

Internet Monitoring Software: Technical Considerations



Abstract

This paper is written primarily for IT personnel who are searching for or evaluating Internet monitoring software. The information provided is intended to help them select a well-designed product that will meet IT's technical requirements as well as the functional requirements of their internal customers.

Section One of this paper provides a brief introduction to employee Internet monitoring software and the general issues surrounding its implementation in the workplace.

Section Two focuses in detail on the technical requirements of an effective Internet monitoring solution. Among other topics, it examines issues of compatibility, interoperability, scalability, security and ease of use. Each topic includes a real-world example of how one product – Wavecrest Computing's Cyfin Reporter – addresses these issues effectively.

Section Three first outlines the functional (informational) requirements of IT's internal audiences, e.g., HR representatives, business managers and IT security personnel. It then briefly describes Cyfin's approaches to those requirements.

Appendix "A" is a tabularized summary of Internet monitoring software requirements and Cyfin's approaches to them.

Internet monitoring software provides information that helps organizations manage the use of their network resources. IT, HR and business managers can use the data to maximize workers' online productivity and minimize misuse of network resources.

Section 1 | Introduction

Numerous organizations use software to monitor their workers' use of Internet and intranet resources. If the software is well-designed, it can help maximize productive use of these resources while minimizing misuse or abuse, and it can do so with no negative impact on network resources or IT personnel.

For optimum results, the software should be viewed and used as one part of a comprehensive network usage management program. Such a program will have a well-designed Acceptable Use Policy (AUP) at its foundation. Enforcing that policy – a function that is critical to the overall program – depends to a large extent on the ability of the software to produce accurate, reliable, comprehensive reports of the workforce's Web activity – reports that can clearly depict users' online behavior and highlight potential problem areas. If the software can do this, it can enable responsible managerial, HR and IT professionals to take corrective actions when necessary.

As well as meeting these functional requirements effectively, well-designed Internet monitoring software needs to work well from a technical perspective – in a variety of network environments and infrastructures. It also needs to be easy to install, integrate, set up and administer.

Wavecrest Computing prides itself on the ability of its Cyfin Reporter product to meet both types of requirements, i.e., functional and technical. Cyfin's approaches to both are discussed in Sections 2 and 3 below.

www.wavecrest.net



Section 2 | Technical Requirements and Wavecrest Approaches

General. Typically, IT personnel are responsible for finding, testing, selecting, installing and administering Internet monitoring software. And in many cases, they are also required to generate and analyze Web-use reports.

These responsibilities give rise to a number of technical questions and requirements. Naturally, the specific nature of these considerations will vary from organization to organization. In this section of the paper we'll examine these issues and see how Wavecrest Computing's Cyfin Reporter product goes about addressing them. Before doing so though, let's take a quick overall look at Cyfin and some of its key attributes.

Functional Considerations. Using logfiles as its raw data, Cyfin monitors and reports on computer users' Web visit activity. From bandwidth usage to detailed user audits, Cyfin's sixteen standard reports include 104 different data views – more than any other Web-use management software. Cyfin's reports provide IT, HR and business managers with the exact information they need to manage employee Internet activity effectively.

Cost Considerations. Cyfin is a best-value, low Total Cost of Ownership (TCO) product. While producing accurate, manager-ready results at a very affordable price, it requires minimal administration and maintenance, and all follow-on support services are provided at no extra charge. Licenses are subscription-based, with pricing based on the number of users to be monitored.

Technical Considerations. Cyfin is very IT-friendly. That is, it works well in – and is totally compatible with – a wide variety of network environments, including those that are heterogeneous and have multiple sources of logfiles. And with its browser interface, it's very easy to administer.

With this background in mind, let's explore six of the more critical technical areas related to Internet monitoring software. We'll first look at the requirements in these areas and then examine Cyfin's approach to them. (Note: Requirements are expressed in generic rather than specific terms so they can apply to a wide variety of organizations.)

2.1 Compatibility

2.1.1 Requirements. Is the software compatible with existing and/or planned network resources? How easily will it integrate with those resources?

2.1.2 Wavecrest's Approach. Let's subdivide Wavecrest's approach into four elements:

A. Ability to Use Logfiles from Multiple Sources. Cyfin's application server uses logfiles as its source of raw data – processing them into a wide variety of accurate, actionable reports. Cyfin is compatible with and can obtain the logfiles from a large number and wide variety of proxy servers, firewalls, and caching appliances. To see a complete list, visit www.wavecrest.net/support/cyfin/specifications/compatibility/compatibility.html.

