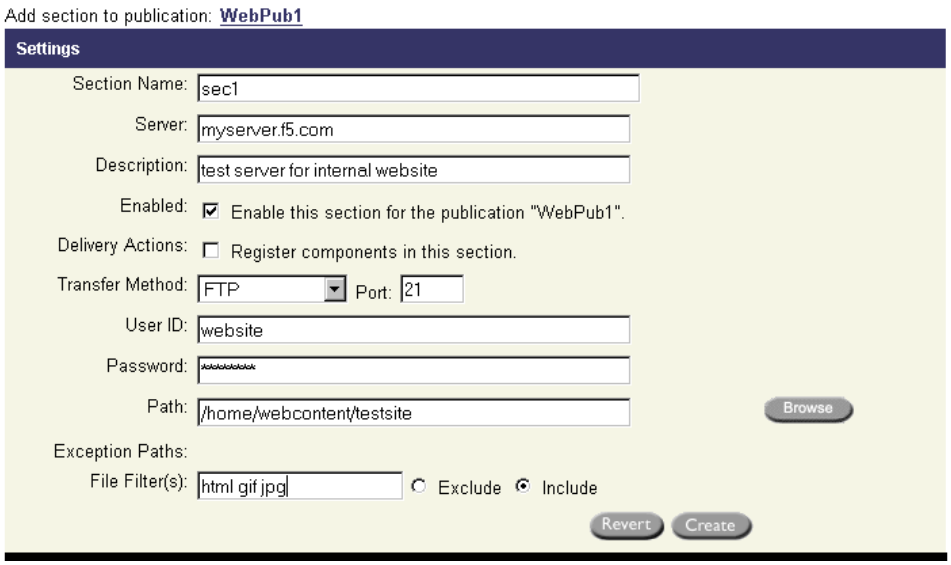

Figure 3.4 *The Create a New Section screen*

To create a section

1. On the navigation pane, click **List Publications**.
   The Publication List screen opens.
2. Click the publication for which you want to create a section.
   The Publication detail screen opens.
3. From the Publication detail page select the Sections tab for your new publication.
   The Publication Sections screen, Figure 3.3, opens.
4. Since there are no sections yet, you need to create one: click the **Create New Section** button.
   The Create a New Section screen opens.
5. Fill in the boxes to create your new section. For more information regarding this screen, click the Help button in the top right corner of the screen.

Once you create the section, if there are no errors, the Section Detail screen opens, showing the source detail for this section.

◆ Note

Section names must be unique for the section that you want to add. The name can only consist of alphanumeric characters, spaces, dashes (-), and underscores (_). You cannot change the name of a section once you have created it.

Testing your section

Once you create the section, we recommend you test the connection to confirm that the GLOBAL-SITE Controller can connect to the server that stores the content using your access method. Testing the connection verifies that the user ID, password, and path information are valid by connecting to the specified system.

◆ Note

*The **Test Connection** button is found in the top right corner of many screens in the GLOBAL-SITE Controller. We highly recommend you use this button to test connections whenever you create or change parameters that affect connectivity (for example: paths, transfer methods).*

To test your section connection

From the Section Detail screen, click the **Test Connection** button.
The Connection Test screen opens.

The Connection Test screen displays the test results in a table. The table lists the name for the source server, the path for the section, and the success or failure of the test connection. The left side of the table shows what is being tested, the right side shows the results of that test. The Connection Test screen content changes, depending on what is being tested.

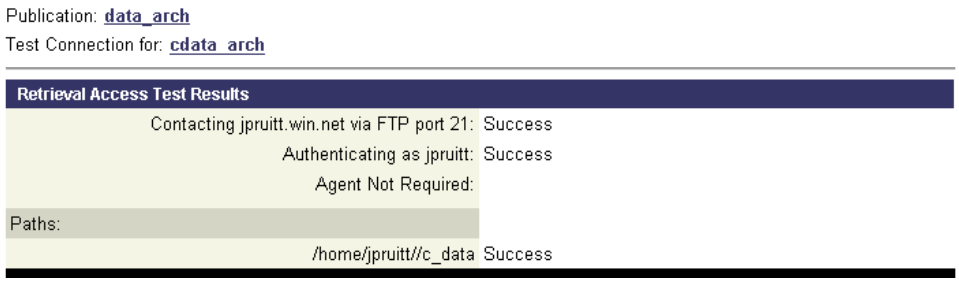
Publication: data_arch
Test Connection for: cdata_arch

Retrieval Access Test Results	
Contacting jpruitt.win.net via FTP port 21:	Success
Authenticating as jpruitt:	Success
Agent Not Required:	
Paths:	
/home/jpruitt//c_data	Success

***Figure 3.5** 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 *Updating the content of sections*, on page 3-16 of this chapter.

Modifying your section settings

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

To change your section parameters

1. From any Publication detail screen, click the Sections tab.
   The Publication Sections screen opens (see Figure 3.3).

   The Publication Sections screen lists all sections for the publication.
2. Click the name of the section you want to work on.
   The Section Detail screen for that section opens.
3. Make the necessary changes to any Section attributes.
   If you make a mistake or change your mind, click the **Revert** button to undo any changes you have made to this section.
4. Click **Save** to save your new settings.

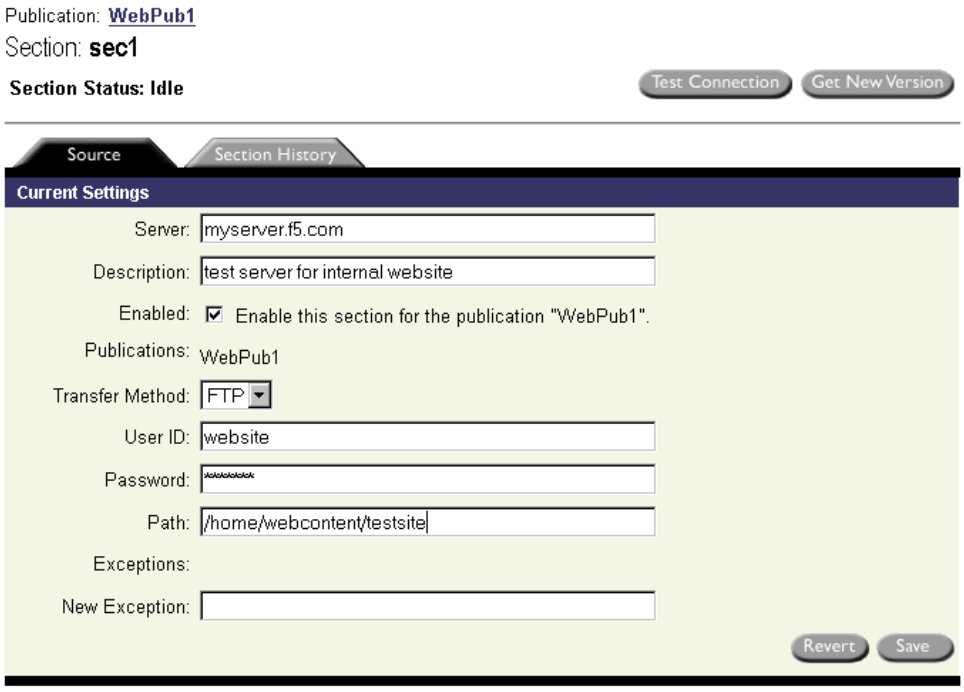

Figure 3.6 *The Section Detail screen, Source tab*

WARNING

*Disabling a section (clearing the **Enabled** check box) causes that section's contents to be removed from the next publication.*

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 each time you publish.

To update the content of a section

1. Start at the Publication Sections screen. You see your sections listed (see Figure 3.3).
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 opens (Figure 3.7), allowing you to monitor the progress of the process.

Publishing status is discussed in more detail in the section *Displaying status*, on page 3-19, later in this chapter.

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

Figure 3.7 *The Publish Progress Display screen*

Specifying the subscribers for section content

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. There are two types of subscribers for your publications: server subscribers, which can a server or a virtual server (also known as ***origin servers*** in EDGE-FX Cache) and cache subscribers. For more information about EDGE-FX Cache and cache subscribers, please see *Working with the EDGE-FX Cache and cache subscribers*, on page 4-10.

Defining server subscribers

The GLOBAL-SITE Controller delivers your publication to a server ***subscriber***. A server subscriber can be a server or a virtual server (see *Setting up a virtual server subscriber*, on page 4-9) if the BIG-IP is part of your network setup. (See *Creating a BIG-IP on the GLOBAL-SITE Controller*, on page 4-8.) For the files to be delivered successfully, you must provide the controller with:

- The specific path (server and directory) where content is delivered
- The transfer method by which the content is delivered

This is done when you add a subscriber, but can be changed at any time from the Subscriber Detail screen.

Example

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 you want to 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.

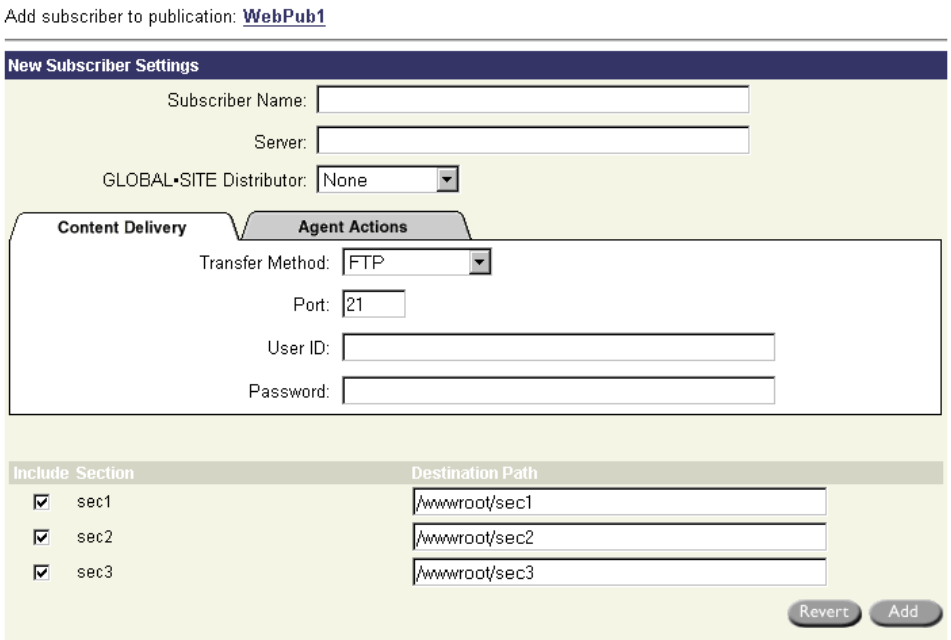

Figure 3.8 *The Add Subscriber screen*

To add a subscriber

1. From any of the Publication detail screens, click the Subscribers tab. The Publication Subscribers screen opens.
2. In the **Add subscriber type** box, leave **Server** (the default).
3. Click the **Add a Subscriber** button.
   The Add Subscriber screen opens.
4. Make the necessary changes as required. For more information regarding this screen, click the Help button in the top right corner of the screen.
5. Once the subscriber is created, you can click 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 any time 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 it up and add it to the primary GLOBAL-SITE Controller as a distributor. Secondary GLOBAL-SITE Controllers also need to be configured so that they are aware of, and can communicate with, the primary controller. To do this, you set up and add the primary controller as a distributor to the secondary controllers as well.

To add a distributor from the primary GLOBAL-SITE Controller

1. On the navigation panel, under Distribution, click **Add a Distributor**.
   The Add a Distributor screen opens.
2. Fill in the boxes as required. For more information regarding this screen, click the Help button in the top right corner of the screen.
3. Repeat these steps for each distributor you wish to add.

To add a distributor from the secondary GLOBAL-SITE Controller(s)

1. On the navigation panel, under Distribution, click **Add a Distributor**.
   The Add a Distributor screen opens.
2. Fill in the boxes as required using the primary controller's name and identifier. For more information regarding this process, click the Help button in the top right corner of the screen.
3. Repeat these steps on each secondary controller you added to the primary controller.

4. Once the distributor is created, you can click 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).

All the new distributors can be found on the GLOBAL-SITE Distribution List; new distributors are added at the bottom of the list.

Once you have both your primary and secondary distributors set up, you can easily specify a distributor for any subscriber that you add. See *To add a subscriber*, on page 3-14 in this chapter for details on adding a subscriber.

To add a distributor with a subscriber

1. Add the subscriber as you would normally. See *To add a subscriber*, on page 3-14 of this chapter for details.
2. In the **GLOBAL-SITE Distributor** box, select the distributor that you want to use for this subscriber.
3. Fill in the boxes as required, and click **Add**.
   This adds the new subscriber and opens the Subscriber Detail screen.
4. On the Subscriber Detail screen, click **Test Connection** to verify that the GLOBAL-SITE Controller can successfully log into the remote server.
   Running a connection test verifies that access can be established with the distributor using the user ID, password, and path information that you provided.

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:

- Updating the content of sections
- Delivering the content to subscribers

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

Updating the content of sections

To gather new content for an archived publication, 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. Each method opens the Publish Progress Display screen.

- From the Publication Sections screen, click the **Get New Version of Selected Sections** button.
- From the Section Detail screen, click the **Get New Version** button.
- From the Publication Editions screen, click **Create Edition**. On the Create a New Editions screen, select the **Get a new version** option from the version box and click the **Create Edition** button.
- From the Publication Editions screen, click the **Deliver Edition** button. On the Deliver screen, select **Create new Edition** from the **Edition** box, and click the **Deliver Edition** button. (This also publishes the newly retrieved files.)

In all cases for an archived publication, the GLOBAL-SITE Controller examines the directory tree specified by the combination of its staging server name, user ID, and path. The GLOBAL-SITE process ignores any exception directories on the path, includes or excludes the specified file extensions, and compares the content to the most current version in the GLOBAL-SITE Controller's versioned section. This process gathers: any new files; any files with different (not necessarily newer) modified timestamps; and any files that are a different size, and stores them in the GLOBAL-SITE Controller. The number of new or changed files is displayed on the Publish Progress Display 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.)

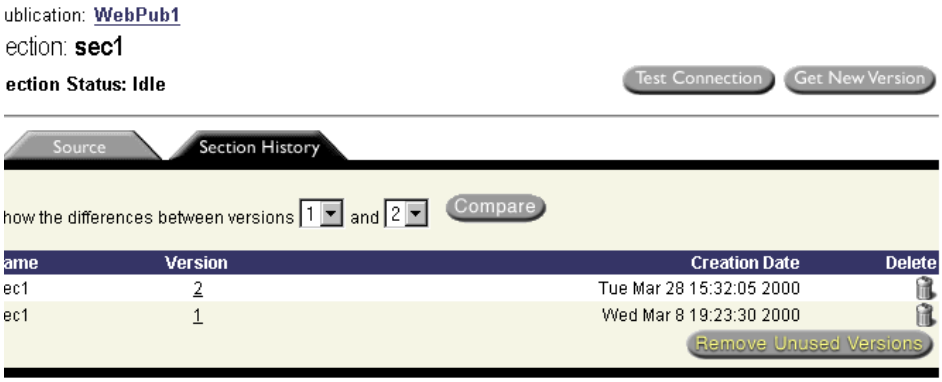

Figure 3.9 *The Section Detail screen, Section History tab*

In the example, we have already stored two versions. You can view the contents of any version by clicking the number in the Version column.

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 server 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.

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.

Delivering the content to subscribers

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

For details of the non-archived publication process and the enhanced scheduling options, refer to the online help or to *Scheduling the publishing process*, on page 4-2 and *Creating non-archived publications*, on page 4-15.

Creating the edition

The preparation phase involves defining exactly which version of each section in the publication is 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. On the navigation pane, click **List Publications**.
   The Publication List screen opens.
2. Click the name of the publication for which you want to create a section.
   The Publication Editions screen opens.
3. From the Publication Editions screen, click the **Create Edition** button.
   The Create a New Edition screen opens.

   The Create a New Edition screen displays all the sections in the publication, with a **Version** box that lists the each section's version. The latest existing version (with the highest version number) is displayed first in the **Version** box. You can click the arrow at the right of the box to view additional versions, create new versions, or remove the existing version. You accept the displayed version of all sections by clicking the **Create Edition** button.
4. Select the action you want for each section and click the **Create Edition** button. To accept the default version for all sections, click the **Create Edition** button without modifying the version boxes.

   The Publication Editions screen opens, with the newly created edition added to the list. The new edition in turn lists each section it contains and which version of that section it delivers.

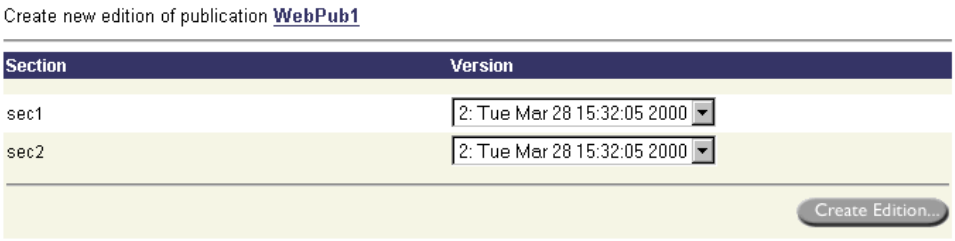

Figure 3.10 *The Create a New Edition screen*

The GLOBAL-SITE Controller automatically assigns both version and edition number, and 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 any Publication detail screen you can see **Status: Idle** in the top left side of the screen just above the tabs. When the publication status is **not** Idle, you can click the Status message at the top of the screen to display the Publish Progress Display screen.

You can also access the Publish Progress Display screen at any time from any Publication detail screen by clicking the **Show Details** button.

During the **Idle** state, the screen refreshes itself infrequently. Once the delivery process has started, 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, or you can use the **Stop** button on your browser to stop the automatic updates.

Delivering the edition

Once you have set up all your sections and decided what versions you want to deliver to your subscribers, the final step is actually delivering the publication.

To deliver the archived edition to the subscribers

1. From the Publication Editions screen, click the **Deliver Edition** button.
   The Deliver screen opens.
   The Deliver screen in our example shows that neither path has anything published to it yet. The **Edition** box defaults to the last edition selected in the Publish Editions screen. Consequently, you should double-check the **Edition** box before delivering, to ensure that you have the right edition selected. If necessary, you can change the edition displayed in the **Editions** box.
2. For each subscriber, select the sections you want to include in this edition by checking or clearing the check box in the Include column next to each listed section.

   Note: If there are two (or more) sections with the same path, and any of these sections are included for this delivery, then all of these sections are delivered to the subscriber(s) even if you cleared the **Include** check box.
3. From the **Start Delivery** list, select **Now** or select the time you want to start delivering the edition.
4. From the **Activate New Content** list, select when you want the content to be available to end users.
5. From the **Publish Option** list, select:
   - **Normal** to publish as usual.
   - **Verify Without Delivery** to check the delivered files against the GLOBAL-SITE database.
   - **Republish All Content** to republish all files regardless of what the subscriber has.
6. Click the **Deliver Edition** button.
   The Publish Progress Display screen opens.

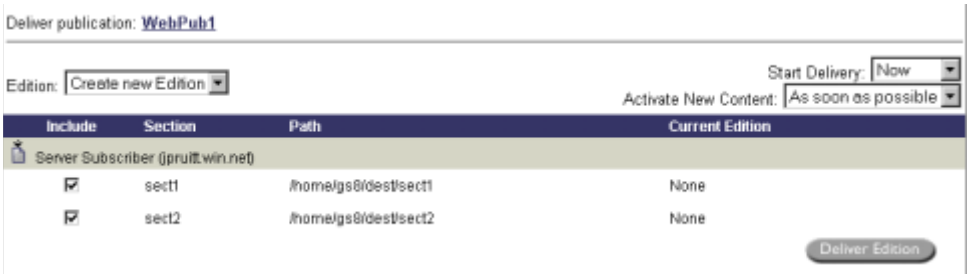

Figure 3.11 *The Deliver screen*

To deliver a non-archived publication to the subscribers

1. From the Publication Sections screen, click the **Deliver** button.
   The Deliver screen opens.
2. In the Action column, select the action to take for each section. Options are **Deliver**, **Ignore**, and **Remove**. Repeat this for each subscriber/each section listed on the screen.

   Note: *If you have two (or more) sections for the same subscriber with the same path, you must select the same Action for all of them. You get an error message if two sections with the same path have different actions.*
3. From the **Start Delivery** list, select **Now** or select the time you want to start delivering the edition.
4. From the **Activate New Content** list, select when you want the content to be available to end users.
5. From the **Publish Option** list, select:
   - **Normal** to publish as usual.
   - **Verify Without Delivery** to check the delivered files against the GLOBAL-SITE database.
   - **Republish All Content** to republish all files regardless of what the subscriber has.
6. Click the **Deliver** button.
   The Publish Progress Display screen opens.

Figure 3.12 *Publication Progress Display, status:* ***Updating Distributors***

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 3.12 shows details from an early phase in the process, and the details are about the sections. For details on all publishing phases and status messages, see *Interpreting phases in the status line*, on page 5-17 and *Using the Phases of Publication table*, on page 5-18 of this guide.

Using the Cancel button

Sometimes either the section update process (getting new content) or publication delivery process needs to be stopped before it is 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 is crucial.

The behavior of the **Cancel** button depends on how you have set up Error Handling on the New Publication or Publication Options page. If you set up the publication to **Pause publishing, log the error and wait for user input**, then clicking **Cancel** pauses the process at whatever point it is, and you can restart the process at that point.

If the publication is not set to pause publishing, clicking **Cancel** stops all processing activity. If the publication has not yet reached the **Copying Updates** phase, this resets the publication to **Idle**.

If the publishing cycle is stopped without any error indication and without having been manually stopped, clicking the **Continue** button should complete the publication cycle for any subscribers that have already had all section files copied to them by.

You can also try delivering again for the subscriber sections that failed, by clicking the **Retry** button. You can stop the publishing cycle with the **Reset** button, which restores the previously published edition 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 5.1, on page 5-19, lists the various activities or transition states, gives information about what should be happening, and includes suggestions on how to proceed if it is stalled and when to start over.

4

Increasing Control over Publishing

- Using enhanced publishing features
- Scheduling the publishing process
- Working with the BIG-IP and virtual servers
- Working with the EDGE-FX Cache and cache subscribers
- Creating non-archived publications
- Understanding file transfer methods
- Controlling servers with the GLOBAL-SITE agent
- Installing and configuring the GLOBAL-SITE agent and scripts
- Transferring files using WebDAV and WebDAV-SSL
- Section updates using FTP-Push
- Defining exceptions to sections

Using enhanced publishing features

Once you have read through instructions on how to set up a basic archived publication, you are 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 Controller provides several features that help you administer and manage your various publications. Some features you can use to augment basic publishing include:

- **Scheduling the publishing process**
  Scheduling the delivery to 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 units and virtual servers**
  The integration of the GLOBAL-SITE Controller with the BIG-IP 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. This reduces the risk of exposing a possibly inconsistent mix of old and new content to web users during the update process.
- **Working with the EDGE-FX Cache and cache subscribers**
  Using the GLOBAL-SITE Controller as the content management appliance for the EDGE-FX Cache eliminates the cache's polling process and keeps the cache informed of new content on the origin servers. This restores network bandwidth and enables the caches to have the most current content as soon as possible.
- **Creating non-archived publications**
  Non-archived publications give 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 faster delivery of content to subscribers and requires less storage capacity.
- **Transferring files using a method that meets your needs**
  The GLOBAL-SITE Controller gives you four different methods to transfer files: FTP, FTP-Push, WebDAV, and WebDAV-SSL. Each method has its pros and cons. For details, please see *Understanding file transfer methods*, on page 4-18 of this chapter.

- **Controlling servers with the GLOBAL-SITE agent**
  The GLOBAL-SITE agent enables limited remote administration of target servers. Using the agent gives you the ability to stop web services, register components, and reboot the server. The agent can also provide safer communication of passwords and content than FTP does, if you use its WebDAV capabilities.
- **Using 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. In addition, extensions allows you to either exclude or include only those files with a specified file extension.

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 content. The 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 you activate the new content at the subscriber for the customer to see.
- You can trigger delivery whenever there is new content using FTP-Push. See *Section updates using FTP-Push*, on page 4-30.

To understand the full extent of the GLOBAL-SITE Controller scheduling features, you should understand a bit about how publications are delivered to 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.

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 content updates (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 publish automatically 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 update intervals, if an update takes 18 minutes, the scheduler misses 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 starts again without delay in response to the immediately preceding scheduled interval (the one missed at the 15-minute mark).
- After you 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.
- You can allow content changes to trigger publishing by using FTP-Push as your transfer method. For more information, see *Section updates using FTP-Push*, on page 4-30 of this chapter.

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.

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**.

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

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 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 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 published manually, you can determine how the new content is 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.

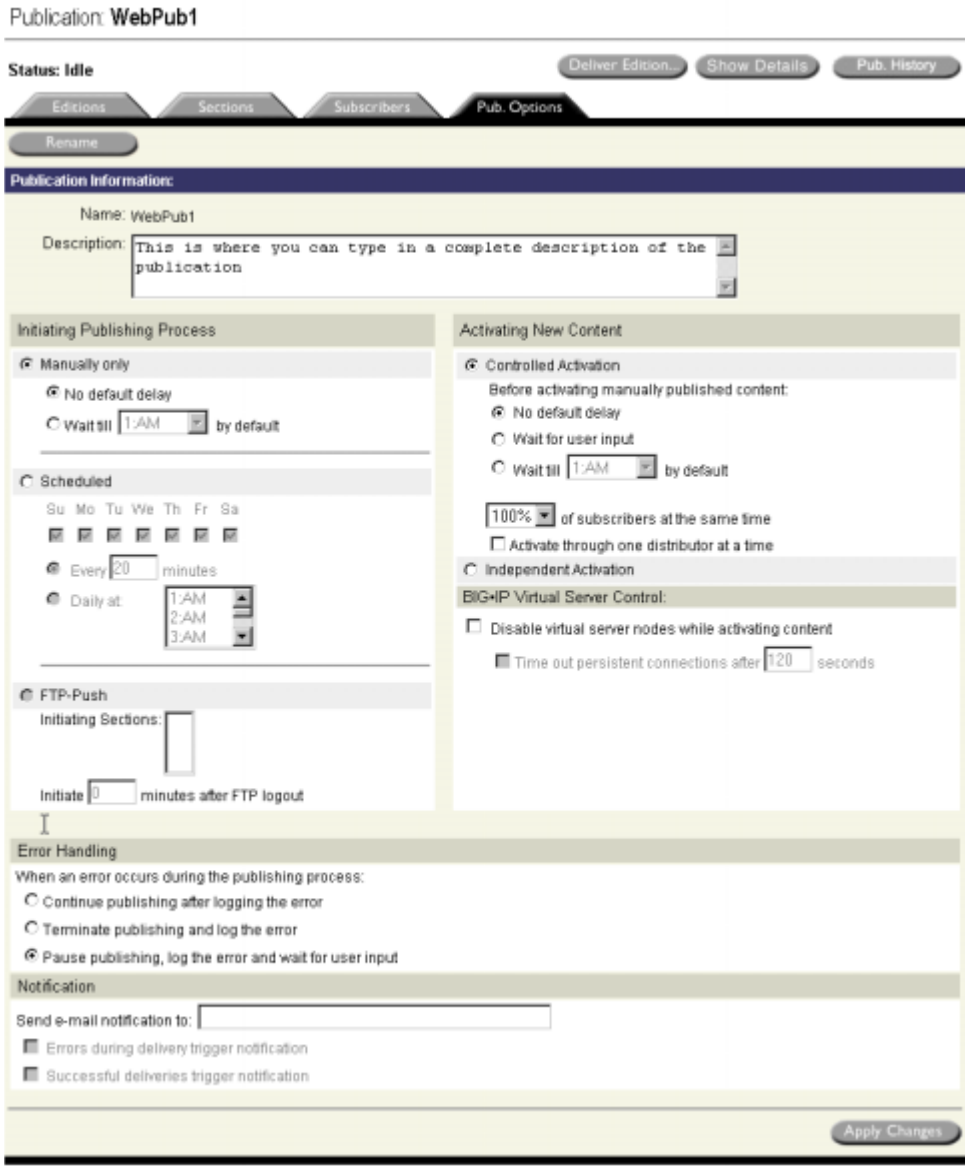

Figure 4.1 *Use the Publication Options screen to activate new content*

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 the content of the other sections have 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 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 the GLOBAL-SITE Controller you control how subscribers are activated. These control options are 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 content 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.

When you are accessing some of the subscribers with 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 and virtual servers

A powerful feature of the GLOBAL-SITE Controller is the integration of publishing with the BIG-IP. (The GLOBAL-SITE Controller can communicate with the BIG-IP, version 3.1 and above.) Typically, the BIG-IP is 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, 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, 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 reduces the risk of exposing a possibly inconsistent mix of old and new content to web users during the activation process.

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

Use the GLOBAL-SITE Controller browser interface to:

- Set up a BIG-IP as a virtual server
- Remove the unused BIG-IP units
- Change settings for the current BIG-IP units
- Mark the BIG-IP nodes as down during publishing

Guidelines for BIG-IP—GLOBAL-SITE Controller cooperation

Before you create a virtual server subscriber on the GLOBAL-SITE Controller, be sure to follow these guidelines:

For the BIG-IP version 4.0 and later, the iControl portal runs by default on the BIG-IP. It is therefore ready to accept connections from the GLOBAL-SITE Controller. Please see the iControl SDK documentation for more information on verifying and controlling the portal process on the GLOBAL-SITE Controller.

For older versions of the BIG-IP, 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 distributes connections). In some releases of the BIG-IP, that interface must have **admin** access.

- Older versions of the BIG-IP must have the **bigorb-server** installed for the GLOBAL-SITE Controller to talk to. (Your vendor may install it for you.) Newer versions of the BIG-IP provide an iControl interface, which uses an administrator account and password.
- The GLOBAL-SITE Controller must have FTP or WebDAV access to the virtual server nodes using the same node IP address that the BIG-IP has configured for the virtual server that the GLOBAL-SITE Controller in turn uses to create a virtual server subscriber.
- You need to define a virtual server on the BIG-IP for port 2792 (for GLOBAL-SITE Controller to GLOBAL-SITE Controller communication, using the F5-GLOBAL-SITE protocol) if you are using it as a distributor for another GLOBAL-SITE Controller.
- If you use the GLOBAL-SITE Controller (that you are using with the BIG-IP) as a publisher, you must configure a virtual server on the BIG-IP to the GLOBAL-SITE Controller for a port for access to the web browser. The default is port 443, but you can configure this in the First-Time Boot utility. You may also need a secure NAT (SNAT) for the GLOBAL-SITE Controller to get FTP access to systems such as section sources by way of the BIG-IP.
- Once you have completed the preceding tasks, you need to inform the GLOBAL-SITE Controller about the BIG-IP using the Add a BIG-IP screen.

Creating a BIG-IP on the GLOBAL-SITE Controller

To add a new BIG-IP 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 only if the BIG-IP is up and available locally or through a distributor.

To add a BIG-IP

1. Type the name of the BIG-IP that you want to add. This is the internal IP address, or alias of the BIG-IP.
2. Type a description for the BIG-IP. This identifier may make it easier for you to recognize each BIG-IP unit. (This step is optional.)
3. Select a BIG-IP distributor from the list of available BIG-IP units.
4. For the BIG-IP version 4.0 or later, type the user name and password used for the CORBA portal on the BIG-IP. For earlier versions of the BIG-IP, you can skip this step.

5. Click the **Add** button to add this BIG-IP distributor.
   The BIG-IP List screen displays.

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

Setting up a virtual server subscriber

If you have one or more BIG-IP units 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 units, your option is to set up a normal server subscriber. To set up any server, start with the Publication Subscriber screen.

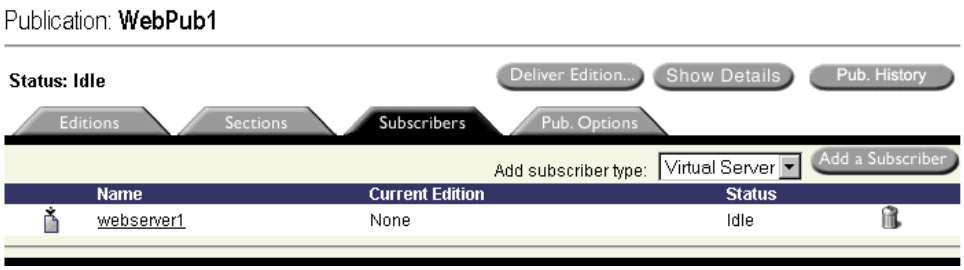

Figure 4.2 *Selecting a virtual server subscriber on the Publication Subscriber screen*

To add a new subscriber

1. From the **Add subscriber type** box, select **Virtual Server**.
2. Click the **Add a Subscriber** button.
   The 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 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 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 is instructed to re-enable web access. Use the BIG-IP Virtual Server Control options on the Publication Options screen for these settings.

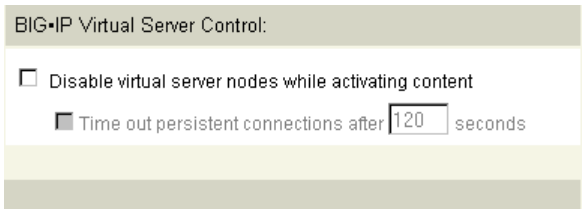

Figure 4.3 *BIG-IP Virtual Server Controls on the Publication Options screen*

If you set up a virtual server subscriber, you can specify that 50% of the subscribers are to be activated together and that access to them be disabled during the activation phase. Using this method, no single virtual server node is enabled with a mix of old and new content at the same time. Because they are accessed by a virtual server address, any single subscriber may show old pages while the first files are being updated, but then has 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 load balances the user from a system with old content to one with new content.

If you choose the **Independent Activation** mode, the BIG-IP 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.

Working with the EDGE-FX Cache and cache subscribers

As an administrator performing routine publication management, you can populate or expire cached content on the EDGE-FX Cache from the GLOBAL-SITE Controller. When the GLOBAL-SITE Controller completes publishing content to its subscribers (which the EDGE-FX Cache calls ***origin servers***), it can immediately notify cache subscribers to update stored content from those servers. Integration between these products:

- **Improves content delivery**
  Newly-published content can be automatically cached before it is requested so content gets to the end-user faster.

- **Consumes less bandwidth**
  The cache is told when content has expired, eliminating the need for constant freshness checks.

You can also create cache subscribers from the GLOBAL-SITE Controller. Cache subscribers receive information on the status of new content from the controller. This eliminates the cache's need to poll the origin servers to check for new content, and consumes less bandwidth. The caches can then get new content prior to any request for that content, or can remove outdated content as soon as it expires.

In order for the GLOBAL-SITE Controller to manage your cache, you must make sure that you have configured the EDGE-FX Cache to external validation mode. Please refer to the EDGE-FX Cache documentation for this information.

Supported URL schemes

The GLOBAL-SITE Controller, and the EDGE-FX Cache manage three different URL schemes: **http://**, **mms://**, and **ftp://**. The exact behavior of the Expire and Populate modes is slightly different for each one.

- **http://**
  The Expire mode purges any new, changed, or removed file from the EDGE-FX Cache. A failure of the purge command to the EDGE-FX Cache is considered a delivery failure by the GLOBAL-SITE Controller. The reason is that until the purge command is known to succeed, an EDGE-FX Cache could be retaining stale data. The GLOBAL-SITE Controller reissues the same purge command until it succeeds, depending on how its normal error handling options have been configured for the publication. The Populate mode causes an immediate **get** for the specified URL. A failure to actually retrieve the URL (a 404 error) is not considered a publication failure, because in that event, the EDGE-FX Cache does not serve up stale content either. Such an error probably represents an overall configuration failure, because it means that the EDGE-FX Cache was instructed to get data not available from the origin servers. Such a scenario is very unlikely if the GLOBAL-SITE Controller is managing the content on multiple origin servers behind a BIG-IP and a 3-DNS Controller, but not all sites are so well architected.
- **mms://**
  This is the scheme used for Windows media files. The Expire mode simply purges data from the media cache, without pre-positioning. The Populate mode actually pre-positions the data to the EDGE-FX Cache. The data is delivered directly by way of iControl without getting it from the origin server. This avoids the need for the EDGE-FX Cache to ever get (potentially) large media files from the origin server at all. The GLOBAL-SITE Controller retrieves the data once from the origin server, and then directly delivers it to the specified set of caches. You can configure multiple GLOBAL-SITE Controllers, as you would with any subscriber, in a primary-distributor relationship to scale up delivery capability to large numbers of subscribers.

- **ftp://**
  The Expire mode behaves the same as the others, simply purging any cached data when new, changed, or removed files are delivered. The Populate mode gets the file from the specified (**ftp://server.com/directory**) using anonymous access only. If anonymous access is not configured for the FTP server, this process fails. The GLOBAL-SITE Controller treats such failures as a delivery failure, and retries delivery in subsequent delivery attempts. Only EDGE-FX Caches that are configured in forward mode have the ability to either expire or populate **ftp://** requests. If you attempt to use this scheme for caches in reverse mode, the process results in delivery failures.

Guidelines for EDGE-FX Cache and GLOBAL-SITE Controller cooperation

In order for you to effectively use the GLOBAL-SITE Controller and EDGE-FX Cache arrangement, you must ensure that both the GLOBAL-SITE Controller and the EDGE-FX Cache are configured to work with each other. In addition, you must be aware of certain details.

- You configure The EDGE-FX Cache to allow the GLOBAL-SITE Controller to manage it. See the EDGE-FX Cache documentation for information on how to do this.
- You must have permission for each cache operation. You set permissions on the EDGE-FX Cache machine. See the EDGE-FX Cache documentation for information on how to do this.
- Cacheable dynamic content cannot be expired, however most dynamic content is not cacheable.
- The GLOBAL-SITE Controller (or other external agent) must manage(expire/populate) all content on the cache for a domain that is in external validation mode. If non-managed content for that domain gets into the cache, it never expires.

Managing a cache subscriber

The GLOBAL-SITE Controller supports the EDGE-FX Cache subscribers using the iControl system. A cache subscriber is generally different from other subscribers in the following ways:

- A cache subscriber exclusively uses the iControl system for data management, unlike other subscribers which may use either FTP or WebDAV for data management. The username and password specified for the cache subscriber must be a valid iControl account on the EDGE-FX Cache system. See the EDGE-FX Cache documentation for details on configuring such accounts.
- Cache subscribers are handled only after all other subscribers have had their content delivered and activated. This is because the EDGE-FX Cache is presumed to be configured to get needed data from an origin server that is represented by one or more of the non-cache subscribers to

a publication. This allows a single publication to deliver updates to origin servers, and update caches with fresh content, in a single delivery cycle.

- A cache subscriber does not specify a file system directory as a destination for file delivery. Instead, a URL base is specified for each section. This URL base specifies the scheme, domain, and a directory path. Each file being updated in the section is appended to this root to form a complete URL for presentation to the cache. This, in turn, presumes a one-to-one mapping between files in the section and the URL to which the origin server responds for that data.
- A cache subscriber supports two mutually exclusive modes, Expire and Populate, for managing cached content.
  - The Expire mode is the closest approximation to typical cache behavior. In the Expire mode, a new or changed file causes the GLOBAL-SITE Controller to send an iControl command to purge the synthesized URL (as explained previously) from the EDGE-FX Cache. This, in turn, forces the next access attempt by a cache client to get the URL from the origin server, even though freshness heuristics may have otherwise determined that previously-cached content could be served. This ensures that the cache never serves stale content.
  - The Populate mode takes the idea further by instructing the EDGE-FX Cache not only to expire any cached copy immediately, but also to immediately get a new copy. This allows the GLOBAL-SITE Controller to effectively pre-position new content into an EDGE-FX Cache so that fresh, cached, content is served from the EDGE-FX Cache on the very next client request, without the need to proxy that initial request back to the origin server. Combine this with the GLOBAL-SITE Controller's scheduled content delivery capabilities, and you can update a large set of EDGE-FX Caches with new content overnight, or whenever network load is light, to minimize surges in demand on origin servers when newly updated content is accessed on a daily cycle.

When setting up a cache subscriber, you have the option of setting it to expire content or populate content. Once you have entered the URL root and the section path, you need to choose between expire and populate:

- Choose **Expire** if you want the cache to delete the content.
- Choose **Populate** if you want the cache to immediately get new content (before a request is made for it).