B. Integration. Cyfin integrates easily and smoothly into all types of networks. It's server-based, it doesn't require a third-party Web server, and it doesn't require installation on the monitored users' computers. Cyfin does not connect to or communicate with any proxy, firewall or appliance that it supports, i.e., the server simply reads and analyzes logfiles. And because Cyfin is not a plug-in, it has no direct impact on network performance.

Employee Internet management software needs to meet the technical requirements of IT and the functional requirements of HR, compliance and business unit managers. This paper addresses key technical issues. For more on functional issues, see "Resources" on www.wavecrest.net.



C. Installation Location. Cyfin is very flexible and can be installed several different ways. For example, in certain environments, it can be installed directly on the firewall or proxy server where the logfiles reside. This approach is illustrated in Figure 1a below, using a proxy server as an example.

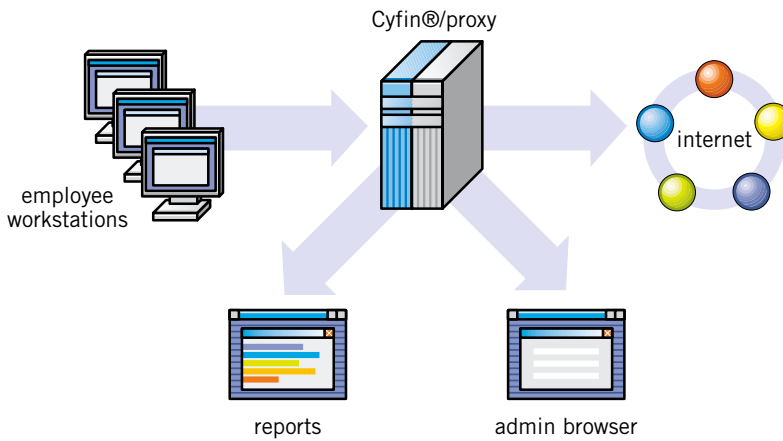


Figure 1a – Cyfin installation on a proxy server

As another example, illustrated in Figure 1b below, Cyfin can be installed on a separate general-purpose server. In this case, the logfiles would be residing on a firewall, proxy server or caching appliance. In this arrangement, the administrator can either (a) transfer the logfiles to Cyfin using FTP or syslog techniques or (b) access them via a network drive.

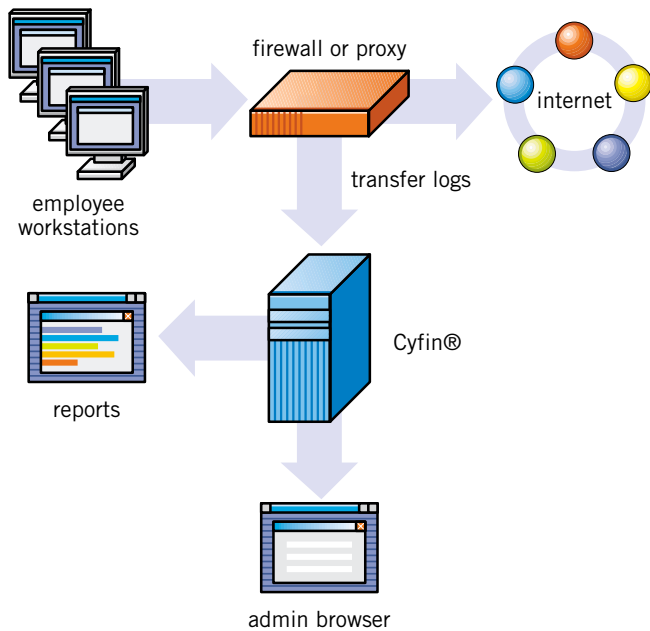


Figure 1b – Cyfin installation on a general-purpose server.

Using either of these approaches, the CPU's processing load will be transferred from the proxy/firewall server to the Cyfin machine.

Can the product integrate easily with your system? In Cyfin's case, the answer is yes. Cyfin Reporter is flexible and can be installed in several configurations as either an on-box or off-box solution.



D. OS Compatibility. Cyfin can be installed on a variety of operating systems, including Windows XP, Windows 2000, Linux, and Solaris.

2.2 Ease of Installation and Administration

2.2.1 Requirements. Is the software easy to install, and will it be easy and efficient to administer?

2.2.2 Wavecrest's Approach. Cyfin is very easy to install and administer, as discussed below.

A. Installation. Cyfin is server-based and requires no installation on the monitored users' computers. Downloading the software and installing it on a server is straightforward and only requires a few minutes.

B. Administration. Cyfin can be administered remotely or locally from any compatible browser (see paragraph 2.5). It requires very few settings for basic operation and only a few more for advanced functionality, as discussed in "C" and "D" below.

C. Required Settings. To function at a basic level, Cyfin only needs to be told where the logfiles are located. The administrator can then create and view reports using several default settings that govern the format of the reports.

D. Advanced, Optional Settings. Advanced settings enable the administrator to group users (IDs) for focused reporting, import user IDs from directories, rate categories for acceptability, establish thresholds to automatically detect abuse, and customize the format of reports. None of these settings is very difficult to configure.

2.3 Adaptability and Scalability

2.3.1 Requirements. Can the product adapt to changes in network architecture or resources? Similarly, is it scalable, i.e., will its performance hold up if more users are added?

2.3.2 Wavecrest's Approach. Cyfin is designed with customer growth and network changes in mind. As discussed in paragraph 2.1, Cyfin is compatible with a wide variety of proxy servers, firewalls, caching appliances and operating systems. This compatibility, plus Cyfin's numerous installation options and use of a browser interface, gives it almost unlimited flexibility to respond to changes in network configuration or IT resources. In addition, Cyfin's scalability, i.e., to 100,000 plus users, enables it to adapt automatically to increases in user population.

2.4 Accuracy of Results

2.4.1 Requirements. Accuracy of information in reports is absolutely critical and is affected by several key issues:

A. Does the software use a URL control list as the basis for categorizing visits, or does it use less accurate "sniffer" or keyword "analysis" techniques?

B. Can it make a clear distinction between true user visits and extraneous hits, e.g., ads, audio, etc.?

C. Can it measure the time that a user spends on a particular site?

Can the product adapt to changes in network architecture or resources?

Is it scalable? A solution like Cyfin allows for customer growth and network changes and is scalable to 100,000+ users.



2.4.2 Wavcrest's Approach. Let's examine Wavcrest's approaches to these three accuracy-related issues:

A. Categorization Technique: Cyfin uses a URL "control list" for categorization of hits and visits. This approach is considerably more accurate and reliable than the use of "sniffer" or "keyword" techniques.

B. Differentiating True Visits from Extraneous Hits: Cyfin makes a clear distinction between true user visits and the numerous extraneous hits, e.g., ads, audio, etc., that are downloaded in response to the visit request. Without such differentiation, users' levels of activity can be greatly inflated. This is an extremely important consideration and is discussed in a bit more detail in Section 3, Functional Requirements.

C. Measurement of Time Spent on Site: No software can measure the time a user spends viewing a Web site. From a technical perspective, when a computer user clicks on a URL, that action merely constitutes a one-time request for a document, and the "connection" with the server is immediately broken. There is no "session," no "end-of-session," and consequently no possibility of measuring duration. From a human behavior perspective, the fact that a Web page has been downloaded doesn't necessarily mean that the user was actually looking at it while it was open. That is, the user could have suddenly received a phone call, gone to lunch, attended a short-notice meeting, been interrupted by a colleague, or gone to the restroom. Also, users often have more than one page open at a time, but only one is truly viewable, and there is no way to know which one.

Some products claim to measure time spent on the Web. This is an appealing but unachievable and misleading assertion. And the use of such products can lead managers to take ill-considered and unjustifiable actions against users.

Although it's not possible to measure time on site, per se, it is possible to obtain an accurate estimate of the overall volume or level of activity engaged in by a user or group of users. Cyfin does this by counting actual visits (clicks) after they have been separated out from extraneous hits, as discussed above.

How accurate is the data provided by the product? Only Cyfin makes a clear distinction between true user visits and extraneous hits (ads, sound files, graphics), giving managers an accurate picture of the volume of user activity.

2.5 Interface Efficiency and Flexibility

2.5.1 Requirements. Does the software feature a Web-enabled browser interface that can be used by authorized non-IT personnel as well as IT administrators? Is the interface secure?