In order for the GLOBAL-SITE Controller to manage the cache, you must have the correct permissions set on the cache, and you must have correctly configured the EDGE-FX Cache. For details, please see the EDGE-FX Cache documentation and *Guidelines for EDGE-FX Cache and GLOBAL-SITE Controller cooperation*, on page 4-12 of this chapter.

Creating a cache subscriber

There are three ways to create a cache subscriber in the GLOBAL-SITE Controller. All are effective, however some are easier to do. Depending on the situation, you can:

- Create a cache subscriber from a server subscriber
- Create a brand new cache subscriber
- Clone an existing cache subscriber

Creating a new cache subscriber

Use this procedure if you are creating a cache subscriber that is not based on a server or virtual subscriber. When you create a new cache subscriber, you have to fill in all the fields and paths manually. You must make sure you enter paths and directories exactly, or the GLOBAL-SITE Controller cannot contact the cache.

To create a new cache subscriber

1. From any of the Publication detail screens, click the Subscribers tab. The Publication Subscribers screen opens.
2. In the **Add subscriber type** box, select **Cache**.
3. Click the **Add a Subscriber** button.
   The Add a Cache Subscriber screen opens.
4. Fill in the information as required. Paths must be URLs.
5. Once the subscriber is created, you can click the **Test Connection** button to verify that the GLOBAL-SITE Controller can successfully log into the specified server using the user account and manage the specified domain. This is a good thing to do at this point and any time you make changes to the subscriber definition.

Cloning a cache subscriber

Before you can clone a cache subscriber, you must already have a cache subscriber created. The list of subscribers is found on the Subscribers tab. You cannot have an exact duplicate of any subscriber, therefore when you clone a cache subscriber, the controller clears the **Subscriber** box and the **Cache** box, and you need to provide new information for these boxes. Change any other boxes on the screen as necessary.

To clone a cache subscriber

1. From the Cache Subscriber Detail screen, click the **Clone** button at the bottom of the screen.
   A new Cache Subscriber Detail screen opens. All the boxes have data except for **Subscriber** and **Cache**.

2. Enter a new name in the **Subscriber** box for this cache subscriber. You can use only alphanumeric characters, spaces, dashes (-), and underscores (_).
3. Enter a new host name in the **Cache** box.
4. Change the data in the other boxes as necessary.
5. Click **Update** to save your new subscriber.

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 the 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 and available to the viewer on the subscriber.

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 then 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 copying 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 the content the subscriber has. It does this by comparing a list of files, sizes, and dates (that are 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 than100 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 improves 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 been removed or becomes empty, 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 the controller updates on the next publishing cycle. The reason 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 *Scheduling the publishing process*, on page 4-2 of this guide.

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. The New Publication screen is the same for archived and non-archived publications, however to make a publication non-archived you must clear the **Archive publication editions and section versions** check box at the top of the screen.

There are a few 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, because non-archived publications have no editions and no versions, there are other small differences on some screens. These differences are noted in the online help pages for each screen.

One important difference is the delivery screen. Non-archived publications allow the following actions: **Deliver**, **Ignore**, or **Remove** for each section on a subscriber. You must choose the same action for all sections that share the same subscriber's path.

◆ Note

Once you have created a publication, you cannot change its archived or non-archived state.

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. At the top of the screen, clear the **Archive publication editions and section versions** check box.
3. Complete the settings as required and click the **Add** button.
   The Publication Sections screen opens.

   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.

Note

The Editions tab does not display, as there are no editions of a 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.

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 sections and the online help pages point out any differences.

Understanding file transfer methods

The GLOBAL-SITE Controller offers four transfer methods, FTP, FTP-Push, WebDAV, and WebDAV-SSL. Each method has its advantages and disadvantages. This section of the guide helps you choose the right transfer method for your needs. Following the transfer method sections is a section detailing how to choose a port based on the transfer method you chose.

WARNING

Changing the transfer method of an existing section can have serious consequences if the services you are talking to have different root directories.
For example, a problem occurs if you changed transfer methods from FTP to WebDAV and your FTP server had a different root directory than your HTTP server.
In particular, changing an existing section to FTP-Push requires copying the section's files to the GLOBAL-SITE Controller. See ***Switching sections in an existing publication to use FTP-Push transfer method****, on page 5-8 of this guide.*

FTP

The GLOBAL-SITE Controller functions as an FTP client and initiates contact with FTP servers where content is picked up. FTP is familiar to most users, is readily accessible, and available on all platforms. However, FTP does have its limitations, which include:

- Opening a new connection for most commands, which increases overhead, and slows the operating system
- Transferring passwords and files with no encryption or error checking
- Requiring manual intervention on each subscriber to run CGI, and other normally executable files, because many FTP servers cannot set the execute permission on transferred files

FTP-Push

There are two main differences between FTP and FTP-Push:

- With FTP-Push, the machine associated with a section contacts the GLOBAL-SITE Controller when content has changed and pushes content to the GLOBAL-SITE Controller (the controller is the FTP server in this case). The content is stored on the GLOBAL-SITE Controller. (For more information, see *To add content to an FTP-Push section*, on page 4-32 of this chapter.)
- Once content is updated on the GLOBAL-SITE Controller, an FTP-Push section can be set to trigger the GLOBAL-SITE Controller to push the new content out to the subscribers. If you set up your FTP-Push this way, you do not need any manual intervention to deliver the publication. FTP-Push has the same inherent disadvantages as standard FTP. (For more information, see *Section updates using FTP-Push*, on page 4-30 of this chapter.)

WebDAV

To use WebDAV (World Wide Web Distributed Authoring and Versioning) to transfer files, you must either have a web server that supports it (IIS5.x or Apache HTTP Server with **mod_dav**), or use the GLOBAL-SITE agent. If you install the GLOBAL-SITE agent, you can use WebDAV to transfer files even if your web server does not support it. (For more information please see RFC 2518.)

◆ Note

When you use WebDAV to transfer files to an IIS server, some file extensions cannot be transferred. This may have to do with the IIS server settings. Please see that product's documentation for more information.

WebDAV reuses connections, resulting in efficient use of the operating system and firewalls. WebDAV with digest authentication encrypts passwords and performs file checking, making it easier to discover files that have changed in transit. (For more information on digest authentication, please see RFC 2617.) WebDAV can maintain the execute permission on files.

To use digest authentication with a Windows 2000 server, the server must be configured as a domain controller. For more information, see *Transferring files using WebDAV and WebDAV-SSL*, on page 4-29 of this chapter.

For file transfers for either sections or subscribers, on Windows NT/2000 machines, you can use WebDAV if you install the GLOBAL-SITE agent. For more information, see *Controlling servers with the GLOBAL-SITE agent*, on page 4-20 of this chapter. For subscribers, if you are going to use any agent actions you must install the GLOBAL-SITE agent, and you must choose WebDAV for your file transfer method.

WebDAV-SSL

WebDAV-SSL adds security to WebDAV file transfers by encrypting both passwords and files using secure socket layer (SSL) features. However, you trade increased security for file transfer performance. For more information, see *Transferring files using WebDAV and WebDAV-SSL*, on page 4-29 of this chapter.

Choosing a port for the file transfer method

You must use the correct port number for successful file transfer. Each transfer method has a default port as follows:

- FTP = Port 21
- FTP-Push = Port 21
- WebDAV = Port 80
- WebDAV-SSL = Port 443

In general you should use the default ports. However, if you are using the GLOBAL-SITE agent, use the port number that you entered when you installed the agent (the default is 50,000).

Controlling servers with the GLOBAL-SITE agent

When you install the GLOBAL-SITE agent on a server, it gives you the ability perform certain necessary tasks remotely from the GLOBAL-SITE Controller. The following are tasks you can perform remotely when you install the GLOBAL-SITE agent:

- **Stop HTTP Server before activating content**
  Allows you to stop the web server on the subscriber when you deliver new content or components. Stopping and restarting the server allows the server to recognize new files. The ability to stop the server also allows you to get around locked files, as stopping the server unlocks the files.
- **Initiate scripts**
  Allows you to run a user-defined script, at a specified time during the publishing process. Scripts reside in a particular directory. When it is time to run the scripts, the agent runs all executables that it finds in the specified directory. Please see *Installing and configuring the GLOBAL-SITE agent and scripts* on page 4-27 for more information about user-defined scripts.
- **Reboot machine after activating content**
  Allows you to reboot the subscriber (server) so it recognizes newly registered components
- **Register components**
  Allows you to register components of your web site that need to register with the operating system or the web server.

The GLOBAL-SITE agent also allows you to use WebDAV as your transport mechanism even if your web server does not support WebDAV. WebDAV provides safer communication of passwords and content than standard FTP.

You must install the GLOBAL-SITE agent on each machine that you want to use it on.

Once you have installed the agent on a section or subscriber machine, you need to modify the settings for the section or the subscriber to enable its use. The Connection Test results have also been expanded to include information on the agent, where appropriate (see *To test an agent for a subscriber or section*, on page 4-26).

Installing the GLOBAL-SITE agent on the subscriber

The Windows agent installer is an InstallShield® application. See *To install an agent on a Windows server*, on page 4-22. On Solaris and Linux you get the GLOBAL-SITE agent in a tar file that has the image and the installation program with which to install and configure the agent. You run the installation program, **gsAgentInstall** from the command line as the **root** user. The installation program also looks for the installation image, **GlobalSite.tar**, in its current directory. See *To install an agent on a Solaris or Linux server*, on page 4-23.

The Solaris and Win32 installation prompts are almost identical. The Solaris installation prompts you for the command to stop the web server and the command to start the web server.

If you enable scripting support, the installation creates a virtual directory, **UserScripts**, under the document root directory. The virtual directory points to the **UserScripts** directory under the agent installation root directory. See *Installing the scripts on the subscriber*, on page 4-28 for more information on scripts.

After you install the agent, you can modify the configuration by running the configuration program, **installationRoot/bin/gsAgentConfig**.

To install an agent on a Windows server

Run the following installation on each Windows server that you want to install the GLOBAL-SITE agent on. You do not install this on the GLOBAL-SITE Controller.

Following are step-by-step instructions for running **global-siteagent.exe**.

1. Log on to the server using an account in the Administrators group, and exit from all other Windows applications.
2. Connect to the setup program on the GLOBAL-SITE Controller. From a browser on the server, use the name of the GLOBAL-SITE Controller in the following URL:

   `https://<GLOBAL-SITE-CONTROLLER>`
3. On the opening screen, click **Agents**.
   The Agent Download screen opens.
4. Click **Windows**.
5. Select **Run this program from its current location**, and confirm it in the next popup screen.
6. Read the Welcome screen and click the **Next** button.
7. Read the license and click **Yes** if you agree to the terms.
8. Read the acknowledgement screen and click **Next**.
9. Accept the default destination folder or browse to select an alternative, and then click **Next**. The default destination folder is

   `c:\Program Files\F5 Networks\Global-Site Agent`
10. Press Enter to accept the default of a non-SSL-enabled agent, or type **y** and press Enter to enable SSL on the agent.

    Note: If you enable SSL for your agent, every section or subscriber that uses this agent must use the WebDAV-SSL transfer method (see *WebDAV-SSL*, on page 4-20). If you disable SSL for your agent, every section or subscriber that uses this agent must use the WebDAV transfer method (see *WebDAV*, on page 4-19).
11. Type the port the agent monitors. The range is 1-65535 and the default is 50000. You need to use the same port number for every section or subscriber that gets this agent, so please make a note of it.
12. Type the directory to which files are delivered. This usually matches your IIS DocumentRoot directory. The default is:

```
c:\inetpub\wwwroot
```

13. Select the hosts that can access the agent (it is safer to limit access to the IP address of the GLOBAL-SITE Controller):
    - **all**
      All hosts with a password can access the agent (the default.
    - **domain.com**
      Only hosts in **domain.com** can access the agent.
    - **IP address**
      Only one host can access the agent (the hostname does not work).
14. Select one or more of the following options to restrict the agent:
    - Restart the web server and reboot the computer
    - Register and unregister components
    - Deliver and activate files on this computer
15. Add users and passwords in the Agent Password Manager popup screen.
16. When you finish typing names and passwords, click **Exit**.

To install an agent on a Solaris or Linux server

Run the following installation on each Linux or Solaris server that you want to install the GLOBAL-SITE agent on. You do not install this on the GLOBAL-SITE Controller.

Following are step by step instructions for running **gsAgentInstall**.

1. From the GLOBAL-SITE Controller home page, (https://<GLOBAL-SITE Controller>), click **Agents**.
   The Agent Download screen opens.
2. Click the appropriate agent for your server: **SolarisAgent2.6.tar** or **LinuxAgent.tar**.
3. Copy the agent tar file to your server.
4. Log in as **root** to the directory on your server where the agent tar file has been copied to.
5. Run the following commands:

   ```
   tar -xvf LinuxAgent.tar
   ./gsAgentInstall
   ```
6. Press Enter to accept the default destination directory, or enter a different directory in which to install the agent. The default destination directory is **/var/f5**.
7. Press Enter to accept the default of a non-SSL-enabled agent, or type **y** and press Enter to enable SSL on the agent.

   Note: If you enable SSL for your agent, every section or subscriber

that uses this agent must use the WebDAV-SSL transfer method. If you disable SSL for your agent, every section or subscriber that uses this agent must use the WebDAV transfer method.

8. Type the port the agent monitors, or press Enter to accept the default of **50000**. The range is **1-65535**. You need to use this same port number for every section or subscriber that uses this agent, so please make a note of it.
9. Type the base directory to which files are delivered (for subscribers) or read from (for sections). Press Enter to accept the default of **/var/www**.
10. To be more secure, limit access to the IP address of the GLOBAL-SITE Controller that uses this agent. Type the hosts that can access the agent or press Enter to accept the default of **all**.
    - **all** - All hosts with a password can access the agent.
    - **domain.com** - Only hosts in **domain.com** can access the agent.
    - **IP address** - Only one host can access the agent (the host name does not work).
11. The following four choices restrict the agent's access to the server. Choose **y** to permit the agent (the default) or **n** to prevent it on the following actions:
    - Reboot the server
    - Restart the Web server
      If you chose to allow the agent to restart the web server, enter the commands to stop and start the web server on the following two prompts:

      ```
      /etc/rc.d/init.d/httpd start
      /etc/rc.d/init.d/httpd stop
      ```

    - Deliver and activate files on this machine
      For subscriber(s), you must accept the default or the GLOBAL-SITE Controller can not deliver files to the subscriber.

      If you are installing the agent for sections only, then you choose **n**.
    - Run custom scripts
12. Your entries are displayed on the screen. (If they are not, press Enter to re-enter the correct values.) Type **y** and press Enter to commit your entries and continue with agent installation.
13. The controller now generates a self-signed certificate.
14. When prompted, enter a user name and password to use for the agent. This must match the user name and password that you use for your section or subscriber.
15. Press Enter to complete installation and start the agent immediately, or type **n** and press Enter to complete installation, but not start the agent.

To enable an agent for a subscriber

This procedure is done on the GLOBAL-SITE Controller from a browser anywhere in the network.

1. In the navigation pane, click **List Publications**.
   The Publication List screen opens.
2. Click the name of the publication you are working with.
   The Publication Editions screen opens (for archived publications) or the Publication Sections screen opens (for non-archived publications).
3. Click the Subscribers tab.
   The Publication Subscribers screen opens.
4. Click a subscriber in the Name column or create a new subscriber (see *To add a subscriber*, on page 3-14).
   The Subscriber Detail screen opens. This screen has two tabs, Content Delivery and Agent Actions.
5. On the Content Delivery tab, select the transfer method **WebDAV.**
   A default **Port** setting appears.
6. Type the same port setting you specified when you ran the **global-siteagent.exe** program.
7. Type a user ID. This should be the user ID you use for WebDAV access to this subscriber.
8. Type a password. The agent uses the user ID and password to authenticate file transfers.
9. Click the Agent Actions tab.
   If you selected WebDAV on the Content Delivery tab, then the box for **Use Content Delivery connection settings** is checked.
10. In the Actions area, check one or more of the following boxes to remotely administer the services the agent provides to each subscriber. Note that the agent restarts all web services when finished.
    - Check **Stop HTTP Server before activating content** to stop the web services on the subscriber, enabling the agent to receive files and register components.
    - Check **Reboot machine after activating content** to allow the agent to reboot the subscriber to make it recognize newly registered components.
    - Check **Initiate Scripts** and choose the phase in which to run the scripts. Scripts are user defined and must reside in the correct directory to work properly. Please see *Creating the scripts*, on page 4-28 for more information about scripts.
    - (For Windows agents only) Check **Register Components** to allow the agent to register newly copied components on the subscriber. If the publication contains sections that have

Register Components checked (on the Section Detail screen), then the **Path** field on the Subscriber Detail screen displays where components should be registered.

11. Click **Update** if you are making changes to an existing subscriber. Complete the settings on the screen if you are creating a new subscriber. Click the Help button for more information on this screen.

To enable an agent for a section

1. In the navigation pane, click **List Publications**.
   The Publication List screen opens.
2. Click the name of the publication you are working with.
   The Publication Editions screen opens (for archived publications) or the Publication Sections screen opens (for non-archived publications).
3. Click the Sections tab if you working with an archived publication.
   For non-archived publications you are already at this screen.
4. In the Name column, click the section name or create a new section (see *To create a section*, on page 3-9).
   The Section Detail screen opens.
5. Select the transfer method **WebDAV**.
   A default **Port** setting appears.
6. Type the same port setting you specified when installing the agent.
7. Click **Save**.
   The Section Detail screen refreshes to display your changes.

To test an agent for a subscriber or section

It is a good idea to check your connections once you have made changes to a section or a subscriber to make sure that everything is working well before you start sending files.

To do this, click the **Test Connection** button found on either the Section Detail or Subscriber Detail screen.
The Connection Test screen opens.

There are several possible test messages, but the most common are:

- The **Contacting Agent on <server> on port <port number>** line displays one of the following test messages:
  - Success
  - Error connecting to server
- The **Authenticating to Agent as <userid>** line displays one of the following test messages:
  - Success

- No agent present
- Stop and start Server (OK or DENIED)
- Reboot Server (OK or DENIED)
- Deliver Files (OK or DENIED)
- Register components (OK or DENIED)

Installing and configuring the GLOBAL-SITE agent and scripts

The GLOBAL-SITE Controller can run custom scripts, during specified phases and on specified publications, on subscribers running Microsoft Windows NT, Microsoft Windows 2000, Sun Solaris, and Red Hat Linux. This feature requires user configuration on both the subscriber and the GLOBAL-SITE Controller. The GLOBAL-SITE Controller configuration is documented in the online help for the Subscriber screen. This section describes how to configure the subscriber to match the GLOBAL-SITE agent script configuration.

Configuring the subscriber

You must install the GLOBAL-SITE agent on every subscriber that has been configured; you can use the steps described in the section *Installing the GLOBAL-SITE agent on the subscriber*, on page 4-21 of this chapter. You configure the agent to run scripts on the server. You must also copy, to each subscriber, the scripts that need to run. The following sections describe the phases at which you can run the scripts and provide an example of the directory structure that you must use.

In the example that follows, **<agent installation directory>** is the directory where the agent software gets installed, and not the directory where data files are accessed. Installing the agent creates the **<agent installation directory>/cgi-bin/UserScripts** directory, however you must create the **<publication name>/<publication phase >** directory before you copy scripts into it. For example, any scripts you want to run in the Before Copy phase must be copied into the following directory:

```
<agent installation directory>/cgi-bin/UserScripts/<publication name>/BeforeCopy
```

Following is the list of publication phases. The name of the directory you need to create for each phase is listed with the phase name.

- **Before Copy**
  The Before Copy publication phase runs the scripts before any files are copied to the subscriber. You can run any number of scripts during this phase. Name the directory **BeforeCopy**.
- **Before Activate**
  The Before Activate phase runs the scripts before any files are activated on the subscriber. You can run any number of scripts during this phase. Name the directory **BeforeCommit**.

- **After Activate**
  The After Activate phase runs the scripts after all files are activated on the subscriber. You can run any number of scripts during this phase. Name the directory **AfterCommit**.
- **After Publication Success**
  The After Publication Success phase runs the scripts after a publication is successful. You can run any number of scripts during this phase. Name the directory **OnSuccess**.
- **After Publication Fail**
  The After Publication Fail phase runs the scripts after a publication fails. You can run any number of scripts during this phase. Name the directory **OnFail**.

Creating the scripts

Any scripts that you want the GLOBAL-SITE agent to use must have the following characteristics.

- The scripts must be executable, in their directory, by the default web server user.
- On UNIX systems, the default user is typically **nobody**.
- The scripts must return a status code value of **zero** on success, and a **non-zero** value on failure.
- Successful script completion is logged. However, the status of the script completion does not affect the status of the publication.

When there are multiple scripts in a directory, the agent tries to run all the scripts. You cannot control the order in which the agent runs the scripts. The controller reports success only when all scripts in the directory return a zero value. If one or more of the scripts exits with a non-zero value, the behavior of the scripts and the returned error value is undefined. If you want to control the execution of multiple scripts in the same directory, we recommend that you create a master script, in a script directory, that executes the other scripts (which are placed in a private directory) in an ordered fashion.

If a script is not successful, you may want to verify the following potential problems: the script is not set as executable, the script has a dependency on a different directory, and whether there are files in the directory that are not scripts.

Installing the scripts on the subscriber

You can manually copy the scripts into the **UserScripts** directory on the subscriber, or you can deploy the scripts using the GLOBAL-SITE Controller.

To copy the scripts to the subscriber from the command line

1. Locate the **installationRoot/cgi-bin/UserScripts** directory.

2. Create a directory for each publication group that requires scripts.
3. Beneath the publication group directory, create a directory for the phase in which to run the script.
4. Copy the script to the newly created directory.
5. The GLOBAL-SITE agent expects to find scripts in the following directories: **BeforeCopy**, **BeforeCommit**, **AfterCommit**, **OnSuccess**, and **OnFail**.

The following is an example of a correct directory structure:

/var/f5/agent/cgi-bin/UserScripts/Pub1/BeforeCopy/precpy.sh

A GLOBAL-SITE publication can deliver scripts to the target directories by using the WebDAV transfer method.

To deliver scripts to the target directories

1. Create a publication group that has a section for each phase that requires scripts.
2. Create subscribers that use the WebDAV transfer method.
3. The path section mapping should look similar to the following example:

Section	Destination path
secPreCpy	/UserScripts/<PubGrp>/BeforeCopy
secPreCommit	/UserScripts/<PubGrp>/BeforeCommit
secPostCommit	/UserScripts/<PubGrp>/AfterCommit
secOnSuccess	/UserScripts/<PubGrp>/OnSuccess
secOnFail	/UserScripts/<PubGrp>/OnFail

4. Run test connections to create the directory hierarchy.

Running the scripts using the GLOBAL-SITE Controller

You can run scripts using the GLOBAL-SITE Controller by checking the **Initiate Scripts** option on the Agent Actions tab on the Subscriber Detail screen. Click the **Detail** button next to the **Initiate Scripts** option to display the phases at which to run the scripts. The details of these screens are described in the online help for those screens.

Transferring files using WebDAV and WebDAV-SSL

WebDAV offers several advantages over traditional FTP:

- WebDAV uses HTTP/1.1, which uses network connections more efficiently than FTP.

- WebDAV can use either basic or digest authentication; digest authentication provides more secure password transfers. (To use digest authentication with a Windows 2000 server, the server must be configured as a domain controller.)
- WebDAV can set the execute permission on files, enabling you to deliver and run scripts on every web server.
- WebDAV on top of SSL can encrypt content as well as passwords.

WebDAV-SSL offers one significant advantage over regular WebDAV: encrypted content files. WebDAV applies secure socket layer (SSL) encryption to both passwords and content files, adding significant security to file transfers. This level of security, however, adds overhead that reduces file transfer performance.

Using the GLOBAL-SITE agent requires that sections use either the WebDAV or WebDAV-SSL transfer method. For more information, see *Controlling servers with the GLOBAL-SITE agent*, on page 4-20.

WebDAV uses HTTP/1.1 to transfer files. If you have installed the GLOBAL-SITE agent, you can use WebDAV even if your web server does not support it. The web servers that do support WebDAV are IIS5.x or Apache HTTP Server with mod_dav. For more information on WebDAV in general, see **http://www.webdav.org**.

Section updates using FTP-Push

FTP-Push is a file transfer method to the GLOBAL-SITE Controller that is based on basic FTP. FTP-Push enables sections to push content to the GLOBAL-SITE Controller. When using FTP-Push, the GLOBAL-SITE Controller acts as an FTP server. You can also configure FTP-Push sections to trigger the delivery of a publication.

FTP-Push has two distinct processes:

- You can use FTP-Push to transfer files to the GLOBAL-SITE Controller to await publication. FTP-Push sections contact the GLOBAL-SITE Controller and transfer files when content has changed.
- You can use FTP-Push sections to trigger the delivery of new content to the subscribers, if you are using FTP-Push to move files to the GLOBAL-SITE Controller.

WARNING

We recommend that you not attempt to use FTP-Push if you are not familiar with FTP commands.

You must follow this sequence to create an FTP-Push section. First, create the publication (see *To add a publication*, on page 3-5) if one does not already exist. You can share FTP-Push sections on more than one publication. Second, create a section and specify **FTP-Push** as the transfer

method. Third, add content to the FTP-Push section (see *To add content to an FTP-Push section*, on page 4-32). (You can use FTP-Push in conjunction with a web development environment to FTP files directly to the GLOBAL-SITE Controller.) Fourth, set up the FTP-Push section to trigger delivery of content to subscribers. This step is optional.

◆ Note

Note that the GLOBAL-SITE Controller derives the user ID and path from the section name.

To create an FTP-Push section

1. From the Create a New Section screen, select **FTP-Push** as your **Transfer Method**. (You must select the transfer method first when setting up FTP-Push.)
   The screen refreshes when you select **FTP-Push** and the only available boxes are **Section Name**, **Description**, and **Password**.
2. Fill in the available boxes. For more information click the Help button on the top right of the screen.
3. Click the **Create** button.
   The Section Detail screen opens with the modified information.

To modify an FTP-Push section

You can modify only **Description**, **Enabled**, and **Password** on the Section Detail screen for FTP-Push.

1. In the navigation pane, click **List Sections**.
   The Sections List screen opens.
2. In the Publication Membership column, click a publication name associated with the FTP-Push section.
   The Publication Sections screen opens.
3. In the Name column, click the section name.
   The Section Detail screen opens.
4. Make any necessary changes to the available attributes. For more information about this screen, click the Help button at the top right of the screen.

To add content to an FTP-Push section

You add content to an FTP-Push section by opening an FTP session on the GLOBAL-SITE Controller using the FTP-Push section's user ID, password, and path. Using standard FTP commands, create and remove directories, add and remove files.

WARNING

We recommend that you not attempt to use FTP-Push if you are not familiar with FTP commands.

1. Start an FTP client (not on the GLOBAL-SITE Controller) on the machine with the content and open a connection to the GLOBAL-SITE Controller.
2. Log in using the user name and password you specified when you created the section.
3. Use the FTP command **PUT** to deliver files to the correct directories on the GLOBAL-SITE Controller machine that is publishing the content. (You may have to use the FTP command **MKDIR** to make directories before you copy files to them.)
   The files are transferred to the GLOBAL-SITE Controller machine.

Note

If you are using a graphical user interface for FTP, follow your software's directions for moving files.

To retrieve a previous file version from an archived FTP-Push session

You can use FTP to retrieve files from a previous version of an archived FTP-Push section.

1. Using FTP, log on to the GLOBAL-SITE Controller using the FTP-Push user ID and password.
2. Create the FTP version directory by appending a plus sign (+) and the version number to the FTP-Push section path.

   ```
   /<section name>+<version#>
   /section+28
   ```

3. Use FTP to change the working directory to the FTP version directory.

Setting up FTP-Push to trigger section delivery

If you set up an FTP-Push section in a publication to trigger delivery, once the FTP file delivery is complete for that section and the FTP session is logged off, the GLOBAL-SITE Controller delivers all the sections of that publication to the subscribers.

You may have five sections with content that comes in at different times, but you want to publish all the sections concurrently after they all have updated content.

To control when FTP-Push sections publish, you can set up a section that has no content. You configure only this section to trigger delivery (see *To use an FTP-Push section to trigger a publishing cycle* following). When you are ready to deliver the publication, you transfer a dummy file to this section, triggering the delivery of all sections for that publication to the subscribers.

Delivery of a publication to subscribers starts after FTP logout. How long after FTP logout depends on how you set **Initiate ___ minutes after FTP logout** on the Publication Options screen. If other FTP sessions log in before the delay times out, the timer starts anew after that session logs out.

If you do not want to use FTP-Push to trigger delivery, you have the option of starting the delivery manually or on a schedule (see *Scheduled activation of new content*, on page 4-4).

WARNING

*FTP-Push does not trigger delivery until after FTP logout. If another FTP session starts before **Initiate ___ minutes after FTP logout** times out, the timer resets and does not restart until after the current FTP session logs out. If FTP sessions keep logging in before the timeout, it is possible to never propagate the section.*

To use an FTP-Push section to trigger a publishing cycle

1. In the navigation pane, click **List Publications**.
   The Publication List screen opens.
2. In the Name column, select an existing publication.
   A Publication detail screen opens.
3. Click the Pub. Options tab.
   The Publication Options screen opens.
4. Under Initiating Publishing Process, click the **FTP-Push** button to choose FTP-Push to trigger publishing.
   The **Initiating Sections** box becomes available.
5. Click the FTP-Push section(s) that you want to trigger a publishing cycle for this publication. To select more than one name, hold down the Shift or Ctrl key as you click.

   If you are using a dummy page to trigger publishing, select only that section. (See *Setting up FTP-Push to trigger section delivery*, on page 4-32 of this chapter.)
6. If you want to delay delivery after the FTP session logs off, type the number of minutes (0-60) you want to wait.
7. Click the **Apply Changes** button.
   The Publication Options screen refreshes.

Defining exceptions to sections

As you set up your basic publications and create various sections, you may want to create ***exceptions*** to the section. Exceptions instruct the GLOBAL-SITE Controller to disregard specific subdirectories. An exception is a way of excluding a subdirectory from being published as part of a section.

The most challenging part of setting up exceptions is planning the publication in advance, so that sections and exceptions work to your benefit. Remember that, when first created, each section includes all files and directories at the given path. To exclude a subdirectory from the section, you must clear the box for that subdirectory. Note that you cannot exclude files, only directories, on the Section Browser screen. See *Within the directories where you store your content, you may also have files that you do not want to publish or deliver to your subscribers. With the GLOBAL-SITE Controller you can exclude files by their file type for this reason. Or you may have all types of files in a directory, but only want to publish the .html files. The GLOBAL-SITE Controller also allows you to include only file types that you specify.*, on page 4-36 for information on listing file extensions to include or exclude.

Creating exception paths

There are two ways to create exception paths. For one you can browse the directories, for the other, you need to know the directory path that you want to exclude.

Example

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

Assuming you have already created your initial section, start at the Section Detail screen to create an exception path (in our example, it is **sec1**).

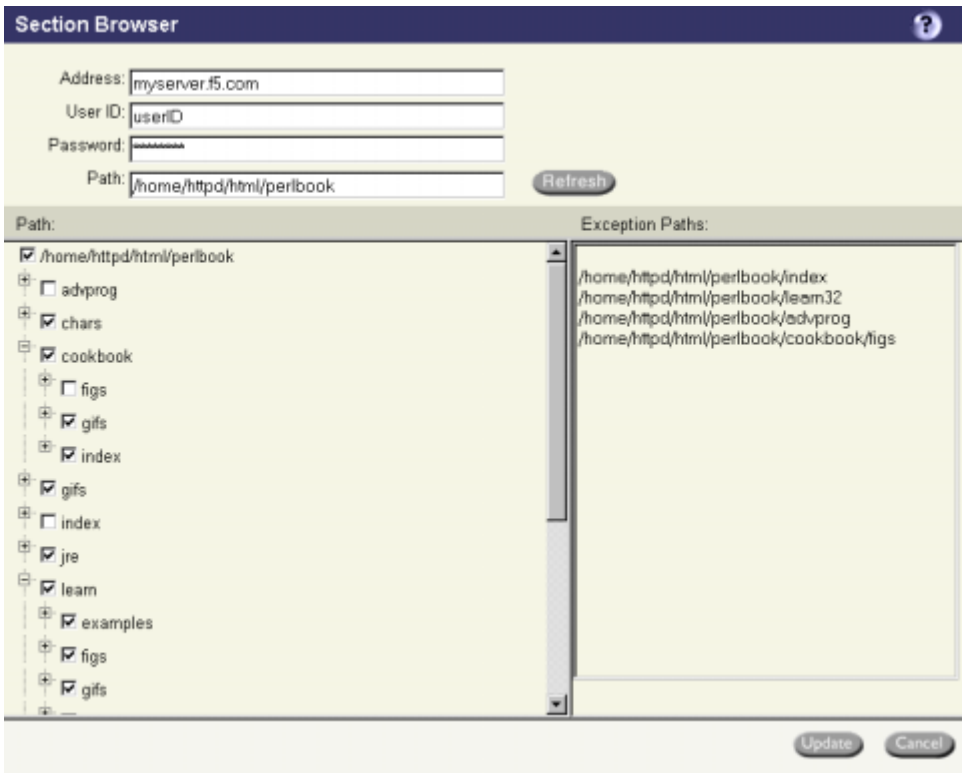

Figure 4.4 *The Section Browser screen, clear boxes for exceptions*

To create an exception using the Section Browser screen

Start at the Section Detail screen.

1. Next to the **Path** box, click the **Browse** button.
   The Section Browser screen opens.

   This example assumes the path is already defined. The four boxes at the top of the screen display the content server address, the user ID and password needed to access the server, and the path.

2. In the Path area at the lower left, the path appears in the top line: **/home/webcontent/testsite**. The remaining lines contain the names of the subdirectories in the path. Clear the box in front of the **images** subdirectory.
   The exception path appears in the Exception area to the right.

3. To save your settings, click the **Update** button.
   The Section Detail screen opens again, with the exception path listed below the **Path** box.

To create an exception using the New Exception box

We recommend you use this method to create exceptions when you have a lot of directories. 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/testsite/images**.
2. Click the **Add** button directly to the right of the **New Exception** box.
3. Click the **Save** button to commit the changes to the section.

◆ 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. For subscribers, it can create 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. Even in that context, that path must already exist, or you must first create it manually on the subscriber.

For 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.

Including or excluding file extensions

Within the directories where you store your content, you may also have files that you do not want to publish or deliver to your subscribers. With the GLOBAL-SITE Controller you can exclude files by their file type for this reason. Or you may have all types of files in a directory, but only want to publish the **.html** files. The GLOBAL-SITE Controller also allows you to include only file types that you specify.

To include or exclude file extensions

1. In the **File Filter(s)** box, type one or more file extensions, separated by spaces. For example, **cgi gif html**.
2. Click **Exclude** to keep files with these extensions out of the path, or **Include** to add files with these extensions to the path.

5

Monitoring Publishing Processes

- Managing publishing
- Making ongoing configuration changes
- Understanding publication logs
- Understanding the GLOBAL-SITE Controller phases

Managing publishing

In addition to enhanced publishing features, the GLOBAL-SITE Controller includes features that can help you control and manage your publications. One of these features is the ability to make ongoing changes to your publications. Another feature is the set of logs that provide insight into the complete publishing process. A third feature outlines the foundations of the GLOBAL-SITE publishing process.

- **Making ongoing configuration changes**
  Once you have your publications set up and working well, you may find that changes in sections, 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 process logs**
  The GLOBAL-SITE Controller has three process logs that provide insight into the publication process at different levels. With these logs, you can see where problems are occurring and take steps to correct them.
  - The Publication log shows you information about either an entire publication or one edition.
  - The System log screen shows the publishing history for all publications and all editions.
  - The Error log shows system-level messages about the publishing stages for a publication, sorted by distributor.
- **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 potential errors. Familiarity with the phases of the publishing process also helps you grasp the significance of the status messages that are provided throughout the publishing process.

Making ongoing configuration changes

After you initially set up a GLOBAL-SITE Controller or define a publication, various situations can arise that require changes to the configuration. This section provides overviews of some 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.
- You want to change a section to use FTP-Push.
- You want to change initial system configuration settings.
- Other configuration changes may be required.

Working with unavailable subscribers

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

Marking a path offline, and back online

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

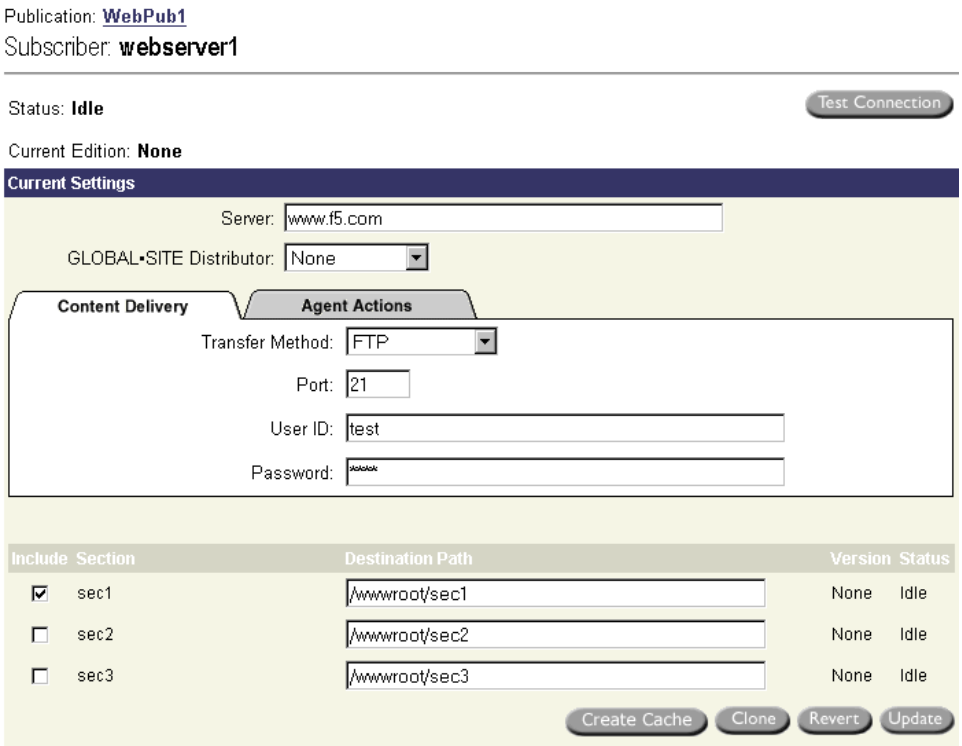

Figure 5.1 *The Subscriber Detail screen*

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 mark this way are ignored by the attempt to publish. And each must be returned to **Include** status (by checking the check box) before you can publish it again.

◆ Note

All sections that share the same subscriber path must have the same Include status, because any sections on the same subscriber path are published regardless of their Include status.

To re-enable a section for a particular subscriber

For each subscriber that was offline, when it returns online, you need to include sections in order to publish to it. Before you can change the publication, you must check its status on the Publication detail screen. You cannot change publication settings unless the publication is in an**Idle** state. For details on checking publication status, see *Displaying status*, on page 3-19 of this guide.

1. On the navigation pane, click **List Publications**.
   The Publications List screen opens.
2. Click the publication you want to work with.
   A Publication detail screen opens.
3. Click the Subscribers tab.
   The Publication Subscribers screen opens.
4. Click the name of the subscriber you want to change.
   The Subscriber Detail screen opens.
5. Under Include, check the box next to each section name you want to include in the next publication.
6. Click the **Save** button.
   All sections with a check mark are 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.

◆ Note

All sections that share a common subscriber path must have the same Include status, because if one subscriber path is enabled, then they are all enabled.

You need to exclude each subscriber that is offline from the publication, for each section, before attempting to deliver the edition. If you do not do this, the delivery gets an error and publication stops. Depending on how you have set your error handling, you may have to restart your delivery from the beginning.

To omit a subscriber from delivery of a particular edition

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 does not receive the edition.