2.5.2 Wavcrest's Approach. Cyfin features a secure, easy-to-use browser interface that supports IE and Netscape versions 6.0 and higher. It can be used by non-IT as well as IT personnel, as discussed below.

A. Efficiency. Cyfin's browser interface can be used by non-IT personnel as well as network administrators. Consequently, if so desired, much of the day-to-day burden of report creation and review can be handled by authorized individuals in other departments, thus easing IT's workload. This approach is illustrated in Figure 2.




Figure 2 – Cyfin's multi-purpose browser interface.



B. Accounts. Non-IT personnel such as HR representatives and business unit managers can be granted “operator” accounts. Operator accounts enable them to create and access routine reports but don’t allow them to input or alter administration settings. In addition, operator accounts can be restricted to receiving reports on limited sets of employees, e.g., only those that are supervised by the account holder.

C. Security. Cyfin’s interface uses fully developed technology such as browser connectivity, HTML forms, CGI, SSL security, etc. Its thin client approach combines and configures several of these proven technologies and can be used to support a relatively simple – yet powerful and secure – decentralized Web-use management concept. Its security features are illustrated in Figure 3 below.

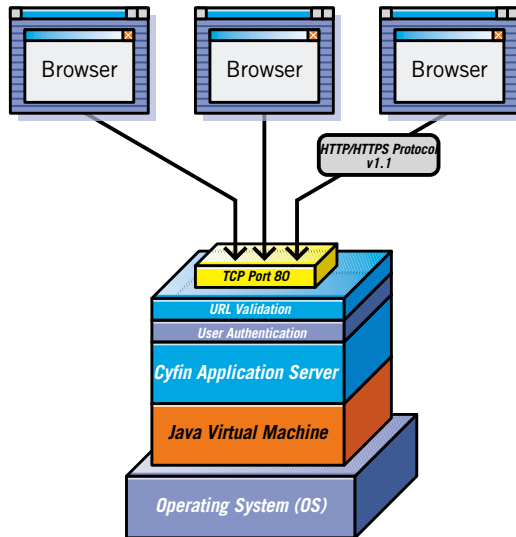


Figure 3 – Cyfin’s browser interface security features

2.6 Vendor Support

2.6.1 Requirements. What type of support does the vendor offer? How expensive will it be? How reliable is the vendor in terms of providing quality support over the long term? Does the vendor specialize in employee Web-use monitoring?

2.6.2 Wavecrest’s Approach. Wavecrest’s online and telephone support services are responsive and free. With a Cyfin license subscription, Wavecrest provides free software maintenance (including upgrades and patches), free technical support and free weekly updates to the built-in categorization list, called Wavecrest URL List. Wavecrest customer support personnel also work one-on-one with customers to enhance and improve reporting results on an ongoing basis. For more info on Wavecrest Computing’s support services, visit <http://www.wavecrest.net/support/>

What type of support does the vendor offer? Wavecrest provides online and telephone tech support as well as participation in our OtherWise list optimization service free of charge. Not all vendors do – find out early what services are included in your contract.

Section 3 | Functional Requirements and Wavcrest Approaches

3.1 General Requirements Internet monitoring software must meet the functional needs of IT's internal customers. Typically, these are HR representatives, business unit managers and IT security personnel. To perform their roles successfully, these customers need reliable information that, as a minimum:

- A. Shows very accurately which users visited which Web sites, when they did so, and how many times.
- B. Categorizes visits by the type of content viewed.
- C. Can indicate if the activity complies with the organization's Acceptable Use Policy (AUP).
- D. Is accessible, accurate, and actionable.
- E. Is clear and easy to use.
- F. Is available in summary-level as well as detailed-level reports.
- G. Can cover an individual user, a particular group of users, or an entire enterprise.

IT's internal customers – HR, legal, business managers, IT security or compliance officers – have a range of functional requirements for Web use management software. Do reports tie in to the company's Web use policy? Are reports categorized clearly? Are results accurate? Can managers access reports themselves?

Information Accuracy. By far, the most important consideration in the reports is the accuracy of information that depicts the level or amount of user activity by category. To prevent serious inflation of this measurement, it is crucial that the software make a clear distinction between actual *visits* (clickable events performed by human beings) and extraneous *hits* (unsolicited ads, audio, etc.). Figure 4 and its accompanying notes explain the need for this distinction.