Any subscribers or 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 omit a subscriber from delivery of 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 or 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 include a subscriber for delivery an edition

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

For archived publications, on the Deliver screen:

1. Check the **Include** box for each section.
2. When you have checked the boxes for each section that you want to include, click the **Deliver Edition** button.

To include a subscriber for delivery of a non-archived publication

All subscriber sections that are selected with a check mark are included in the next published publication.
For non-archived publications, 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.

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.

Removing a subscriber from a publication

The publication retains information about the last published section version on a subscriber unless the subscriber is deleted.

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.
   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 on the GLOBAL-SITE Controller. There are two ways to remove old versions of your publications from the GLOBAL-SITE Controller: by group or individually.

Removing unused version by group

If you have a series of editions to remove from your archived publication, you can now do it faster with the Remove Unused Editions button than by using the delete button (trashcan icon) on the Publication Editions screen. Use this feature to remove groups or series of old editions that are no longer being used in a current publication. This option cleans up the listing on the Publication Editions screen, and makes it easier for you to focus on the current editions.

Before deleting any editions, the GLOBAL-SITE Controller considers your criteria, and applies its own criteria to the editions specified. To be removed, editions must meet these conditions:

- The edition cannot be published anywhere; it cannot belong to any currently published subscriber.
- The publication must be in an Idle or Unavailable state.
- There can be no parent/child relationship for any edition you are deleting.
- The edition must be within the range you specified to remove.

Regardless of your choices, the GLOBAL-SITE Controller does not remove your last published edition.

◆ Note

Be careful when specifying removal by days, as the GLOBAL-SITE Controller removes files that are as little as 1 minute older than the time you specified to keep.

To remove unused editions from the controller

Use this option when you want to delete several editions in a series. This button is available only when there are two or more editions in an archived publication.

1. Click the **Remove Unused Editions** button
   The Remove Editions popup window opens.
2. Click the button next to the option you want to use to remove the editions. The choices are:
   - **Remove editions __ through __**
     You can choose, by number, which editions you want to remove.
   - **Remove all but the last __ editions.**
     You can choose to remove all the editions except the last few or several depending on your needs.
   - **Remove all but editions from the last __ days.**
     You can use time to determine which editions to remove.
3. Click the **Remove Editions** button.
   The Remove Unused Editions screen closes and the removal of the unused editions begins.

Removing unused editions individually

You do not have to delete a series of versions, you can elect to delete them one at a time.

To purge versions

Purging versions is a two-step process.

1. 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; therefore its listing does not show a Delete button.
   - You should not remove any other editions that you may want to roll back to or publish to other subscribers at some point.
2. Once you have deleted all unnecessary editions, click a section name.
   The Section Detail screen opens.
3. Click the Section History tab, and then click the **Remove Unused Versions** button.

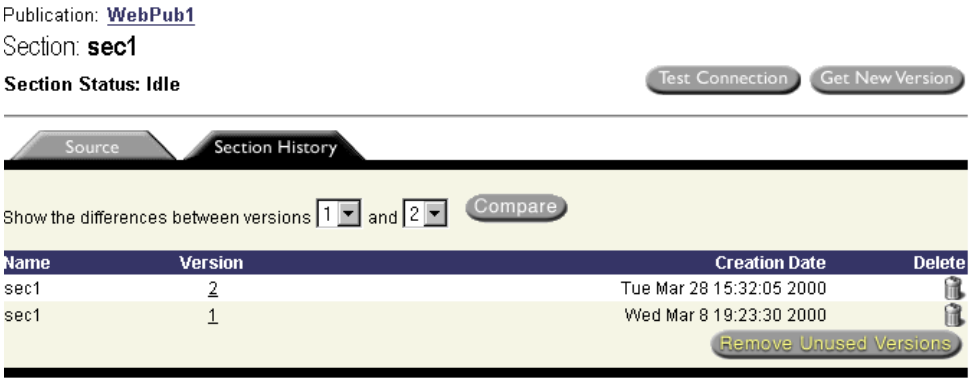

Figure 5.2 *The Section Detail screen, History tab*

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.

Switching sections in an existing publication to use FTP-Push transfer method

A section that uses the FTP-Push transfer method is different in several ways from a section that uses the FTP or WebDAV transfer method. FTP and WebDAV sections store their data on distributed servers, while an FTP-Push section stores its data on the GLOBAL-SITE Controller. If set up to do so, only an FTP-Push section can trigger a publishing cycle. The GLOBAL-SITE Controller creates an FTP user account for an FTP-Push section.

You can convert an existing section (archived or non-archived) to an FTP-Push section using a two-step process. First, on the Section Detail screen, change the Transfer Method to **FTP-Push**. Second, move the section's existing data to the GLOBAL-SITE Controller using a UNIX tar file or a GLOBAL-SITE Controller. We recommend moving the section files in a tar file because it preserves the file dates and publication history for the section.

Before you start, verify that you have a UNIX command line shell, the tar utility, and permissions to the gsite account on the GLOBAL-SITE Controller.

Starting the section move

To start the section move, you must change the section to an FTP-Push section, and move the files for the section to the GLOBAL-SITE Controller.

To change a section to an FTP-Push section

Note the server name, user ID, password, path, and exception paths for the existing section. You need these if you use the GLOBAL-SITE Controller to move the section data.

1. On the Section Detail screen, select **FTP-Push** in the **Transfer Method** box.
   The screen refreshes.
2. Type a new, unique section name.
3. Type a description, if desired. This step is optional.
4. Type a password.
5. Click the **Create** button.
   The screen refreshes again. Note that the GLOBAL-SITE Controller fills in the user ID and path, which it derives from the section name.
6. Type a description and password for the new user and click the **Save** button.

Note the new user ID, password, and path for the FTP-Push section. You need this information for FTP access to the FTP-Push section.

To move section data using a tar file

Verify that you have a UNIX command line shell, the tar utility, and permissions to the gsite account on the GLOBAL-SITE Controller.

1. Use the command line to set the working directory on the source server to the section's old path.

   ```
   cd <path to working directory>
   ```

2. Compress the contents of the directory to a tar file, avoiding the old exception paths. (Copying the directory in a compressed format preserves the file date markers, enabling the GLOBAL-SITE Controller to preserve the publication history.)

   ```
   tar cf /usr/tmp/<tar file name>.
   ```

3. Copy the tar file to the GLOBAL-SITE Controller using a protocol like FTP.

   ```
   ftp <global-site controller>
   User: gsite
   Password: <gsite password>
   cd /gSITE/Data/GSID<gsid number>/gsite/DZ/<ftp-push section name>
   put /usr/tmp/<tar filename>
   by
   ```

4. Log on to the GLOBAL-SITE Controller gsite account and set the working directory to:

```
/gSITE/Data/GSID<gsid>/gsite/DZ/<section name>/<section
name>
```

where <gsid> is the GLOBAL-SITE ID for that server, and <section name> (used twice) is the name of the FTP-Push section.

5. Extract the files to the new directory.

```
cd /gSITE/Data/GSID<gsid number>/gsite/DZ/<ftp-push
section name>/<ftp-push section name>
```

```
tar xf ../<tar filename>
```

To move section data using a GLOBAL-SITE Controller

Note that this method is much slower than the tar file method (see *To move section data using a tar file*, on page 5-9 for more information). You can use this method for both archived and non-archived publications. The procedures are the same except where noted.

1. Add a publication from the New Publication screen.
   - You can accept all the defaults.
   - See *To add a publication*, on page 3-5, or click the Help button on this screen for more information about adding a publication.
2. Create a new section, which becomes the source for your new FTP-Push data.
   - Choose **FTP** for the Transfer Method.
   - See *To create a section*, on page 3-9 of this guide or click the Help button on this screen for more information about creating a new section.
3. Add a new subscriber, which becomes the destination for your FTP-Push data.
   - In the **Transfer Method** box choose **FTP**.
   - Create a new server, user ID, password, and path for the FTP-Push section.
   - Do not specify a distributor.
   - See *To add a subscriber*, on page 3-14 of this guide or click the Help button on this screen for more information about adding a new subscriber.
4. Initiate the publishing cycle to transfer the original section data to the new FTP-Push section. Click the publication name (created in step 1 of this procedure) at the top of the screen.
   A Publication detail screen opens.
5. Click the **Deliver Edition** (for archived publications) button or the **Deliver** button (for non-archived publications) at the bottom of the screen.
   The Deliver screen opens.

6. For archived publications, in the **Edition** box, select **Create New Edition**. Leave the other boxes with their default values and click the **Deliver Edition** button.
   For non-archived publications, leave the defaults and click **Deliver** at the bottom of the screen.
   The Publish Progress Display screen opens. If the Status column says **Success**, you have successfully moved the section data.

Completing the section move

To complete the section move, you must delete the temporary subscriber, edition, section, and publication that you just created. Then, you can publish the new FTP-Push section to subscribers.

To delete the temporary subscriber

1. Click the publication name at the top of the screen.
   A Publication detail screen opens.
2. Click the Subscribers tab.
   The Publication Subscribers screen opens.
3. Click the delete button  next to the temporary subscriber.
4. Click the **OK** button to confirm the deletion.
   The Publication Subscribers screen refreshes. This deletes only the subscriber, leaving the data in its new location.

To delete the temporary edition

1. Click the publication name at the top of the screen.
   The Publication Editions screen opens.
2. Click the delete button  next to the temporary edition.
3. Click the **OK** button to confirm the deletion.
   The Publication Editions screen refreshes.

To remove the temporary section from the publication

1. Click the publication name at the top of the screen.
   The Publication Editions screen opens.
2. Click the Sections tab.
   The Publication Sections screen opens.
3. Click the **Edit Section List** button.
   The Edit Section List screen opens.
4. Select the temporary section from the **Sections in** list on the right, and click the **Remove Section** button.

5. Click the **Save** button.
   The Publication Sections screen refreshes, the section is removed, and your changes are saved.

To delete the temporary section

1. In the navigation pane, click **List Sections**.
   The Sections List screen opens.
2. Click the delete button  next to the temporary section.
3. Click the **OK** button to confirm the deletion.
   The Sections List screen refreshes and the temporary section is deleted.

To delete the temporary publication

1. In the navigation pane, click **List Publications**.
   The Publication List screen opens
2. Click the delete button  next to the temporary publication.
3. Click the **OK** button to confirm the deletion.
   The Publications List screen refreshes.

To publish the new FTP-Push section

1. In the navigation pane, click **List Sections**.
   The Sections List screen opens
2. In the Publication Membership column, click the name of a publication associated with the new FTP-Push section.
   The Publication Editions screen opens (for archived publications) or the Publication Sections screen opens (for non-archived publications).
3. Click the **Deliver Edition** button (the **Deliver** button for non-archived publications).
   The Deliver screen opens.
4. In the **Edition** box (only for archived publications), select **Create new Edition**.
5. Click the **Deliver Edition** button.
   The Publish Progress Display screen opens. When the **Status** message in the upper left says **Idle**, the publishing cycle is finished.

For archived publications, the GLOBAL-SITE Controller compares each file's content with the previous archived version. Only the changed content is published.

For non-archived publications, the GLOBAL-SITE Controller publishes all content.

Changing system settings using the System Settings screen

Clicking **Show Settings** on the navigation pane opens the System Settings screen. From this screen you can stop or start the scheduling process and the FTP server, and you can set the system date and time.

The System Settings screen displays the publication scheduler status, the FTP server status, the GLOBAL-SITE Controller date and time, and the RAID status.

The time is displayed to the minute and the screen refreshes automatically every 60 seconds so that it is never more than a minute off from the time the GLOBAL-SITE Controller is using.

For detailed instructions on how to change the system settings, click the Help button on the top right of the screen.

The System Settings screen now displays the status of the GLOBAL-SITE Controller disks and partitions, for systems that have multiple disk configurations.

You can use the lower part of the System Setting screen to do a quick check on the disk status and drive status if your system is configured with multiple disks. The top row of the table clearly shows any applicable partitions, by name, along with the Type and Status of each. The bottom row of the table shows icons indicating whether the disk is active or inactive.

If the partition is showing a status of sync error, you probably should resynchronize the disks. Contact your F5 Networks representative, who can correct this for you.

Changing system settings using the ITCMconsole command line interface

From the command line, you can run the ITCMconsole command line interface to change any of the attributes you set using the First-Time Boot utility. For more information about running command line, please see *ITCMconsole commands*, on page 6-1 of this guide. To start the ITCMconsole, type **ITCMconsole** at the command prompt.

Understanding publication logs

The log files in the GLOBAL-SITE Controller are easy to read and provide more information than ever before.

- The Publication Log screen provides a single place to look for logs about an individual publication.

- The System Log screen provides a single place to look for logs about all your publications.
- The Error Log screen provides a single place to look for logs about a specific delivery of an individual publication. In addition, on the Publication Editions screen, you can click any date in the Last Published column to open the log for that edition.
- The Verification log shows any differences between the database and the subscriber files. (This log is only created if you click the **Verify Content** button on the Deliver or Deliver Editions screen.)

Publication log

The Publication Log screen shows the publishing history for either one publication or one edition. The editions are listed in reverse chronological order, with the most recent logs at the bottom of the screen. Up to four weeks of data is available from this screen.

To see the log for a publication

1. On the navigation pane, click **List Publications**.
   The Publication List opens.
2. Click the publication you want to see log files for.
   The Publication Editions screen (for archived publications) or the Publication Sections screen (for non-archived publications) opens.
3. Click the **Pub. History** button at the top right of the page.
   The Publication log for that publication opens.

Note

*There is a **Pub. History** button at the top of all the Publication detail screens (Publication Editions, Publication Sections, Publication Subscribers, Publication Options).*

To see the log for one edition

On the Publication Editions screen, click the date for that edition in the Last Published column. The Publication Log screen displays.

System Log

The System Log file displays the log files for all publications associated with the GLOBAL-SITE Controller. The System Log file keeps an historical view of approximately four weeks of activity for all publications sorted by date and time. The information here is the same data in the Publication Log, but files are shown by individual publication.

To see the System Log

Click **System Log** on the navigation pane.

Error Log

The Error Log is a good place to start for troubleshooting delivery problems. It is made up of two separate screens: Retrieval Log and Delivery Log. The log screens display information by publication, and the screen is divided into segments by primary GLOBAL-SITE Controller (listed first) and then each distributor. Each time a publication is delivered (whether delivery is successful or not), a new log file overwrites the old log file.

For archived publications, the Retrieval Log is created only when you create a new version of a section for a publication. This is important to note because it is possible to have a Delivery Log that references a different (later) delivery than the Retrieval Log. The **Last Retrieval Log** link always takes you to the most recent Retrieval log, but this may not be the retrieval log that corresponds to the most recent Delivery log.

For non-archived publications, the Retrieval Log and the Delivery Log are created each time you deliver a publication, therefore both logs reference the same delivery.

Reading the Error Log

When a phase is completed without error, the line in the log file says **Successful** followed by the message about the phase, for example, **Successful return from DoCreateCDirs**.

Error messages in the log file contain a status identifier (which starts with a pound sign [#] and is followed by the message number in hex), and then the error message itself. For example, **#00040501:Could not login to remote system**.

Sample error messages

- **FTP cannot log in**
  You could have an invalid authentication user ID/password pair.
- **Not enough space on the subscriber server**
  The GLOBAL-SITE Controller cannot deliver all the files because the subscriber server is out of disk space.
- **Lost network connection**
  You could be having network problems that are preventing the GLOBAL-SITE Controller from communicating with a subscriber or distributor.

To see the Error Log

1. On the navigation pane, click **List Publications**.
   The Publication List screen opens.

2. Click the publication about which you want to see log files. A Publication detail screen for that publication opens.
3. Click the **Show Details** button on the top right of the screen. The Publish Progress Display screen opens
4. Click the **Show Logs** button at the lower right of the screen. The Error Log screen opens.

Verification log

With the Verification log you can see any differences between the database and the subscriber files. The Verification log is created only if you use the **Verify Without Delivery** option from the Publish Options menu.

To read the Verification log

The following may help you in understanding the Verification log.

- The log is separated into the following hierarchy: Distributor, Subscriber, Section. The log displays subscribers as: **<section name> for <path> on <node> for <user> (<GSidentifier of distributor>)**.
- Files found to have differences are labeled: **Changed**, **Absent**, or **Extra**. **Changed** (files that have a different file size) and **Absent** files (files that are missing from the subscriber) are replaced if you choose to reconcile files on the Publish Progress Display screen. **Extra** files (files that exist on the subscriber but not in the database) are ignored (not deleted nor altered).
- The letter **F** indicates a file with a discrepancy.
- The letter **D** indicates a directory with a discrepancy.

◆ Note

Changes in the case of directory names on case-insensitive systems, such as Windows NT, can be a source of spurious errors. The GLOBAL-SITE Controller is case-sensitive and when the case of the directory name has changed, reports that a change has occurred. This is generally harmless, except in the Verification log, which shows that the subscriber has an extra directory and the original directory is absent. The directory is actually there.

To see the Verification log

After delivery is complete, click the **Show Differences** button at the bottom of the Publish Progress Display screen.

If you have not chosen to verify content, this button does not appear.

Understanding the GLOBAL-SITE Controller phases

When creating publications and running the publishing process, you see publication status indicators. These indicators of where the publication is, or how the publishing process is going, relate to various states and phases. This section discusses the publishing phases in relation to what you may see while monitoring your publishing process.

Interpreting phases in the status line

Various GLOBAL-SITE Controller 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.
- **Getting New Content**
  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.
- **Preparing Publication**
  A brief state that prepares some data structures for the rest of the process and checks that there is new data to publish.
- **Updating Distributors**
  A very brief state, unless distributors are involved. 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.
- **Determining Updates**
  Files are prepared for quick and efficient transmittal to subscribers.
- **Copying Updates**
  Changed files are copied to temporary subdirectories for the subscribers. Archived publications copy data from the distributor to the subscribers. Non-archived publications copy data from the section servers to the subscribers by way of the GLOBAL-SITE Controllers. The process pauses after this phase if the publication has scheduled activation.

- **Activating New Content**
  Changed files are moved from the temporary subdirectories to their final location at the path you specified, and obsolete files are removed from subscribers.
- **Cleaning Up**
  After the publishing process has completed, the GLOBAL-SITE Controller attempts to clean up any files or directories created to support publishing and, if necessary, to rollback the subscriber(s) to the previous state.

Note

*If the publication is set to **Pause publishing, log the error and wait for user input**, the publication pauses in the phase in which it had the error.*

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 is new or 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 **Copy** phase, you can then click the **Continue** button to complete the publication cycle for any subscribers that have already had all section files copied to them.

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

Using the Phases of Publication table

Table 5.1 describes what is happening to the publishing process during several overall publication state and section or path state combinations.

Keep in mind that in all phases up to the **Copying Updates** phase, the status is reporting on sections for each distributor. After that, the **Copying Updates** and **Activating New Content** 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.
Creating Edition, Edition Ready	N/A	Only in these states briefly during the creation of a new edition.
Edition Ready Failed	N/A	The process has failed while creating a new edition from the latest section versions.
Getting New Content	Updating	In this state when performing the potentially lengthy process of determining which files to fetch from section source and fetching them.
Idle	Successful	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.
Getting New Content Incomplete	mix of **Successful, Failure**, and **Cancelled**	When multiple sections are being updated, not all completed.
Getting New Content Failed	Failure	All section updates failed.
Publication Preparation	N/A	A brief phase that constructs a work list for each GLOBAL-SITE publisher and distributor from the configuration database.
Pub Prep Done	Idle	Successful completion of **Preparing Publication**.
Pub Prep Failed, Pub Prep Suspended	N/A	An error message displays, typically a distributor was not reachable. The delivery cannot proceed. **Reset** resets the pub state to **Idle** so it can be retried.
Updating Distributors	Determining Distributor Updates	Creates the changes on the publisher GLOBAL-SITE Controller necessary to bring each distributor up-to-date.
Updating Distributors	Determining Distributor Updates Done	The set of changes is transmitted to each distributor.
Updating Distributors	Updating Section on Distributor	Each distributor updates its section(s) with the received changes.
Updating Distributors Done	Updating Section on Distributor Done	Ready to proceed to **Determining Updates.**
Updating Distributors Canceled	some **Canceled**	The user clicked **Cancel** during this phase. Cannot proceed with delivery. **Retry** retries the operation. **Reset** takes the publication into **Cleaning Up**.
Updating Distributors Failed	Determining Distributor Updates Failed	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 **Cleaning Up**.