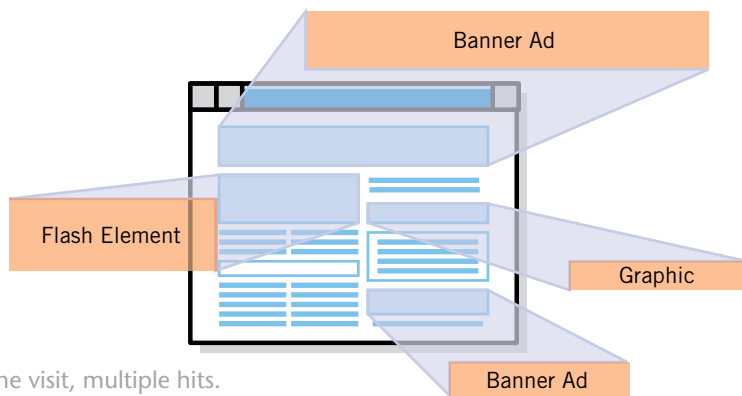


Figure 4 – One visit, multiple hits.

Figure 4 represents one requested Web page. It was downloaded and displayed with one click. This equates to one true visit. However, each highlighted element constitutes an additional hit and was downloaded separately from the text that accompanied the original page request. In the logfiles, this one Web page would be represented by five URLs, i.e., the original request and four other extraneous hits. Listed below are some typical extraneous hits:

- Graphics
- Banner ads
- Audio clips
- Flash elements
- Java Applets
- SSL encryptions

Unfortunately, many Internet monitoring tools lump extraneous hits with visit requests, seriously exaggerating the level of activity by a factor of three, four or five.



Time Measurements. This subject was discussed in Section 2 but it bears repeating here. Most organizations would like to have software that could measure and report on the exact amount of time that specific users spend on specific Web sites. Unfortunately, for technical as well as human behavioral reasons, this just isn't possible, i.e., logfiles only record the time that URLs are requested, and workers can leave Web pages open for long periods of time without actually looking at them. On the other hand, while actual time-on-site cannot be measured, it is possible to roughly gauge download time (i.e., page loading time) associated with each true visit. This can serve as a rough estimate of the absolute *minimum* amount of time that each visit consumed. Wavecrest uses this technique, as mentioned later.

3.2 How Cyfin Meets Functional Requirements.

3.2.1 General. Wavecrest's Cyfin Reporter product is designed to provide reliable, accurate Web-access reports that help organizations reduce misuse of network resources and control labor and legal costs. These same reports can also help preclude embarrassing and costly lawsuits and help demonstrate a proactive approach to Web-access control.

To meet these requirements and objectives, Cyfin's reports indicate which online workers visited which Web sites, when they did so, how many times they did so, the types of content they were seeking, and the "acceptability" of the content. This information can be used in a variety of ways to:

- quickly gauge compliance with the organization's Acceptable Use Policy.
- provide information with which to adjust policies, processes and training.
- accurately gauge the effectiveness of Internet, intranet and extranet activity.
- develop evidence to support employee termination or other disciplinary action.

3.2.2 Meeting Specific Functional Requirements. To meet the functional requirements discussed above, Cyfin Reporter:

- lists results in 67 content-labeled categories.
- enables customers to add custom categories if needed.
- lets customers rate categories for "acceptability."
- can be configured to detect abuse automatically.
- enables customers to schedule reports to run automatically.
- enables authorized users to generate reports from any computer with browser access.
- automatically downloads categorization list updates.
- enables customers to disable nonessential categories.
- lets customers import IDs from directories.
- lets customers exclude VIPs from reports if desired.
- delivers "exception-only" reports.
- can report on individual users, specific work groups, or entire enterprises.
- can produce 16 reports with 107 different data views.
- makes a clear distinction between actual user visits from extraneous hits.

Wavecrest's Cyfin Reporter product is designed to provide reliable, accurate Web-access reports that help organizations reduce misuse of network resources and control labor and legal costs.

Summary. Cyfin's summary-level and detailed reports provide accurate, reliable information that helps organizations maximize employee productivity while minimizing bandwidth costs and legal liability risks. HR and legal personnel use the reports to identify and deal with personnel management and legal liability issues. IT personnel use them to spot patterns of activity that have bandwidth implications. And business unit managers use them to deal with productivity issues.