Table 5.1 *The GLOBAL-SITE Controller Phases of Publication*

Overall Publication State	Section or Path State	Description
Updating Distributors Failed	Updating Section on Distributor Failed	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 **Cleaning Up**.
Determining Updates	Determining Subscriber Updates	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.
Determining Updates Done	Determining Subscriber Updates Done	Ready to proceed to **Copying Updates.**
Determining Updates Canceled	Canceled	The process was canceled by the user. The publication cannot proceed further. **Retry** retries the operation and **Reset** takes the publication into **Cleaning Up**.
Determining Updates Failed	Determining Subscriber Updates Failed	An error, typically lack of disk space, is displayed. The delivery process cannot proceed. **Retry** retries the operation and **Reset** takes the publication into **Cleaning Up**.
Copying Updates	Copying Updates	Each GLOBAL-SITE Controller copies new and changed files to temporary subdirectories on each subscriber.
Copying Updates Done	Copying Updates Done	Ready to move to **Activating New Content.**
Copying Updates Suspended	Copying Updates Done	The publication has a scheduled activation. It proceeds to **Activating New Content** after clicking **Continue** or at the scheduled time. **Retry** retries the operation and **Reset** takes the publication into **Cleaning Up**.
Copying Updates Canceled	**Canceled** plus various	The user clicked **Cancel**, stopping subscriber path copies in progress, as well as those not started yet. If you click **Continue**, only subscribers that already successfully completed the **Copying Updates** phase for all sections can proceed to **Activating New Content**. Clicking **Reset** sets the publication state to **Cleaning Up** and **Retry** retries the operation from the point at which it was cancelled.
Copying Updates Failed	Copying Updates Failed	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 for the failed path from the point they first failed. **Reset** takes the publication into **Cleaning Up**.
Activating New Content	Activating New Content	The copied new and changed files are being renamed from their temporary subdirectories to their final location.
Activating New Content Done	Activating New Content Done	The publication proceeds to update the configuration database and set its state back to **Cleaning Up**.
Activating New Content Suspended	Activating New Content Suspended	Proceeds to **Cleaning Up** after you click **Continue**. Click **Retry** to retry the **Activating New Content**; click **Reset** to cancel the publishing.

Table 5.1 *The GLOBAL-SITE Controller Phases of Publication*

Overall Publication State	Section or Path State	Description
Activating New Content Canceled	**Canceled** plus others	The **Cancel** control was clicked, which stopped all activates in progress, as well as any not yet started. Avoid this if possible, as it leaves incomplete subscribers in an indeterminate state. Click **Retry** to retry the activate from the point at which it was canceled and **Reset** to move into **Cleaning Up**.
Activating New Content Failed	Activating New Content Failed	Error message indicates the problem, usually communication failure. Click **Retry** to retry the activate for the failed paths from the point they first failed. Click **Reset** 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.)
Cleaning Up	various phase examples, all with options [..] of **Done** / **Failed** **Deleting Subscriber Updates**, .. **Copying Updates**, .. **Cleaning Up From Aborted Publish**, .. **Activating New Content**, .. **Cleaning Up Outdated Files**, .. **Rollback Done/Failed** **Rollback Done Clean**.. **Rollback Failed Clean**..	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.
Cleaning Up Canceled	as in **Cleaning Up**	The **Cancel** button was clicked while **Cleaning Up** was in progress. Click **Reset** to quit **Cleaning Up**.
Cleaning Up Done	as in **Cleaning Up**	The **Cleaning Up** phase has completed.
Cleaning Up Failed	as in **Cleaning Up**	The process has failed during the **Cleaning Up** phase. Click **Reset** to quit the **Cleaning Up** phase.
Cleaning Up Ready	as in **Cleaning Up**	The controller is ready for the **Cleaning Up** phase, which begins soon.
Delete Versions	N/A	The section is deleting versions.
Test Connection	N/A	The system is currently testing the connection.
Removing From Distributors	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 was tried; the deletion may have failed and should be tried again.

Table 5.1 *The GLOBAL-SITE Controller Phases of Publication*

Understanding cleanup phases

All publishing goes through a cleanup phase. In this phase, the GLOBAL-SITE Controller attempts to clean 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, cleanup is always done for sections, regardless of how far the publishing progressed. At the end of publishing, section states move into the cleanup phase, and progress until the cleanup phase is done. If the section's final state is **Failed**, the next publishing process attempts to clean up before continuing.

The cleanup phase moves through many stages, and these can be reflected on the detail screen. It can be helpful 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 5.1, on page 5-19.

Details on the rollback state are provided in Table 5.2 following.

Rollback state	Explanation
Rollback Done	Rollback completed successfully, but cleaning the path was not attempted, probably because the cleanup phase was cancelled.
Rollback Failed	Rollback failed and path cleanup was not attempted, probably because the cleanup phase was canceled.
Rollback Done, Cleanup Done	Both rollback and path cleanup were successful.
Rollback Failed, Cleanup Done	Rollback failed but path cleanup was successful.
Rollback Done, Cleanup Failed	Rollback was successful, but path cleanup failed.
Rollback Failed, Cleanup Failed	Both rollback and path cleanup failed.

Table 5.2 *Description of the rollback state*

6

ITCMconsole Command Line Interface

ITCMconsole commands

The ITCMconsole command line interface adds new functionality to the GLOBAL-SITE Controller. With these new commands you have an even greater ability to configure and run the controller.

This chapter lists the various ITCMconsole commands, including syntax requirements and functional descriptions. Table 6.1 outlines the conventions used in the command line syntax.

Item in text	Description
\	Continue to the next line without typing a line break.
< >	You enter text for the enclosed item. For example, if the command has **<your name>**, type your name.
|	Separates alternate options for a command.
[ ]	Syntax inside the brackets is optional.
...	Indicates that you can type a series of items.

Table 6.1 *Command line conventions*

Using the ITCMconsole

To run the ITCMconsole, you have to be logged in as the root user. You can use the ITCMconsole in two ways, batch mode and interactive mode. For batch mode, if you are at the command prompt, and you know the entire command, you can type it out in a string (starting with **ITCMconsole**) on one line and then press the Enter key. For interactive mode, you have to be inside the ITCMconsole shell, which you open by typing **ITCMconsole** at the command prompt and pressing the Enter key. Interactive mode is very useful if you are unsure of the exact syntax of the commands, or if you want to see the online help for the commands. You use interactive mode by pressing the Enter key after each command, which allows the system to query you as to what you would like to type next.

For example, the interactive process shown in Figure 6.1 gets the same results as this string:

```
ITCMconsole user logadmin enable
```

```
[root@user /root]# ITCMconsole
 Press ? for help at any time
 Press Ctrl-U to clear the line

ITCMconsole$>user
Commands:
 root                 Modify root user account
 support              Modify the support user account
 logadmin             Modify FTP user account
 gsite                Modify gsite user account

ITCMconsole$>user logadmin
Commands:
 disable              Disable the FTP user (logadmin)
 enable               Enable the FTP user (logadmin)
 password             Change the password for the FTP user
(logadmin)

ITCMconsole$>user logadmin enable
Logadmin account enabled
```

Figure 6.1 *Sample interactive mode for the ITCMconsole*

In interactive mode, you can use auto-complete (also known as tabbed complete). If you type the first few letters of a command and press the Tab key, the command completes itself without your having to type the entire command.

Ctrl + U clears all the text in front of the ITCMconsole prompt. Ctrl + X breaks you out of an interactive command prompt and returns you to the ITCMconsole prompt. The **quit** command closes the ITCMconsole shell.

We present the commands and examples in this chapter assuming that you are using interactive mode. If you are using batch mode, you need to preface each command with the prefix **ITCMconsole**. For example, to use the command **show user** in batch mode, you type:

```
ITCMconsole show user
```

◆ Tip

*You can type **ITCMconsole** or **gsconfig** to open the ITCMconsole command line interface.*

Commands

The following lists in alphabetical order the individual **ITCMconsole** commands, with a concise definition, and a page reference where you can find the detailed description.

?

```
<command> ?
```

The **?** command displays the list of commands or subcommands and the corresponding online help for each.

To see the list of initial commands, at the ITCMconsole prompt, type:

?

Figure 6.2 demonstrates the results when you type **alert ?**

```
ITCMconsole$>alert ?
Commands:
 raise                Raise a named alert with optional data
 add                  Add an alert with an associated action
 delete               Delete an alert and its associated action
 modify               Modify an existing alert
 named_action         Add/Delete/Modify a named action for alerts
 setting              Configure default settings for alerts

ITCMconsole$>alert
```

Figure 6.2 *Example of subcommand list text*

Figure 6.3 demonstrates the results when you type **alert raise ?**.

```
ITCMconsole$>alert raise ?
Raise a named alert with optional data.

Usage: alert raise [name][args]
Where
        [name] is the name of the alert and
        [args] are any action overrides in the form of setting=value pairs.
        Whether overrides are used depends on which actions are configured.

Available settings are:
        description (used for trap messages and email subject lines)
        toaddress (override for To: address field for email actions)
        fromaddress (override for email From: address field)
        attachment (file name for overriding email attachment field)
        body (default body text for email actions)
        command (overrides command to run for exec actions)
```

Figure 6.3 *Example of command help text*

acl

```
acl syntax
acl migrateacl
```

The access control list (ACL) tells the operating system what access rights each user has for a particular object. The .acl file is **/config/default.acl**.

Working with the default.acl file

You should only edit the **/config/default.acl** file if you understand the rules associated with it. Using the **acl** command, you can check the file syntax and migrate the **default.acl** file to the LDAP database. To see a sample ACL file, go to **/config/sample1.acl**.

To check the syntax of the acl file

If you modify the **default.acl** file, you should check the syntax of your changes. To do this, type:

```
acl syntax
```

To migrate the **default.acl** file to the LDAP database, type:

```
acl migrateacl
```

alert

```
show alert [all | named | action | setting]
alert raise [name][args]
alert add [name][action][args]
alert delete [name]
alert modify add_action [name][action][args]
alert modify delete_action [name][action_name]
alert named_action add [action_name][subaction][args]
alert named_action delete [action name]
alert named_action add_subaction [name][subaction][args]
alert named_action delete_subaction [name][subaction]
alert setting add [variable=value]
alert setting delete [name]
```

You use the **alert** command to configure default alerts (see *Default alerts*, on page 7-3) and user-defined alerts (see *Creating user-defined alerts*, on page 7-4). Most default alerts have SNMP traps information associated with them. Alerts with OIDs can be triggered by specific events (see Figure 6.4 for a complete list of the default alerts and their triggers). For more specific information see Chapter 7, *SNMP, MIBs, traps, and alerts.*

To view all the alerts, actions, and global settings use the **show alert all** command.

Creating actions and subactions

User-defined alerts are made up of actions and/or subactions. Default alerts can be modified to use actions and/or subactions. Actions and subactions use one of the following: **exec**, **email**, or **named_action**. Actions and subactions tell the alert what to do when it triggers. **Description** is a required field and should be used to describe the triggering event for an alert.

exec

Use **exec** when you want to run a command when the alert triggers. Commands run as root when triggered. The **exec** command requires a command to run when an alert triggers.

email

Use **email** when you want to send email notification of an event (modifying a default alert) or when you raise an alert. The **email** command requires a **To** address and may have one or more of the following:

- additional **To** addresses
  You can add as many addresses as you want. Separate each address with a semi-colon (**;**).
- **From** address
  The address you want to appear in the email **From** box. This variable is optional, press the Enter key to choose the default [**none**].
- **body**
  The text you want to appear in the body of the email. This variable optional, press the Enter key to choose the default [**none**].
- **attachment**
  The file name of any attachment you want sent with the email alert. This variable is optional, press the Enter key to choose the default [**none**].
- **description**
  The text for the subject line of the email.

named_action

A **named_action** groups together a set of email, exec, or other **named_actions** into a single entity. Once you create a **named_action**, you can reuse it in any number of alerts just by adding its name to the alert. If you need to use the same action more than once, then using a **named_action** is the most efficient way to do this.

Initially the **named_action** is created with a single **subaction**. Additional **subactions** may be added once the **named_action** has been created. The information you need to create a **named_action** depends on the type of **subaction** you use: **email** or **exec**, and the **args** you use (if any). Global settings provide values for any variables (**args**) that you do not give a value to. You use the following command to create a **named_action**.

```
alert named_action add [action_name][subaction][args]
```

Creating user-defined alerts

From the ITCMconsole, you can create your own user-defined alerts using the **alert** commands. User-defined alerts are not triggered by an event, they can however be issued by using the **alert raise** command. (If you want to send an alert based on an event, you need to modify a default alert. See *Default alerts*, on page 7-3 of this chapter.) User-defined alerts do not have an OID and therefore cannot be used to send SNMP traps.

The **show alert all** command displays both default alerts and user-defined alerts. The main information that you need is the name of the alert and then only if you plan on modifying it in some way. You can also see what named actions or subactions you have assigned to an alert. You need the **named_action** (or **subaction**) if you wanted to change or delete it.

To create or modify a user-defined alert

1. From the ITCMconsole, type:

   ```
   alert add
   ```

2. At the prompt for **Alert name**, type a name for the alert.
3. Choose to add either a **named_action**, **email**, or **exec** as your action.
   - A **named_action** is an user-defined action that you created previously. (Please see *named_action*, on page 6-7 for information on creating a named_action.) At the prompt, type the name of the **named_action** you created.
   - An **email** action sends an email for the alert. At the prompt, type **email** (see *email*, on page 6-6 for more information). The steps that follow ask you for the to address, from address, body text, and if you want to add an attachment, the attachment name.
   - An **exec** action runs scripts or executables for the alert (see *exec*, on page 6-6 for more information). At the prompt, type **exec**. The steps that follows asks you for a command to execute when the alert triggers and for a description of the alert.

Modifying default alerts

You can add or delete user-defined actions to the default alerts, but you cannot delete default alerts, change the name of the default alerts, or change the OID. Actions can be either **named_action**, **email**, or **exec**. For example, upon successful publication of a site, the **snmpd** sends the **GLOBALSITE.globalsiteTrapPublicationSucceeded** trap "Publication succeeded" to the destination IP address. You can modify the default alert and add an email action that sends mail to additional people, to inform them of the successful publication as well.

```
/*Alerts:*/
   COMMON.trapDiskDriveFailure
        /*OID*/ ".1.3.6.1.4.1.3375.1.5.110.2.1"
        /*DESCRIPTION*/ "A disk drive has failed."
   COMMON.trapFreeVirtualMem
        /*OID*/ ".1.3.6.1.4.1.3375.1.5.110.2.2"
        /*DESCRIPTION*/ "There is not enough memory to complete an operation."
   COMMON.trapExcessiveCpuUsage
        /*OID*/ ".1.3.6.1.4.1.3375.1.5.110.2.3"
        /*DESCRIPTION*/ "The CPU usage exceeds its threshold."
   GLOBALSITE.globalsiteTrapDataDiskFull
        /*OID*/ ".1.3.6.1.4.1.3375.1.4.110.2.1"
        /*DESCRIPTION*/ "The data disk is full."
   GLOBALSITE.globalsiteTrapConfigurationDiskFull
        /*OID*/ ".1.3.6.1.4.1.3375.1.4.110.2.2"
        /*DESCRIPTION*/ "The configuration disk is full."
   GLOBALSITE.globalsiteTrapPublicationSucceeded
        /*OID*/ ".1.3.6.1.4.1.3375.1.4.110.2.3"
        /*DESCRIPTION*/ "Publication succeeded."
   GLOBALSITE.globalsiteTrapPublicationFailed
        /*OID*/ ".1.3.6.1.4.1.3375.1.4.110.2.4"
        /*DESCRIPTION*/ "Publication failed."
   GLOBALSITE.globalsiteTrapPublicationCheckedForNew
        /*OID*/ ".1.3.6.1.4.1.3375.1.4.110.2.5"
        /*DESCRIPTION*/ "There is no new version of this publication."
```

Figure 6.4 *Default alerts*

To modify a default alert

To modify a default alert, you first need to know the name of that alert. Default alert names can be fairly long and must be typed accurately. You can also find the names of the default alerts in the GLOBAL-SITE MIB, which can be accessed from the GLOBAL-SITE Controller home page or from the Documentation section of the product CD splash screen.

◆ Tip

If you are accessing the GLOBAL-SITE Controller from an SSH application, and this application allows you to copy and paste, we recommend you do so.

1. To get the name of the alert from the ITCMconsole, type:
   `show alert all`
2. If possible copy the name of the alert you want to modify.
3. From the ITCMconsole, type:
   `alert modify add_action`
4. At the prompt for **Alert name** paste or type the name of the alert.

5. Choose to add either a **named_action**, **email**, or **exec** as your action.
   - A **named_action** is an user-defined action that you created previously. (Please see *exec*, on page 6-6 for information on creating a named_action.) At the prompt, type the name of the named_action.
   - An **email** action sends an email when the alert triggers (please see *email*, on page 6-6 for more information). At the prompt, type **email**. The steps that follow ask you for the to address, from address, description (subject line), body text, and if you want to add an attachment, the attachment name.
   - An **exec** action runs scripts or executables for the alert (please see *exec*, on page 6-6). At the prompt, type **exec**. The steps that follows asks you for a command to execute when the alert triggers and a description of the alert.

Changing global settings

There are six global settings that you can assign a value to. These global settings become the default values in the alerts or actions you create. When you use the global settings you do not have to retype the same information over and over for each alert or action you create. If you leave a global setting without a value, the setting is ignored.

The global settings are:

- **description**
  used for default alert messages and email subject lines
- **toaddress**
  the **To** address box for email actions
- **fromaddress**
  the **From** address box for email actions
- **attachment**
  the file name for the email attachment
- **body**
  the default body text for email actions
- **command**
  the command to run for **exec** actions

To set the global settings

Changing or setting the global settings requires that you know the name of the setting you want to change. The names are listed here, but you can also find them by typing:

```
alert setting add ?
```

Then, to change a setting, you type:

```
alert setting add [setting name]
```

For example, to change or set the **description** variable, type:

```
alert setting add description This is a GLOBAL-SITE alert.
```

When the setting changes, you see a confirmation message.

WARNING

*You are not given any warning if you are overwriting a global setting that you already gave a default value. Use the **show alert all** command to see any set global settings.*

Overriding global settings

You use **[args]** to override the global settings within an individual **alert**, **named_action**, or **sub_action**. For example, if you create a new alert, named **globalsite_action** and you want to override the **toaddress** global setting, you type:

```
alert add globalsite_action email toaddress=bob@yourcompany.com
```

date

```
date
date [MMDDhhmm[[CC]YY][.ss]]
```

The **date** command allows you to reset the date and time (using a 24 hour clock) on your machine. The format for the date command is: [MMDDhhmm[[CC]YY][.ss]].

Setting the date and time

You set the date and time using the following system:

- **MM** is the two-digit month number
- **DD** is the two-digit date
- **hh** is the two-digit hour (24 hour clock)
- **mm** is the two-digit minute
- **CCYY** (optional) is the four-digit year
- **.ss** (optional) is for seconds (0 - 61)

So to set the clock for October 10, 2001 at 4 p.m., you could type:

```
date 101016002001
```

or

```
date 10101600.21
```

gs_identifier

```
gs_identifier <identifier>
show gs_identifier
```

The GLOBAL-SITE identifier is a unique, pre-assigned numeric identifier provided by your vendor that distinguishes one GLOBAL-SITE from all others.

A host name must already be set prior to running this command.

The **gs_identifier** command resets the GLOBAL-SITE identifier on the controller. The identifier can be any integer, but it does not accept letters or special characters.

To show the current GLOBAL-SITE identifier, type:

```
show gs_identifier
```

WARNING

*If you change the GLOBAL-SITE identifier on a controller, you must update its **gs_identifier** information on all the GLOBAL-SITE Controllers that use that controller as a distributor. This is done from the web-based user interface from the Distributor List screen.*

ldap

```
show ldap | host | admin | status
ldap sethost <hostname>
ldap setpassword <password> <password>
ldap checkhost [<hostname> [<userdn> <password>]]
ldap checklogin
ldap installhost <LDAP admin dn> <LDAP admin pwd>
ldap enable
ldap disable
ldap start
ldap stop
```

This group of commands configures the LDAP server. The GLOBAL-SITE Controller uses the LDAP server for iControl and web user authentication purposes.

Running the LDAP server

To allow the LDAP server to start when you reboot the computer, type:

```
ldap enable
```

If you do not want the LDAP server to start when you reboot the computer, type:

```
ldap disable
```

To start or stop the LDAP server, type:

```
ldap start
ldap stop
```

Administering the LDAP server

The command **ldap setpassword** modifies the LDAP root password.

You can change the LDAP host used for authentication with the **ldap sethost** command.

Checking the host name for the LDAP server

The **ldap checkhost** command shows whether the host has an LDAP server available for authentication. You can use the fully qualified domain name or the IP address as the **<hostname>** variable.

To see if a login is valid on the LDAP server, you can add the **<userdn>** and **<password>** pair to the command. The **<userdn>** must be a distinguished name, for example, **cn=Manager,dc=f5,dc=com**. To find the **userdn** for the local LDAP server, type:

```
show ldap admin
```

Changing the IP address of the controller

If for some reason you need to change the IP address for the GLOBAL-SITE Controller, you need to run the following command after you make the change:

```
ldap installhost <AdminDN> <LDAP admin password>
```

Changing the LDAP server

If you change the LDAP server you use for authentication, you need to stop and restart the CORBA daemon using the **service corba stop** and **service corba start** commands. For information on the CORBA daemon, see *service*, on page 6-28 of this chapter. Changes to the LDAP server are any of the following:

- remote host to local host
- local host to remote host
- remote host to a different remote host

To change from a remote LDAP server to a local LDAP server

If your controller currently points to a remote LDAP server and you need to set up a local LDAP server, use the following script. Note that when you type the **new LDAP admin password**, you are creating a new administrator password for the local LDAP server.

```
/usr/local/bin/setuplocalldap.sh <new LDAP admin password>
```

To change from a local LDAP server to a remote LDAP server

If your controller currently has an LDAP server, and you want to change to a remote server that is already set up, use the following script:

```
/usr/local/bin/useremoteldap.sh <hostname> <remote LDAP admin password>
```

To change from one remote LDAP server to a different remote LDAP server

If your controller currently points to one remote LDAP server and you want to change it to another remote LDAP server, use the following two commands:

```
ldap sethost <hostname>
```

```
ldap installhost <cn=manager,dc=f5,dc=com> <remote server's ldap
admin pwd>
```

Checking status on the LDAP server

To see if the LDAP server is **enabled** and **running**, type:

```
show ldap status
```

To see what the host name is for the LDAP server, type:

```
show ldap host
```

license

```
license [check] <product>
```

Use **license check** to see if you are licensed for the product you purchased. Type:

```
license check
```

If you are properly licensed, the response will read: Authorized for GLOBAL-SITE. If you are not authorized, please contact support to get a valid license key.

netinet

```
show netinet [config | dns | domains | hostname | defaultrouter | interfaces]
netinet defaultrouter <router> | none
netinet dns <dns> | none
netinet domains <domains> | none
netinet ntp <ntp server> | none
netinet hostname <host name>
netinet hostsfile
netinet interface [add_sub | remove_sub | eth0 | eth1]
netinet route
```

The **netinet** command configures your network devices. You first configured these devices during the First-Time Boot utility. Use this command to make any changes to the configuration of those devices.

For all **netinet** commands except the **netinet ntp** command, you need to reboot for your changes to take effect. To reboot, type **reboot** from the ITCMconsole. For **netinet ntp**, you need to stop and restart the NTP daemon for your changes to take effect. At the ITCMconsole, type the following to stop and restart the NTP daemon:

```
service ntpd stop
service ntpd start
```

Configuring network devices

Table 6.2 lists the devices you can configure with the **netinet** command.

Command	Description
netinet defaultrouter	Set the default router for the system. Use **netinet defaultrouter none** to clear the current settings for the default router.
netinet dns	Set the DNS for the system. Use **netinet dns none** to clear the current settings for DNS.
netinet domains	Set the search domains for the system. Use **netinet domains none** to clear the current settings for the domains.
netinet ntp	Set the NTP server for the system. Use **netinet ntp none** to clear the current settings for the NTP.
netinet hostname	Set the host name for the system.
netinet hostsfile	With normal use of the controller, you should never need this command. If for some reason you change your host name manually (without using the **netinet hostname** command), use **netinet hostsfile** to synchronize software with the current host name.
netinet interface	Configure the network interface(s). See *Configuring interfaces*, on page 6-19 for more information on the **netinet interface** command.
netinet route	Use the **netinet route** command to add, delete, or move static routes. In some network environments you need additional network routes besides the default route. This is common when the GLOBAL-SITE Controller uses multiple network interfaces, or the physical placement of the controller is in a multi-homed network. To add static routes, use the following command: **`netinet route add <INTERFACE> <IP NETWORK> <IP NETMASK> <GATEWAY>`** This example instructs the controller to forward all packets with a destination IP address inside the **192.168.12.0** network to the router at **192.168.11.254**. **`netinet route add eth0 192.168.12.0 255.255.255.0 192.168.11.254`** In order for the newly added network route to take effect, you must reboot the GLOBAL-SITE Controller.

***Table 6.2 netinet** commands*

WARNING

Be aware that changing the host name on a GLOBAL-SITE Controller resets the host name in the GLOBAL-SITE database. If you use this controller in conjunction with other GLOBAL-SITE Controllers, you need to update these controllers with the new host name of the changed controller.

Configuring interfaces

You can change a variety values on the interface using the **netinet interface <interface>** command. Table 6.3 lists these values.

command	Description
netinet interface eth0 broadcast	Set the broadcast IP address for the system interface.
netinet interface eth0 down	When you use this command, the interface is shut down immediately, dropping all current connections. You are not given a confirmation when you press the Enter key. Warning: Do not run this command from a remote interface.
netinet interface eth0 duplex	Configure the duplex mode for the interface. The command returns a choice of duplex modes. Choose the number that corresponds to the duplex mode for your interface.
netinet interface eth0 ip	Set or change the IP address for the system interface. To clear the IP address, use: `netinet interface eth0 ip none`
netinet interface eth0 up	Activate the configured interface. You are not given a confirmation for this action.

***Table 6.3 netinet interface** subcommands*

NIC negotiation

NIC negotiation is the setting of the duplex mode for the network interface card (NIC). You set the duplex during the First-Time Boot utility, but if you change or add a NIC, you need to reset the duplex mode.

To set the duplex mode for an interface

1. You need to choose which interface you want to configure. To display a list of available interfaces, from the ITCMconsole command line, type:

   `netinet interface`

2. From the new prompt, type the name of interface you want to configure.
3. From the new prompt, type **duplex**.
   The ITCMconsole displays a list of available duplex modes.
4. From the list of available duplex modes, type the corresponding number for the mode you want to use.
   You get a confirmation of what mode you set.

Viewing your configuration

You can view your entire current configuration or the configuration of your network devices using the **show netinet** command. For example, to see, the entire configuration, type:

`show netinet config`

The display looks like Figure 6.4.

```
hostname=machine.company.net
defaultrouter=192.168.11.254
dns=192.50.100.1
search_domains=win.net
interface=eth0 ip=192.168.11.240 netmask=255.255.255.0
interface=eth1 ip=0.0.0.0 netmask=0.0.0.0
```

Table 6.4 *Sample output for* ***show netinet config***

Changing the host name of the controller

You may need to change the host name of the controller, for instance if you move it from a test environment into production.

To change the host name of the controller

To change the host name for the GLOBAL-SITE Controller, type the following command:

`netinet hostname <hostname>`

ratelimit

```
show ratelimit [aggregate | classes | class | stats | status]
ratelimit aggregate [max | set| clear]
ratelimit custom [add | delete | modify | promote | demote]
ratelimit on
ratelimit off
ratelimit enable
ratelimit disable
```

The GLOBAL-SITE Controller regulates available Internet connection bandwidth with aggregate and custom rate limiting. If publication delivery degrades server performance, you can use the controller's rate limiting capability to slow down its output to particular servers and at particular times of day. This allows you to better manage available bandwidth and meet your traffic needs.

WARNING

The GLOBAL-SITE Controller cannot limit the rate at which it receives data from clients and servers. It can only limit the rate at which it sends data to clients and servers.

Note

When both aggregate and custom rate limiting are in use at the same time, the custom rates take priority over the aggregate limit, although each individual custom class is limited to the aggregate limit.

Configuring aggregate rate limiting

To configure aggregate rate limiting, you need to know what type of network interface card (NIC) you have in your machine. If you opted to purchase the Gigabit NIC, that interface has a total capacity of 1000 Mbps. The standard NIC is 100 Mbps.

To set or modify the maximum possible bandwidth

Maximum bandwidth is set at installation time. If you are using rate limiting, maximum bandwidth must be modified if a new device is added or if the interface is connected to a hub that is slower than 100 Mbps. (If you are not using rate limiting, the setting for maximum possible bandwidth has no bearing on traffic flow.) The only valid values for maximum possible bandwidth are: 10, 100, and 1000. The command is:

```
ratelimit aggregate max <interface> <max value>
```

For example, to set the maximum bandwidth to 1000, on the **eth0** interface, type:

```
ratelimit aggregate max eth0 1000
```

To add an aggregate rate limit to an interface

Use the **ratelimit aggregate set** command to add an aggregate rate limit to an interface. The command is:

```
ratelimit aggregate set <interface> <Mbps>
```

For example, to set a rate limit of 20 (on a 100 Mbps NIC), type:

```
ratelimit aggregate set eth0 20
```

To clear an aggregate rate limit

The **clear** command removes an aggregate rate limit from an interface. You only need to provide the interface name with the command, not the rate limit. The command is:

```
ratelimit aggregate clear <interface>
```

For example, to clear the rate limit from interface **eth0**, type:

```
ratelimit aggregate clear eth0
```

Configuring custom rate limiting

You use custom rate limiting to limit response data for protocols that use known ports, such as HTTP, when responding to clients. From the ITCMconsole, you can create or delete custom rate limit classes. These rate limit classes allow rate limiting based on specified characteristics of the outbound IP packets, such as source and destination IP addresses and ports.

To add a custom rate limit class

Custom rate limit classes have several variables which you can use to further limit traffic across the interface. If you choose not to use these variables to limit traffic rates, you opt out of using them. The command is:

```
ratelimit custom add <class name> <interface> <limit in Mbps> <non-optional variable(s)>
    [optional variables]
```

Configuring custom rate limiting involves three types of variables, required, non-optional (meaning you have to use at least one), and optional.

1. From the ITCMconsole command prompt, type:

   ```
   ratelimit custom add
   ```

2. At the Class name prompt, type a name for the class.
3. At the Interface prompt, type which interface you would like this class to be on.
4. Enter the limit in Mbps that you want to use.
5. For the non-optional variables, you must enter one or more IP addresses or ports. At the appropriate prompt, type a destination or source IP address, and/or a destination or source port for this class. The default is any, but you cannot have any for all the non-optional values. You must enter at least one actual IP address or port.

6. Choose the relative priority for this class. The default is **5**. The system uses priority to give one class a priority over another class when sending packets. It does not resolve filtering conflicts (see *To promote or demote a class*, on page 6-25 for more information on class order). Filtering traffic into classes is done in the order the classes appear in the show ratelimit classes list.

   You can use priority as a quality of service tool in that you can give one class a higher priority than another, and that gives the customer with the higher priority more bandwidth than another customer.

7. Type **y** or **n** to choose which optional variables you want to use. See Table 6.5 for the optional variables and their descriptions. If you do not want to use optional variables, type **n** for each of the options presented.

Optional variable	Description
Is this class bounded [y]	If this class is ***bounded***, meaning that it cannot borrow capacity from its parent class, choose the default **[y]**. Borrowing from a parent class means that if the overall bandwidth limit is not being exceeded, a class can use more than its limit to meet temporary demand. The parent class in this case is the whole device (as the controller does not allow sub-classes of traffic). Borrowing might adversely affect some other class in the short term, which is why you can prevent a class from borrowing (bound it). If it is not bounded, type **n**.
Is this class isolated [n]	If this class is isolated, meaning that other classes cannot use any excess bandwidth from this class, choose the default **[n]**. Isolating a class prevents other classes from borrowing so much of the excess bandwidth that the short-term needs of this class may not be met. If it is not isolated, type **y**.
Do you wish to add a time-based rate [n]	A time-based rate sets a different Mbps for a specified time frame that you choose. To allow a time-based rate, type **y**.
Starting time (HH:MM using 24 hour clock)[8:00]	This is the time you want the temporary rate limit to start. Type a start time using the 24 hour clock.
Ending time (HH:MM using 24 hour clock)[17:00]	This is the time you want the temporary rate limit to end. Type an end time using the 24 hour clock.
Temporary rate limit (Mbps)	Type the rate limit you want to use during the time you specified.

Table 6.5 *Optional variables for custom rate limit classes*

To modify a rate limit class

You can modify any of the rate limit variables using the following command:

```
ratelimit custom modify
```

The interactive mode asks you for: class name, the variable to modify, and then it displays the variable and asks for a new value. If you are unsure of the variable name or value, type:

```
show ratelimit class <class name>
```

You can use this command change the priority of a class.

```
ratelimit custom modify prior <number 1 – 8>
```

To promote or demote a class

You use the **promote** or **demote** commands to move a particular rate limit class up or down in the list of classes. Ranking the classes allows you to choose in what order traffic hits the filters, thereby controlling how the controller filters that traffic. This is important when a particular set of traffic may match multiple rules, such as an HTTP rule and a rule for traffic to a particular client. To see the order that the classes are in, type:

```
show ratelimit classes
```

If you decide you want to move a rate limit class so that it filters packets before another rate limit class, you use the following command:

```
ratelimit custom [promote | demote] <class name>
```

This command promotes (or demotes) a rate limit class by increments of one. If you want to move a class more than one spot up or down, you need to run this command again.

Viewing rate limit statistics

You use the **show ratelimit** command to view statistics for rate limits. You can view the statistics for aggregate rate limits by typing:

```
show ratelimit stats aggregate <interface>
```

You can view the statistics for a particular rate limit class by typing:

```
show ratelimit stats <class name>
```

Using the batch mode to create custom rate limits

If you choose to create custom rate limits in the batch mode (starting at the system prompt and starting each line with **gsconfig**) you need the short-hand names for the variables. See Table 6.6 for the shorthand variable name and its definition. Following is the usage of the commands for batch mode:

```
gsconfig ratelimit custom add <class> <interface> <limit in Mbsp> <non-optional
    variable(s)> <optional variables>
```

Shorthand variable	value
dest	destination IP address
src	source IP address
dport	destination port
sport	source port
priority	priority level (range of 1 through 8)
bounded	0 (not bounded) or 1 (bounded)
isolated	0 (not isolated) or 1 (isolated)
hours	HH:MM-HH:MM@rate

Table 6.6 *Shortcut variable syntax for custom rate limit classes*

For example, to set a time-based rate limit for HTTP response traffic (port 8080), that is 100 Mbps by default, but 10 Mbps from 8:00 a.m. to 5:00 p.m., type:

```
gsconfig ratelimit custom add http_filter eth0 10 sport=8080 hours=8:00-17:00@100
```

reboot

```
reboot
```

The **reboot** command reboots the GLOBAL-SITE Controller immediately once you press the Enter key. You are not given any confirmation notice.

service

```
show <service name>
service alertd [enable | disable | start | stop]
service big3d [enable | disable | start | stop]
service corba [disable | dualmode | enable | ip | passphrase | start | stop]
service ITCMgswatch [enable | disable | start | stop]
service ntpd [enable | disable | start | stop]
service snmpd [enable | disable | start | stop | traps]
service sshd [enable | disable | start | stop]
service telnet [enable | disable]
service ftpd [enable | disable]
```

Use the **service** command to configure services on the GLOBAL-SITE Controller.

[enable | disable | start | stop]

You can enable and disable all services using the **service enable | disable** command. For FTP and Telnet, enable and disable also start and stop the service. You can also start and stop all other services with the **service start | stop** command.

- **start**
  Start the service now. On reboot, the GLOBAL-SITE Controller restarts the service only if it is enabled.
- **stop**
  Stop the service now. On reboot, the GLOBAL-SITE Controller restarts the service only if it is enabled.
- **enable**
  The service does not start now (except for FTP and Telnet). On reboot, the GLOBAL-SITE Controller starts the service.
- **disable**
  The service does not stop now (except for FTP and Telnet). On reboot, the GLOBAL-SITE Controller does not start the service.

Table 6.7 lists the services that are configurable with the **service** command and what they are for.

Service	Description
alertd	The alert daemon watches for specified events and takes any appropriate actions associated with alerts when those events occur. If you stop **alertd**, no actions (email, exec, or traps) trigger for events configured in **alert.conf**.
big3d	The **big3d** daemon answers 3-DNS system queries. The 3-DNS uses **big3d** to collect information about the network path between the BIG-IP and the client requesting a connection. The daemon collects performance data, and returns the data to the requesting 3-DNS. The 3-DNS uses the path information for its own dynamic load balancing.
corba	The CORBA daemon provides the iControl portal service. See the section *Configuring the CORBA service*, on page 6-29 for more information.
ITCMgswatch	**ITCMgswatch** is a system event monitor. It detects events such as excessive CPU usage, excessive memory usage, and low configuration and data disk capacity. It uses thresholds and timeouts to determine when alerts will be sent for each of the items.
ntpd	**ntpd** is used to configure the NTP daemon. The NTPd automatically updates the system time from and NTP (Network Time Protocol) server.
snmpd	With this command, you configure the SNMPd and the SNMP trap community, destination, and port. See *SNMPd*, on page 7-1 for more information on SNMP.
sshd	When running, the SSH daemon provides secure remote access to your controller.
telnet	Telnet allows non-secure access to your controller. You enable or disable Telnet with the **service telnet** command.
ftpd	When enabled **ftpd** permits FTP access to the controller.

Table 6.7 *Configurable services*

Configuring the CORBA service

The CORBA daemon provides the iControl portal service on a GLOBAL-SITE Controller. It is enabled by default so that the GLOBAL-SITE Controller can be an application-aware appliance on the CDN. If you stop the CORBA daemon, other products or iControl client programs cannot communicate with the GLOBAL-SITE Controller. Table 6.8 lists the available **service corba** commands and their use.

Command	Description
service corba disable	Use the **service corba disable** command so that the CORBA daemon does not restart on reboot. This command does not stop the daemon if it is currently running. Use **service corba stop** to stop the daemon.
service corba dualmode	The CORBA daemon can run in dual mode, which means that it allows iControl client programs to make connections on either a secure port (Secure Socket Layer [SSL]) or a non-secure port (Internet Inter-ORB Protocol [IIOP]). The command **service corba dualmode [on | off]** turns dual mode on or off. If turned off, the iControl client can only connect to the iControl service through the secure port channel.
service corba enable	Use the **service corba enable** command so that the CORBA daemon does restart on reboot. This command does not start the daemon if it is currently running. Use **service corba start** to start the daemon.
service corba ip	Set the IP address on which the CORBA daemon listens for client traffic with the **service corba ip** command. Multiple IP addresses are comma delimited.
service corba passphrase	Set the CORBA SSL certificate passphrase.
service corba start	Start the CORBA daemon.
service corba stop	Stop the CORBA daemon.

Table 6.8 *CORBA commands*

Displaying service status

Use the **show** command to see the status of the various services.

- **show alertd [config | status]**
- **show big3d [config | status]**
- **show corba [config | dualmode | ip | port | status]**
- **show ftpd**
- **show ITCMgswatch [config | status]**
- **show ntpd [status | config | server]**
- **show snmpd [config | status | trap]**
- **show sshd [config | status]**
- **show telnet**

show

```
show ?
show <command>
```

The **show** command displays different kinds of information depending on what command it is paired with. You can see configuration information, status, and statistics. Using the **show ?** command lists the subcommands for the `show` command.

Using the show command

Use the **show** command to display pertinent information about a system in the GLOBAL-SITE Controller. The **show** command syntax is displayed in the syntax summary box at the start of each section of this chapter with four exceptions: **acl**, **license**, **reboot**, and **stats**.

The **show** command can also be paired with commands that do not appear as first level commands. For example, to configure FTPd you type **service ftpd enable**. However, to see if FTP is enabled, you type: **show ftpd**.