I. Cyfin's Approach to Technical Requirements

CUSTOMER REQUIREMENTS

CYFIN'S APPROACHES

1. Compatibility

- A. Support multiple logfile sources simultaneously?
- B. Integrate into the network infrastructure with minimal effort and impact?
- C. Offer multiple installation options?
- D. Work with customer's Operating System?

1. Compatibility

- A. Works with numerous proxy servers, firewalls and caching appliances – simultaneously if necessary.
- B. Doesn't require a 3rd party Web server. Is not a plug-in. Only needs to read logfiles; does not impact performance.
- C. Can be installed on any available server. Logfiles can be FTP'd or accessed directly.
- D. Supports Windows XP, Windows 2000, Linux and Solaris.

2. Installation and Administration

- A. Install easily and smoothly?
- B. Easy to administer?

2. Installation and Administration

- A. Server-based. No client installation. Only requires a few minutes.
- B. Browser interface. Very few required settings. Advanced settings easy to configure.

3. Adaptability and Scalability

- A. Accommodate network equipment and software changes?
- B. Scalable for growth of user population?

3. Adaptability and Scalability

- A. Supports dynamically changing heterogeneous environments, e.g., multiple logfile sources and operating systems.
- B. Scalable to >100K users.

4. Accuracy of Results

- A. Categorize Web sites accurately?
- B. Distinguish true visits from extraneous hits?
- C. Measure amount of time user spends on Web site?

4. Accuracy of Results

- A. Uses URL database (list) approach. More accurate than sniffer and keyword techniques.
- B. Separates true visits (human activity) from the more numerous extraneous hits, e.g., banner ads, graphics, audio, etc.
- C. No software can do this. Cyfin counts clicks (visits) to accurately measure amount of activity.

5. Interface: Efficiency and Flexibility

- A. Usable by non-IT as well as IT personnel?
- B. Permit establishment of access accounts?
- C. Provide security protection?

5. Interface: Efficiency and Flexibility

- A. Usable by authorized non-IT personnel for report creation and access only.
- B. Permits establishment of Operator accounts for non-IT personnel and Administrator accounts for IT personnel.
- C. Provides multiple layers of security, e.g., SSL, Java Virtual Machine, URL validation, user authentication, and more.

6. Vendor Support

- A. Provide robust support services?
- B. Provide support services at no extra charge?

6. Vendor Support

- A. Services include tech support, upgrades, updates, usage tips, personalized reporting optimization, and more.
- B. All support services provided within the subscription fee.



II. Cyfin's Approach to Functional Requirements

CUSTOMER'S REQUIREMENTS

1. General

- A. Provide comprehensive reporting?
- B. Correlatable with customer's AUP?
- C. Permit users to be grouped by department or location?
- D. Usable by non-IT personnel to create and access reports?

2. Accuracy of Information

See Accuracy of Results under Technical Requirements.

CYFIN'S APPROACHES

1. General

- A. Provides 16 reports with 107 data views. Provides summarized high-level and detailed low-level reports.
- B. Lets customers:
 - (a) rate categories for acceptability and
 - (b) set thresholds to automatically detect abuse.
- C. To compartmentalize reports, customers can group users by organizational elements or geographic locations.
- D. Usable by authorized non-IT personnel for report creation and access.

2. Accuracy of Information

See Accuracy of Results under Technical Requirements.

III. Cyfin's Approach to Cost Requirements

CUSTOMER REQUIREMENTS

- A. Initial price affordable?
- B. Follow-on costs affordable?

CYFIN'S APPROACHES

- A. Priced lower than competitors while providing better, more accurate results.
- B. Support included in subscription price.

IV. Additional Resources and Contact Information

RESOURCES

- A. **Cyfin Reporter Evaluation.** Free, 30-day evaluation of Wavecrest's Cyfin Reporter, the Internet monitoring software reviewed in this paper.
- B. **How to Buy.** Wavecrest's sales team is available to answer questions regarding pricing or licensing.
- C. **Technical Questions.** Wavecrest's technical support is available to answer questions regarding system integration and installation.
- D. **White Papers.** Additional white papers and resources on Internet monitoring issues.

CONTACT INFORMATION

- A. Register for a Cyfin evaluation at:
www.wavecrest.net/products/cyfin/freetrial.html
- B. Email: sales@wavecrest.net
Phone: 321/953.5351
Toll-free (USA): 877/442.9346
- C. Online: www.wavecrest.net/support/
Email: support@wavecrest.net
- D. Visit
www.wavecrest.net/editorial/resources.html