Viewing the software version

The **show version** command displays the version of the GLOBAL-SITE Controller you are running. It displays:

- Product name and version number.
- Build number
- Crypto Edition, which indicates that the unit is encrypted
- The machine ID number

timezone

```
show timezone
timezone config
timezone options
```

The **timezone** command lets you set the timezone for the controller.

Setting the time zone

If you do not know the specific name of a time zone that you want to set, use the **timezone config** command. This command steps you through setting the time zone and you set it by selecting a country and city (for example America, Los Angeles). To get a list of actual time zones, use the **timezone options** command, which displays a complete list of time zones.

To set the time zone using country and city

1. At the ITCMconsole prompt, type **timezone config**.
   A display of countries with pre-assigned numbers appears.
2. Type the assigned number by the country you want and press Enter.
   A display of cities with pre-assigned numbers appears.
3. Type the assigned number by the city you want and press Enter.
4. Type **Y** to confirm the time zone you selected, or type **N** if the time zone is not correct. Accepting the time zone here sets the new time zone for the controller. Rejecting the time zone (typing **N**) returns you to the ITCMconsole prompt.

To set the time zone using a time zone name

1. At the ITCMconsole prompt, type **timezone options**.
   A list of all available named time zones appears.
2. Find the time zone you want to use and note it so that you can type it exactly as it appears.
3. Once you find the time zone you want to use, you can either use Ctrl + X to escape to the ITCMconsole command prompt, or you can continue paging through the list until you come to the prompt.
4. At the ITCMconsole prompt, type, **timezone config <timezone>**.
   For example to set the time zone for Los Angeles, California, you type:

   ```
   timezone config US/Pacific
   ```

user

```
show user [root | support | logadmin | gsite]
user root [password]
user support [enable | disable | password]
user logadmin [enable | disable | password]
user gsite [enable | disable | password]
```

The **user** command lets you set passwords for the system accounts: **root**, **support**, **logadmin**, and **gsite**. You can also enable or disable any of those accounts except for root. The root account is always enabled. System user accounts are not the same as webuser accounts (see *webuser*, on page 6-36). System user accounts have access to the internal configurations of the controller and the command line.

Configuring user accounts

WARNING

If you use regular command syntax to specify a password, the GLOBAL-SITE Controller command line displays the password on-screen, in clear text, and logs the password in the operating system history file. If you interactively specify a password, it does not appear on-screen or in the history file.

The **root** account is for managing command line access.

The **support** account manages technical support access.

The **logadmin** account configures the FTP log file access user account.

The **gsite** user is the system user on the GLOBAL-SITE Controller.

webserver

```
show webserver [allow | config | port | status]
webserver allow <IP address(es)> | all
webserver cert (same as FTBU)
webserver disable
webserver enable
webserver gencert
webserver port
webserver start
webserver stop
```

The **webserver** command configures the Apache webserver for the web-based user interface.

Configuring the webserver

You can configure your webserver using the **webserver** commands on Table 6.9, on page 6-35.

command	description
webserver allow	You can limit who has access to the web-based user interface by only allowing certain IP addresses access. Use the **webserver allow <ip address(es)> l all** command to restrict access to the user interface. You can type a space delimited string of IP addresses or domains, an allowable subnet (for example, 192.168.0.0), or **all** to allow all users access. For example: `webuser allow 192.168.0.0`
webserver cert	Generate an authentication certificate using the **webserver cert** command. This is the same certificate that you generate during the First-Time Boot utility and it calls for the same information: port, country code, state, locality, organization, name, email address. webserver cert <port> <country code> <state> <locality> <organization> <name> <email> **<port>** The port for the management web server. You can type, **none** to not use port in the certificate. **<country code>** The two-letter country code **<state>** Your state or province **<locality>** Your locality or city **<organization>** Your organization or department name **<name>** The network common name for the web server **<email>** The email address of the web server administrator. This must be a valid email address. You can choose to type **none**.
webserver disable	This command disables the webserver, which means that on reboot, the webserver does not automatically start. Disabling the webserver does not stop it.
webserver enable	This command enables the webserver, which means that on reboot, the webserver does automatically start. Enabling the webserver does not start it.
webserver gencert	The **webserver gencert** command re-generates an authentication certificate using the authentication information that was specified either using the **webserver cert** command or through the First-Time Boot utility.
webserver port	This command configures the webserver port. The default is 443. If you change the webserver port, you need to use the new port number when you access the web-based configuration utility.
webserver start	**webserver start** immediately starts the webserver.
webserver stop	**webserver stop** immediately stops the webserver.

Table 6.9 *The* ***webserver*** *commands*

webuser

```
show webusers
webuser add icontrol <username> <password> <password>
webuser delete icontrol <username>
webuser migrate
webuser icontrol [enable | disable | admin | regular]
```

The **webuser** command creates iControl interface accounts. (Please see iControl documentation for more information on iControl and the iControl SDK.) To create web user accounts for the GLOBAL-SITE Controller, please use the Add User page on the web-based user interface. See the online help for that page for more information.

Configuring accounts

You create a user account using the **webuser** command in that you create one user/password pair for the Apache web server (to access the GLOBAL-SITE Controller) and for the iControl interface.

To add a user with access to the iControl interface

1. From the ITCMconsole, type:

   ```
   webuser add icontrol
   ```

2. At the **Username** prompt, type the user name you want to assign.
3. At the **Password** prompt, type the password for the user account. Passwords must be between 6 and 15 characters long. The interactive mode masks the password on the screen.
4. Confirm the password at the **Confirm Password** prompt. When you complete this process you receive confirmation.

Migrating htpasswd contents to the LDAP server

WARNING

*The **webuser migrate** command runs automatically when you upgrade from an earlier version of the GLOBAL-SITE Controller. You should never need to run it again.*

The **webuser migrate** command migrates the contents of **.htpasswd** to the LDAP server. You need to use this command if you upgraded from an earlier version of the GLOBAL-SITE Controller that did not use LDAP authentication.

Configuring LDAP users

Use the **webuser icontrol** command to configure LDAP users. Once you create a user, you can enable or disable the user as an LDAP user, give the user LDAP administrative privileges, or return a user to regular LDAP privileges. Following is a procedure for giving LDAP administrative privileges, but you can use it for any of the **webuser icontrol** commands.

To give a user administrative privileges

Prior to giving a user account administrative privileges, the user account must already exist in the database.

1. From the ITCMconsole, type:

   ```
   webuser icontrol admin
   ```

2. At the **User to give LDAP Administrator privileges** prompt, type the user name.
   The system confirms that the user now has administrative privileges.

wrapper

```
show wrapper [in.telenetd | big3d | slapd | sshd]
wrapper allow <service> <string>
wrapper allow_append <service> <string>
```

The **wrapper** command sets the permissions for access to services that are listed in the **etc/hosts.allow** file. By default those services are: **in.telenetd**, **big3d**, **slapd**, and **sshd**.

Setting permissions

You can set access permission for in.telenetd, big3d, slapd, and sshd with the **wrapper** command. Once you have set permissions, if you use the **wrapper allow** command, you overwrite existing permissions. Instead, use **wrapper allow_append** to append new permissions to the original string in the **etc/hosts.allow** file.

To set access permission

To set permission for specific IP addresses, type:

```
wrapper allow slapd 192.168.0.0
```

To set permission for all users, type:

```
wrapper allow slapd all
```

7

SNMP, MIBs, traps, and alerts

- Working with SNMP and alerts
- SNMP agent
- Default and user-defined alerts

Working with SNMP and alerts

Alerts push system information to administrators by sending messages to a user-defined destination IP address or through email. Alerts can let administrators know the status of the machines so that the administrator does not have to monitor all the systems on the machine. Alerts can also be configured to automatically start scripts or executables based on system events so that the administrator does not have to.

Default alerts work with an SNMP agent (snmpd). The GLOBAL-SITE Controller ships with a customized simple network management protocol (SNMP) agent and management information base (MIB). The MIB contains the GLOBAL-SITE Controller's default alerts. The SNMP agent listens for events and triggers alerts when a preconfigured event occurs.

User-defined alerts that can be manually triggered or can be triggered using a Perl script.

This chapter describes the configuration tasks with which you can configure the GLOBAL-SITE Controller SNMP agent and how to customize default and user-defined alerts.

WARNING

We assume that you have prior knowledge and experience with SNMP to understand some of the concepts and terminology from this chapter.

You configure the SNMP agent and alerts through the ITCMconsole. If you are unfamiliar with the ITCMconsole, please see Chapter 6, *ITCMconsole Command Line Interface* for more information on the ITCMconsole and a complete listing of SNMP and alert commands and subcommands.

Note

You configure alerts and the snmpd using the ITCMconsole. While some of the pertinent commands are listed in this chapter, you can find all the commands and subcommands in ***ITCMconsole Command Line Interface****, on page 6-1 of this guide.*

SNMP agent

By default, the GLOBAL-SITE Controller runs the SNMP agent (snmpd). You can enable and disable, start and stop the agent, and set the SNMP trap values from the ITCMconsole (see *Using the ITCMconsole*, on page 6-1). If you are going to use default alerts, you have to add values to the SNMP trap:

- **community**
  In this case, **public**

- **dest**
  The destination IP address that the snmpd sends the trap to when it triggers
- **port**
  The port that the default alert uses to communicate with the destination IP address

To configure snmpd

The following commands configure the trap values. At the ITCMconsole prompt, type:

```
service snmpd trap community public
service snmpd trap dest <ip address>
service snmpd trap port <port>
```

MIBs

The GLOBAL-SITE Controller uses several open source MIBs as well as the GLOBAL-SITE MIB, which is a product-specific MIB that polls pertinent data from your GLOBAL-SITE Controller. You can find a text file of the MIB on the Documentation section of the documentation CD, or from the GLOBAL-SITE MIB link on the GLOBAL-SITE Controller main web-based user interface page.

Default and user-defined alerts

There are two kinds of alerts on the GLOBAL-SITE Controller: default alerts and user-defined alerts. Since the snmpd runs by default on your GLOBAL-SITE Controller, the default alerts are already enabled. User-defined alerts are alerts you create to meet your needs. To see all alerts, from the ITCMconsole prompt, type:

```
show alert all
```

Figure 7.1 is a sample of the output you receive when you type **show alert all** from the ITCMconsole prompt.

Using global settings

The alert system was designed with global settings to make modifying default alerts and creating and modifying user-defined alerts faster and easier. For more information on global settings for alerts, please see *Changing global settings*, on page 6-10 of this guide.

```
/*Alerts:*/
   COMMON.trapDiskDriveFailure
        /*OID*/ ".1.3.6.1.4.1.3375.1.5.110.2.1"
        /*DESCRIPTION*/ "A disk drive has failed."
   COMMON.trapFreeVirtualMem
        /*OID*/ ".1.3.6.1.4.1.3375.1.5.110.2.2"
        /*DESCRIPTION*/ "There is not enough memory to complete an operation."
   COMMON.trapExcessiveCpuUsage
        /*OID*/ ".1.3.6.1.4.1.3375.1.5.110.2.3"
        /*DESCRIPTION*/ "The CPU usage exceeds its threshold."
   GLOBALSITE.globalsiteTrapDataDiskFull
        /*OID*/ ".1.3.6.1.4.1.3375.1.4.110.2.1"
        /*DESCRIPTION*/ "The data disk is full."
   GLOBALSITE.globalsiteTrapConfigurationDiskFull
        /*OID*/ ".1.3.6.1.4.1.3375.1.4.110.2.2"
        /*DESCRIPTION*/ "The configuration disk is full."
   GLOBALSITE.globalsiteTrapPublicationSucceeded
        /*OID*/ ".1.3.6.1.4.1.3375.1.4.110.2.3"
        /*DESCRIPTION*/ "Publication succeeded."
   GLOBALSITE.globalsiteTrapPublicationFailed
        /*OID*/ ".1.3.6.1.4.1.3375.1.4.110.2.4"
        /*DESCRIPTION*/ "Publication failed."
   GLOBALSITE.globalsiteTrapPublicationCheckedForNew
        /*OID*/ ".1.3.6.1.4.1.3375.1.4.110.2.5"
        /*DESCRIPTION*/ "There is no new version of this publication."
```

Figure 7.1 *Default alerts*

Default alerts

Default alerts consist of a name and a description. Some default alerts also have an OID (object identifier). Default alerts that have an OID are part of the GLOBAL-SITE MIB and are triggered by events. There are eight default alerts that the following events trigger:

- Disk drive failure
- Insufficient memory to complete an operation
- CPU usage exceeding its threshold
- Full data disk
- Full configuration disk
- Successful publication
- Failed publication
- No new version of a publication

Modifying default alerts

You can add actions to or delete user-defined actions from default alerts, but you cannot delete default alerts, change the name of the default alert, or change the OID. Actions can be either **named_action**, **email**, or **exec**. For example, upon successful publication of a site, the snmpd sends the

GLOBALSITE.globalsiteTrapPublicationSucceeded default alert "Publication succeeded" to the destination IP address. You can modify the default alert and add an action to send email to additional people informing them of the successful publication, as well.

To modify a default alert

To modify an alert, you first need to know the name of that alert. Default alert names are fairly long and must be typed accurately. You can also find the names of the default alerts in the GLOBAL-SITE MIB. (You can access this from the GLOBAL-SITE Controller home page or from the Documentation section of the product splash screen.)

◆ Tip

If you are accessing the GLOBAL-SITE Controller from an SSH application, and this application allows you to copy and paste, we recommend you do so.

1. To get the name of the default alert from the ITCMconsole, type:
   ```
   show alert all
   ```
2. If possible, copy the name of the default alert you want to modify.
3. From the ITCMconsole, type:
   ```
   alert modify add_action
   ```
4. At the prompt for **Alert name**, paste or type the name of the default alert.
5. Choose to add an action.
   - A **named_action** is an user-defined action that you created previously. (Please see *named_action*, on page 6-7 for information on creating a named action.) At the prompt, type the name of the named action.
   - An **email** action allows you to send an email (please see *email*, on page 6-6 for more information). At the prompt, type **email**.
   - An **exec** action allows you to run scripts or executables at the time the alert is sent (please see *exec*, on page 6-6). At the prompt, type **exec**.

Creating user-defined alerts

From the ITCMconsole, you can create your own user-defined alerts using the **alert** commands.

The **show alert all** command displays both default alerts and user-defined alerts. If you modify an alert, you need its name which you get with the **show** command.

To create or modify a user-defined alert

1. From the ITCMconsole, type:

   ```
   alert add
   ```

2. At the prompt for **Alert name**, type a name for the alert.
3. Choose to add either a named action, email, or exec as your action.
   - A **named_action** is an user-defined action that you created previously. (Please see *exec*, on page 6-6 for information on creating a named action.) At the prompt, type the name of the named action you created.
   - An **email** action allows you to send an email. At the prompt, type **email** (see *email*, on page 6-6 for more information). The steps that follow ask you for the to address, from address, body text, and if you want to add an attachment, the attachment name.
   - An **exec** action allows you to run scripts or executables at the time the alert is sent (see *exec*, on page 6-6 for more information). At the prompt, type **exec**. The step that follows asks you for a command to execute when the alert is raised and for a description of the alert.

Triggering user-defined alerts

The **alert raise <alert name> <args>** command triggers alerts. If you type the command at the ITCMconsole command prompt, the alert triggers right away. You can also put the **alert raise** command in a bash or a Perl script and have it run that way.

A

Reconfiguring the Controller

Changing the host name 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:

- Change the host name or IP address using the ITCMconsole command **netinet**. See *netinet*, on page 6-18 for more information.
- 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. You should only do this using the **ITCMconsole netinet** command. Please see *netinet*, on page 6-18 for more information.

Resetting the GLOBAL-SITE identifier

You only need to reset the GLOBAL-SITE identifier if, for some reason, the **netinet** command fails to do so. Follow these instructions in cases where the network name changed without resetting the GLOBAL-SITE identifier.

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.
   You should see a line that shows the installation ID as an integer, like this:

   ```
   Installation ID:  74565
   ```

   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 is 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 can 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 that has just been 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 publication

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.

authentication

Authentication is a method of confirming the identity of a user or client attempting to gain access to your system. There are different methods of authentication, including the confirmation of passwords, certificates, or information stored on an ID card.

cache subscriber

A cache subscriber is the GLOBAL-SITE Controller's way of recognizing a cache and the URLs it caches content for. The GLOBAL-SITE Controller sends information to the cache subscriber regarding the availability of new content on an origin server. Usually based on a server or virtual server subscriber, a cache subscriber is represented by the host name or IP address of an EDGE-FX Cache in your network.

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.

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.

distributor

A distributor is a secondary GLOBAL-SITE Controller that can be used to deliver publications that it receives from the primary controller. A distributor may be located in a remote data center. The distributor can provide the primary controller access to additional subscribers that are geographically closer to the distributor.
Compare to *publisher*.

edition

In archived publications, an edition is a complete collection of specific versions of selected sections. An edition is an instance of a publication that indicates two things: which versions of selected sections are to be included, and 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 only exclude directories, not files.
See also *file filter*.

expire

The GLOBAL-SITE Controller tells the EDGE-FX Cache to expire content. To the cache this means that the content is stale and it retrieves the new content and forwards it to the end user.

file filter

This GLOBAL-SITE Controller feature either includes or excludes all files in a path that have a specified file extension, such as gif, jpg, or html. File extensions are more consistent in Windows, where they often identify specific file types, but this feature can work equally well in UNIX systems, where naming conventions allow you to group related files by extension. For example, html is the extension in this file authentication, basic and

digest.

Name: home.banner.html

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 the GLOBAL-SITE Controller uses to collect and publish file-based content to Internet sites.

FTP-Push

FTP-Push is a file transfer method used by the GLOBAL-SITE Controller that is based on basic FTP. FTP-Push sections can push files from the content source (FTP client) to the controller (FTP server), where they are stored. If a section uses FTP-Push as its transfer method, you can configure the update of content to be used to trigger the delivery of that publication.

GLOBAL-SITE agent

The GLOBAL-SITE Controller agent allows a GLOBAL-SITE Controller to start and stop web services, reboot the server, and register components, as well as transfer content files.

GLOBAL-SITE identifier

The unique, pre-assigned numeric identifier provided by your vendor 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 publication

A 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 customers can access it. However, non-archived publications create and keep version lists of dated files so that the GLOBAL-SITE Controller can detect changed files. Compare to *archived publication*.

origin server

Used in a cache context, the web server referenced by a cache, on which all copies of your content reside and which your cache is caching content for.

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 does not load balance the connections; instead, it returns the clients to the node that they previously connected to.

populate

The GLOBAL-SITE Controller tells the EDGE-FX Cache to populate content. To the cache, this has the same effect as a user entering a URL in the address box of the browser: the cache goes out to its origin server and gets the new content. Therefore, the cache always has the most current content before an end user requests it. Compare to *expire*.

private key

One of two keys used in asymmetric cryptography. The private key can be used to encrypt data that 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 that 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 virtual servers, how it handles errors, and error email notification.

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*.

RAID (Redundant Array of Independent Drives)

RAID allows you to store information in multiple hard drive locations, providing the safety of redundancy or the performance of writing data simultaneously to multiple locations. In the case of redundant disks, if one disk fails, another can immediately take over, reducing the total failure rates and providing fault-tolerance. Performance improvements result from using simultaneous operating system processes to write to multiple disks at the same time.

repeater

See *distributor*.

repository

See *section*.

section

The source content retrieved by 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.

signed certificate

A signed certificate verifies a person's identity online to another user or application. Once that person's identity is verified with the other user or application, those parties can do a private key exchange to establish an encrypted session.

subscriber

The location reached by 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*.

WebDAV (Web Distributed Authoring and Versioning)

WebDAV (and WebDAV-SSL) is a file transfer method used by the GLOBAL-SITE Controller. It is an Internet Engineering Task Force (IETF) standard for collaborative authoring on the Web. WebDAV uses a set of extensions to the HTTP protocol. Therefore, to use WebDAV as your transfer method, you need to have a web server that is compatible with WebDAV (IIS5.x or Apache HTTP Server with mod_dav) or you need to install the GLOBAL-SITE agent on the remote server. WebDAV also adds write access to the read access provided by HTTP. See also *GLOBAL-SITE agent*.

Index

D

E

W